

Horizons and soft quantum information

Daine L. Danielson^{1,2,3,*} and Gautam Satishchandran^{4,†}

¹*Center for Theoretical Physics—a Leinweber Institute, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

²*Black Hole Initiative and Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*

³*Leinweber Institute for Theoretical Physics, Enrico Fermi Institute, and Department of Physics, The University of Chicago, Chicago, Illinois 60637, USA*

⁴*Princeton Gravity Initiative, Princeton University, Princeton, New Jersey 08544, USA*



(Received 26 February 2026; accepted 17 June 2026; published 6 October 2026)

It was recently shown that black holes decohere any quantum superpositions in their vicinity. This decoherence is mediated by soft radiation through the horizon, and can be understood as the result of the fact that quantum states in the exterior source distinguishable states of long-range fields in the interior. To study this phenomenon and others, we extend Tomita-Takesaki theory to accommodate states of soft radiation such as arise in the electromagnetic and gravitational memory effects, and provide a general framework for computing the distinguishability of general coherent states. Applying these tools, we use the methods of unambiguous state discrimination and approximate quantum error correction to prove some new relations regarding the distinguishability of quantum states, and the quantum information content of soft radiation, and thereby show that a black hole (or, under the conditions stated below, any causal horizon) decoheres its environment as though its interior were full of optimal observers.

DOI: [10.1103/mk9b-3qzz](https://doi.org/10.1103/mk9b-3qzz)

I. INTRODUCTION

The fact that black holes have a finite temperature [1] and, thereby, a nonvanishing entropy [2] has provided some of the deepest clues into the nature of these objects in a quantum theory of gravity. In the case of black holes, this entropy has been argued to arise from “internal degrees of freedom” (described by microstates) and, for certain quantum gravitational realizations of black holes, the counting of microstates can directly reproduce their entropies [3]. Despite these successes, the origin and nature of these microstates remain poorly understood. Indeed, classically, a black hole has no internal degrees of freedom. What is their semiclassical imprint? What is their effect on observables in the exterior? These questions lie at the heart of many paradoxes and recent investigations regarding the semiclassical nature of black holes [4,5]. In these investigations, quantum information theory has proven to be an indispensable tool in providing a bridge between the

quantum nature of black holes and their classical spacetime geometry [6–9].

The fact that, in the semiclassical limit, *any* horizon must behave as though it contains internal degrees of freedom directly follows the basic principles of quantum mechanics and causality [10,11]. To illustrate this we consider the following gedanken experiment involving two experimenters: Alice and Bob. For definiteness, we consider the case where they perform their experiments in a black hole spacetime. Alice controls a quantum superposition which she attempts to keep coherent in the exterior of a black hole. Despite her best efforts, she cannot possibly control the long-range gravitational field of her superposition. The resulting superposed gravitational field penetrates the horizon of the black hole where an observer, Bob, could choose to measure it. If Bob can gain any information about Alice’s superposition, then her superposition must be decohered by a complementary amount. However, by causality, Bob’s actions cannot have any effect whatsoever on the state of her superposition. Furthermore, while the interior of a black hole is not an optimal place to perform such a sensitive experiment, we could imagine equipping Bob with better apparatuses and more assistants. From this gedanken experiment we conclude that Alice’s superposition must become entangled with—and subsequently decohered by—the black hole *as if* the black hole contains “maximally entangling” degrees of freedom. We note that

*Contact author: daine@mit.edu

†Contact author: gautam93@mit.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI. Funded by SCOAP³.

the crucial ingredient in this argument was the existence of an event horizon which causally separated Alice and Bob. Thus, these conclusions apply whenever such a horizon bounds the causal past of Alice’s worldline—e.g., Rindler horizons, cosmological horizons.

Motivated by this thought experiment we, together with R. M. Wald, have shown in a series of papers that, indeed, any (Killing) horizon will decohere a quantum superposition in its vicinity [12–14]. To explicitly show this, we considered the case where Alice creates a quantum spatial superposition of a massive (or charged) body—e.g., a nanoparticle—which, to a good approximation, produces a superposition of coherent states of the gravitational field. In the case of a black hole it was shown that, semiclassically, the black hole acquires low-frequency, fluctuating gravitational (and electromagnetic) multipole moments [14] which stimulate the emission of entangling, soft radiation into the black hole [12]. Furthermore, these multipole fluctuations closely mimic those of a material body—though it would appear that, in the gravitational case, extraordinary properties of matter would be required to reproduce the effect [15]. In total these results suggest that the decoherence caused by black holes is, in some sense, “extreme” given its size and temperature. However, these analyses only partly validate the conclusions of the above gedankenexperiment. How does this decoherence compare to the decoherence inflicted by an “optimal” Bob in the interior of the black hole? If a black hole behaves like a material body when viewed from the outside, then it should be possible to show that any family of observers—or “degrees of freedom”—in the interior can gain (at most) as much “which-path” information about Alice’s particle as is allowed by the degree of decoherence of Alice’s superposition in the exterior.

As a warm-up, in Sec. II A, we first address this question by revisiting the specific gedanken experiment considered in [12–14] where Alice creates a quantum spatial superposition. From Bob’s point of view, he is attempting to perform a “one-shot” experiment to distinguish the two possible states of Alice’s particle via their long-range field. In quantum information theory, there are many different measures of the distinguishability of states (e.g., the Holevo and Uhlmann fidelities, the relative entropy, the trace distance), each with its own operational meaning. The relevant notion of distinguishability in this case is the Uhlmann fidelity. For Alice’s particle at a finite time, the Uhlmann fidelity computes the “optimal purification” of her particle so as to minimize her decoherence [16]. In this sense, it is precisely a measure of “how decohered” her particle is up until that time. In the interior, the Uhlmann fidelity also bounds the probability that Bob, by performing a single optimal measurement, can become absolutely certain of the state of Alice’s particle. We thereby obtain the following statement: the decoherence of Alice’s superposed particle at any finite time is always greater than or

equal to the decoherence due to an optimal one-shot, unambiguous state discrimination test in the interior up until that time. At late times this inequality is saturated, as if the black hole (or interior observers/degrees of freedom) is performing an optimal measurement of Alice’s superposed long-range fields. Indeed, this relationship between the information accessible to Bob and the decoherence of Alice by soft quanta has been recently used as a “teleportation protocol” to send information into a black hole with negligible energy cost [17].

The gedanken experiment suggests, however, that such a relationship, between Alice’s experiment in the exterior and the decoherence due to optimal measurements made in the interior, should hold for *any* experiment performed in the vicinity of a black hole. In Sec. II we consider this more general relationship. In quantum information theory, a general experiment performed by Alice is a “quantum channel.” We introduce definitions for the intrinsic decoherence of a quantum channel, and the information of its complementary channel. The “channel decoherence” quantifies the failure to be able to recover her superposition by performing measurements strictly in the black hole exterior. The complementary “channel information” quantifies the information accessible to Bob, and bounds the probability of unambiguously distinguishing the output of his channel from a totally uncorrelated state. Section II gives precise definitions for the channel decoherence and information. As is often the case in quantum information theory [18–20], there is a trade-off between the information carried by a channel into the environment and the disturbance of the original system by that same environment. In Sec. II, however, we show that the channel decoherence and channel information enjoy an *exact* “information-decoherence trade-off.” Namely, we show the two quantities are precisely equal to one another.

Applied to black holes, the channel decoherence and channel information provide an *intrinsic* definition of the exterior decoherence, and interior information, of the black hole in the semiclassical regime, irrespective of any experimental protocol. We recall that, in ordinary experience, decoherence due to entanglement generated with any material body is, to some extent, reversible in principle. In Sec. II C we show that this is not the case for any semiclassical black hole (with any deviations from being suppressed in G_N). Rather, a black hole exhibits what might be called a “second law” for exterior decoherence. Additionally, we comment on how these ideas, being fundamentally quantum-information-theoretic in nature, might plausibly be extended beyond the semiclassical limit. In total, these results lend support to the perspective that a black hole affects its environment as though it were composed of maximally entangling degrees of freedom.

What of the more familiar notions of distinguishability, such as the Holevo fidelity, relative entropy, and the Rényi entropy? While less operationally relevant to the

decoherence effects considered in Sec. II, these notions of distinguishability have played a prominent role in analyzing the quantum mechanical properties of horizons ranging from the generalized entropy to the black hole information paradox [6,21–23]. How does the growth of these other notions of distinguishability compare to that of the Uhlmann fidelity on the horizon? In Sec. III, we address this question in the context of the gedanken experiment considered in [12–14,24] where the relevant states are coherent states on the horizon.

We answer these questions in Secs. III and IV. To explain the answer we first recall that the growth of distinguishability on the horizon is due to the emission of soft radiation. These soft radiative effects are mathematically equivalent to the soft radiative effects found in flat spacetime [25–29]. These effects result in well-known infrared divergences which have plagued scattering calculations¹ since the earliest days of quantum field theory [30]. Indeed, both effects are controlled by a “memory tensor” which on the horizon is given by

$$\Delta_{AB} = h_{AB}|_{V=\infty} - h_{AB}|_{V=-\infty}, \quad (1)$$

where $h_{AB}(V, x^A)$ are the radiative degrees of freedom on the horizon, V is the affine time, and x^A are angular coordinates. The memory captures the failure of the radiation to decay at asymptotically early or late times. States with nonvanishing memory contain an *infinite* number of soft quanta and lie in Hilbert space representations unitarily inequivalent to the “standard” Fock space of radiation. Analogously, the IR divergences encountered in scattering theory arise from the failure to include states with memory in the theory.

In Sec. III, we consider the practical calculation of distinguishability measures of radiation states on the horizon and the extent to which these measures capture effects due to soft radiation. The emission of an infinite number of soft quanta and the corresponding IR divergences occur only in the strict infinite-time limit if h_{AB} fails to decay. At any finite time, only a finite number of soft quanta can be produced. Therefore, one might expect that the IR issues highlighted above play no role in calculating the distinguishability of radiation states at any finite time. However, as we now explain, incorporating states with memory enables one to calculate finite time distinguishability measures in a simple and efficient manner.

We note that the practical computation of any of the above distinguishability measures for any two states requires the technology of “Tomita-Takesaki theory” and is an extremely difficult and unsolved problem. However we can, for our purposes, restrict to coherent states of

radiation since those were the radiation states most relevant to our original gedanken experiment² [10,11,13,14,24,31]. If h_{AB} is a classical radiation field (with no memory) on the horizon, we denote its corresponding coherent states as $U(h)|\Omega\rangle$ where $U(h)$ is the unitary operator generating the coherent state and $|\Omega\rangle$ is the vacuum state on the horizon. It was pointed out by Casini *et al.* that if the operator $U(h)$ factorizes across $\mathcal{C}$, then one can straightforwardly apply the technology of standard Tomita-Takesaki theory. One can obtain an explicit computation of the Tomita operator and the relative modular operator in this case and one can straightforwardly compute any of the above distinguishability measures.

However, $U(h)$ factorizes if and only if h_{AB} vanishes at the cut. If h_{AB} is nonvanishing, then there does not appear to be a simple description of the Tomita operator or modular operator within standard Tomita-Takesaki theory. However, h_{AB} can be shifted by a smooth V -independent trace-free tensor, whose electric-parity part is a supertranslation,

$$h_{AB}(V, x^A) \rightarrow h_{AB}(V, x^A) + k_{AB}(x^A). \quad (2)$$

By choosing k_{AB} appropriately, we show that the unitary can be essentially factorized about the cut. While this removes the above obstruction, the shifted radiation state fails to decay and the unitary operator factors now have nonzero horizon memory. Therefore, to obtain a simple description of distinguishability measures at a finite time, we must extend Tomita-Takesaki theory to include states with memory. We obtain such an infrared extension of Tomita-Takesaki and obtain explicit formulas for the Tomita operator and the relative modular operator between any two coherent states for any cut $\mathcal{C}$. Using this construction we explicitly compute the fidelity, relative entropy, and Rényi entropies of general coherent states on the horizon, with or without memory. We use this to analyze the quantum information content of soft radiation. We find that this quantum information content is well-captured by the fidelity, as might be expected from the original thought experiment. In that context, the fidelity is sensitive to the absolute number of soft gravitons emitted into the black hole (which may be arbitrarily large), as opposed to their energy (which may be arbitrarily small). The relative entropy is, by contrast, only very weakly sensitive to the number of soft photons, and is much more sensitive to the *energy* of the radiation emitted into the horizon.

Throughout this work, we develop several new techniques that may find wider utility. For example, the equivalence of channel decoherence and complementary channel information shown in Sec. II is a general result for

¹In flat spacetime, the IR divergences occur in four dimensions with a massless field. On the horizon, they occur in any spacetime dimension with a massive or massless field [13].

²Coherent states of soft radiation also describe the soft emission of scattering phenomena.

quantum channels, which are the most general models of physical processes acting on quantum states. For a particular class of channels, in which the system becomes entangled with two states of another system, Appendix C also proves a bound between the channel decoherence and the fidelity between those two quantum states. To analyze other distinguishability measures, Secs. III and IV extend the Tomita-Takesaki theory of coherent states so that it includes soft radiation, thus putting questions pertaining to the quantum information content of soft modes on rigorous mathematical footing. We believe these methods may find utility in contexts beyond those considered here.

The reader who is merely interested in the statement of the new results in Tomita-Takesaki theory may safely skip ahead to Secs. III and IV. There, as in much of the present article, we simplify to the case of electromagnetic radiation, and the gravitational case follows by close analogy. Section III extends the horizon algebra to include soft modes, and Sec. IV gives the formula for the relative Tomita operator in general coherent states. We use this to bound the more general and abstract information-decoherence relations of prior sections in terms of explicit fidelities. Section V extends the discussion of the present work to the context of scattering theory, where directly analogous considerations apply. Finally, Sec. VB applies the techniques of Sec. II to show how a generic causal horizon can give rise to fundamental decoherence in its vicinity.

Throughout this paper we work in units where $\hbar = c = G_N = \epsilon_0 = 1$. Lowercase Roman indices denote indices in spacetime, while uppercase Roman indices denote indices tangent to the spatial cross sections of the horizon. Greek indices label the numerical components of a tensor in a basis. We denote quantized observables in a $*$ -algebra by the boldface version of the corresponding classical observable. For example, the classical electric field E becomes $\mathbf{E}$ in the quantum theory. We denote von Neumann algebras in fraktur as $\mathfrak{A}$, and restrictions of algebras to the past of a horizon cut $\mathcal{C}$ by $\mathfrak{A}_{\mathcal{C}}$.

II. INFORMATION-DECOHERENCE TRADE-OFF

In this section we show that the decoherence inflicted by a horizon is a consequence of an “information-disturbance trade-off” between the coherence of any quantum mechanical system in the exterior and the information about that body accessible to any observer in the interior. In Sec. II A we explain how one can already see a version of this relationship in context of the specific gedanken experiment involving the coherence of a quantum spatial superposition considered in [12–14,24]. To prove the more general relationship suggested in the Introduction, in Sec. II A we generalize the notions of decoherence and information accessible in the interior to any experiment and prove their equality. In the case of horizons, this proves the statement that the decoherence inflicted by a black hole—or any horizon satisfying the assumptions stated below—is related

by the exact information-decoherence trade-off to the information available in its interior. For horizons in particular we show that the decoherence of the exterior grows monotonically with time, as would also be expected from a body comprising maximally entangling degrees of freedom.

A. Analysis of the original gedanken experiment

In this section we revisit the original gedanken experiment considered in [12–14,24] involving a quantum spatial superposition in the vicinity of a horizon. From this example, we will identify (1) the appropriate notion of decoherence of such a superposition at any finite time and (2) a precise operational meaning of this decoherence in terms of an “unambiguous state discrimination test” performed in the interior. While our considerations will be valid for any Killing horizon [13]—e.g., a stationary black hole horizon or de Sitter horizon—we will, for simplicity, restrict attention to Alice performing her experiment in the vicinity of a Rindler horizon in Minkowski spacetime.

An experimenter, Bob, assembles his lab equipment on the opposite side of the Rindler horizon of a uniformly accelerating experimenter, Alice. Now suppose that Alice prepares a charged nanoparticle in a definite x -spin state. She then uses a Stern-Gerlach apparatus to adiabatically manipulate the particle. If the Stern-Gerlach were oriented along the x -direction, then it would simply displace her particle. If, however, the apparatus is oriented in the z -direction, she can use it to prepare a quantum spatial superposition of her particle’s wave function. After doing this, her superposed particle sources a superposition of coherent states of the electromagnetic field (nearly identical considerations apply for gravity, with mass replacing charge, and quadrupole radiation replacing dipole radiation). The global Hilbert space is now

$$\mathcal{H} = \mathcal{H}_{\text{Alice}} \otimes \mathcal{H}_{\mathcal{H}^+}. \quad (3)$$

We denote the coherent states $|\Psi_1\rangle$ and $|\Psi_2\rangle$ in the horizon Hilbert space $\mathcal{H}_{\mathcal{H}^+}$, while denoting the z -spin states (and corresponding spatial wave functions A_i) of Alice by $|A_1, \downarrow\rangle$ and $|A_2, \uparrow\rangle$ in $\mathcal{H}_{\text{Alice}}$. In our thought experiment, the resulting entangled state at late times will be

$$\frac{1}{\sqrt{2}}(|A_1, \downarrow\rangle \otimes |\Psi_1\rangle_{\mathcal{H}^+} + |A_2, \uparrow\rangle \otimes |\Psi_2\rangle_{\mathcal{H}^+}), \quad (4)$$

where $\mathcal{H}^+$ denotes the Rindler horizon and labels the corresponding, entangled states of radiation. Meanwhile, suppose that Bob is interested in distinguishing between these two field configurations by performing a measurement on his side of the Rindler horizon.

Previous analyses have focused on the decoherence of Alice’s superposition, as quantified by

$$1 - |\langle\Psi_1|\Psi_2\rangle_{\mathcal{H}^+}|, \quad (5)$$

using the inner product of the final states of entangling radiation that her experiment sources on the horizon. In other words, the final decoherence is controlled by the off-diagonal elements of Alice’s final reduced density matrix.

In addition to the final decoherence of her particle, it is natural to ask for the decoherence of her superposition as a function of time. Such a refined definition of decoherence would be useful for monitoring the dynamical buildup of Alice’s decoherence outside a horizon, but this is complicated by the fact that the “off-diagonal elements of Alice’s reduced density matrix” are not well-defined while Alice’s superposition is still held open. Formally, such a quantity would involve “tracing out” and taking the inner product between two different Coulomb fields, which are not independent degrees of freedom from Alice’s system. It could be argued that the two fields should be considered to be orthogonal, but the resulting apparent, total decoherence of Alice’s superposition is “false decoherence” [32] that in no way obstructs the recovery of Alice’s initial state.

A natural candidate at finite times, which avoids the problem of false decoherence, is to compute the off-diagonals of Alice’s reduced density matrix if, after the chosen time, she were to perform the remainder of her experiment so as to maximize the recovery of the coherence of her original state. More specifically one may ask, what is the minimum decoherence that the horizon will inflict on Alice, if, in region $\mathcal{A}_{II}$ of Fig. 1, she is allowed to manipulate her particle so that it sources radiation in region $\mathcal{A}_{II}$ that minimizes her final decoherence? This protocol gives one operational notion of the “minimum decoherence” (or “maximum recoverable coherence”), and forms the basis of the analysis carried out in [24]. This seems like a reasonable candidate, but it is by no means

clear that this is the distinguished, operationally relevant notion of decoherence at a finite time.

Even the idea of using the off-diagonals of Alice’s density matrix to diagnose coherence can be questioned, because while the quantities clearly control the coherence of her superposition, their precise operational meaning is not clearly the right one. In our case, however, we have already argued that the decoherence should have an operational meaning in terms of distinguishability by Bob. To single out the operationally pertinent signature of decoherence, therefore, we will not look to Alice; we will look to Bob.

While singling out an operationally distinguished measure of decoherence appears challenging, what is clear from the thought experiment is that there should be a complementary notion of the which-path distinguishability information available to the “observer,” Bob, in the interior. We will see that this notion of distinguishability is operationally transparent even at finite times, and it gives a natural complementary notion of decoherence for Alice that will, in fact, correspond directly to the notion of [24]. To see this, we will put ourselves in Bob’s position.

Suppose Bob plans to complete his measurement in a particular finite time interval, so that his measurement concludes before he enters the future light cone of a given horizon cut $\mathcal{C}$ as illustrated in Fig. 1. If we allow Bob (and any number of assistants) to perform any conceivable measurement or joint measurement in his spacetime region $\mathcal{B}|_c$, what is the upper bound on the probability for him to successfully determine the position of Alice’s particle?

The precise answer to this question depends on the precise specification of the task assigned to Bob. Suppose Bob is allowed to perform his measurement anywhere outside the future light cone of a horizon cut $\mathcal{C}$. The resulting states of the quantum field, restricted to Bob’s region, will be mixed states $\Psi_1|_c$ and $\Psi_2|_c$. Now suppose that, after performing some measurement, Bob is allowed to guess whichever state would be more probable to give the outcome he has observed: we say “guess,” because even if his measurement yields an outcome that leaves some ambiguity as to which state he has measured, Bob might be allowed to guess the most probable state, given that outcome. The success probability of this “state discrimination” task will be bounded in terms of the trace distance between the two states [33]. But Bob’s ability to *guess* the right state has just as much to do with classical correlation as it does with entanglement and decoherence, so it cannot be the right operational task.

If, on the other hand, we disallow guessing, we isolate a different quantity. If we set the rules such that Bob may only declare success if he is *certain* of the state, then the so-called “Uhlmann fidelity” is the relevant quantity. He could obtain certainty of the state in such a “one-shot” test if he were to obtain a measurement outcome that is *only* possible in one of the two states he must distinguish between. When this additional “no guessing” rule is imposed, Bob’s task is known as an unambiguous state discrimination task, and




FIG. 1. Bob can perform any possible measurement in region $\mathcal{B}_c$, defined as the region on the $-z$ side of the horizon that lies to the past of the future light cone of the horizon cut $\mathcal{C}$.

the Uhlmann fidelity $f_U(\Psi_1|_c, \Psi_2|_c)$ bounds his success probability to unambiguously distinguish between the states $\Psi_1|_c$ and $\Psi_2|_c$ on the horizon, restricted to the past of $\mathcal{C}$, as illustrated in Fig. 1.

Given two mixed states ρ and σ , the Uhlmann fidelity is the supremum of the inner products between purifications of the states. We denote these purifications as $|\rho\rangle_\pi$ and $|\sigma\rangle_\pi$ with respect to the Hilbert space representation π . In other words, the Uhlmann fidelity is defined as [16]

$$f_U(\rho, \sigma) := \sup_{\pi} |\langle \rho | \sigma \rangle_{\pi}|^2. \quad (6)$$

Here the supremum runs over all Hilbert space representations π of a given $*$ -algebra such that $|\rho\rangle$ and $|\sigma\rangle$ are vectors in the corresponding Hilbert space $\mathcal{H}$ and are purifications of the states ρ and σ , respectively.

The probability that Bob will successfully distinguish the state in this “unambiguous” sense is then bounded by

$$P_{\text{Bob}}^{\text{certain}}(\mathcal{B}|_c) \leq 1 - \sqrt{f_U(\Psi_1|_c, \Psi_2|_c)}. \quad (7)$$

As we shall see, when we relax the time constraints on Bob, this bound in fact becomes an equality.

It is this unambiguous notion of state discrimination, which measures Bob’s ability to achieve *certain*, that is most closely tied to the decoherence of Alice’s particle. This is easily seen from the fact that the Uhlmann fidelity has another, complementary interpretation from the point of view of Alice in the exterior: it gives the infimum of Alice’s unavoidable decoherence discussed above. Now, this measure of decoherence stands on firm operational footing, in terms of the knowledge that can be obtained by an observer coupling to Alice’s environment.

The Uhlmann fidelity gives the infimum of the decoherence $D_{\text{Alice}}|_c$ that the horizon will inflict on Alice, if, in region $\mathcal{A}_{II}$ of Fig. 1, she is allowed to manipulate her particle so that it sources radiation in region $\mathcal{A}_{II}$. The resulting final decoherence of Alice is

$$D_{\text{Alice}}|_c := 1 - \sqrt{f_U(\Psi_1|_c, \Psi_2|_c)}. \quad (8)$$

From Alice’s point of view, the “optimal trajectory of her particle” can be seen as sourcing precisely the “optimal purification” of Eq. (6) on the horizon [24]. Thus, we learn that in general, the probability that Bob succeeds is bounded by the supremum of Alice’s recoverable coherence,

$$P_{\text{Bob}}^{\text{certain}}(\mathcal{B}|_c) \leq D_{\text{Alice}}|_c. \quad (9)$$

In fact, when we place no time restrictions on Bob and allow him access to the entire region on his side of the horizon,³ it is

³In the case of a black hole in an asymptotically flat spacetime, the saturation is never achieved, because of the entanglement in the vacuum between the horizon and null infinity.

guaranteed that Bob can saturate the inequality by some (possibly very complicated) measurement [34], and the probability that an optimal measurement by Bob⁴ will unambiguously reveal the state of Alice’s system is precisely equal to her decoherence:

$$P_{\text{Bob}}^{\text{certain}}(\mathcal{B}) = 1 - \sqrt{f_U(\Psi_1, \Psi_2)} = D_{\text{Alice}}. \quad (10)$$

At least in this scenario, one may view the horizon as carrying out an unambiguous state discrimination test on the particular quantum states prepared by Alice in its exterior. We have assumed that in Eq. (4), Bob must distinguish between two equiprobable states. In more general situations, the states may have unequal prior probabilities p_+, p_- with $p_+ > p_-$. If $f_U \geq p_-/p_+$, the probability that Bob will succeed is bounded by the fidelity itself, $P_{\text{Bob}}^{\text{certain}} \leq p_+(1 - f_U)$, rather than the square root [35], and the interpretation in terms of unambiguous state discrimination again holds.

This result, and the results of [12–14, 24] are, however, incomplete: they describe the decoherence due to horizons as inflicted on a particular class of experiments of the form of Eq. (4). This is far from the general result claimed in the Introduction. The original gedanken experiment suggests that there should be a general equivalence between the decoherence in the vicinity of a horizon, and the distinguishability of the resulting states behind the horizon. In fact, this equivalence suggests that the causal structure of the horizon should be intimately connected with this operational fact and, taken together, causality and complementarity dictate that this principle should hold regardless of the particulars of Alice’s experiment. Equations (9) and (10) only pertain to when Alice creates a superposition of two states, as in Eq. (4), so this is far too restrictive a notion of decoherence to capture the original idea of the gedanken experiment. To capture the general relationship, we will need a more general characterization of decoherence, and of the information available to Bob. Nevertheless, we will remain rooted in the operational interpretation of unambiguous state discrimination, and the generalizations we define will recover the familiar definition of the final decoherence of Alice’s particle when restricted to the experiment of Eq. (4). In doing this we will arrive at a definition of the decoherence intrinsic to the horizon itself.

B. Equivalence of information and decoherence for general experiments

In this subsection we will no longer assume anything about the particular form of Alice’s experiment or Bob’s. We will formulate an equivalence between the decoherence of

⁴We here neglect any backreaction of Bob’s experiment on the environment. It is, for example, possible that a practical realization of the optimal protocol might produce a dramatic change to the spacetime in Bob’s region such that a singularity forms. This could prevent Bob from achieving the optimal protocol in practice. Nevertheless, the physics in Alice’s region will proceed *as if* Bob were performing an optimally discriminating measurement.

Alice and the information available to Bob or arbitrary cuts of the horizon, and for arbitrary quantum experiments by Alice and Bob. These more general notions of decoherence and information will recover the original notions when applied to the specific experiment of Alice and Bob, and will carry suitably generalized operational definitions in terms of unambiguous state discrimination. No clarity will be lost if the reader continues to assume that the quantum state is of the form given Eq. (4), but in fact the following discussion will hold for any state Alice can create on her side of the horizon and for any experiment Bob can do on his.

A suitable notion of decoherence should satisfy some essential properties. It should be the following:

- (D.i) Intrinsic to the horizon in that it is independent of the exterior state.
- (D.ii) Localized to the past of a given horizon cut.
- (D.iii) Equivalent to the ordinary notion of decoherence in the appropriate limit.

The complementary notion of information should be the following:

- (I.i) Independent the choice of experiment used to measure the information behind the horizon.
- (I.ii) Localized to experiments performable behind the horizon and to the past of the causal future of a given horizon cut.
- (I.iii) Equivalent to an ordinary notion of distinguishability behind the horizon in the appropriate limit.

After developing these general definitions, we will apply them to formulate channel-level bounds on unambiguous state discrimination tasks involving general quantum states in the exterior.

The key to achieving properties D.i and I.i is the notion of “quantum channel” $\mathcal{N}$, which is the most general description of any physical experiment performed by either Alice or Bob. A channel is a linear, completely positive, trace-preserving map, but is otherwise unconstrained. Every channel can be realized by coupling some “environment” (in our case, the horizon) to the system and carrying out the subsequent open evolution [36]. In our case,

$$\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+}: \mathcal{S}(\mathcal{H}_{\text{Alice}}) \rightarrow \mathcal{S}(\mathcal{H}_{\text{Alice}}), \quad (11)$$

where $\mathcal{S}(\mathcal{H}_{\text{Alice}})$ is the state space (density matrices) on Alice’s Hilbert space. $\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+}$ will be responsible for the decoherence of Alice’s state outside the horizon. For every channel $\mathcal{N}$ there is a “complementary” channel $\mathcal{N}^c$ (unique up to isometric equivalence) corresponding to the change in this environment as the joint evolution applies $\mathcal{N}$ to the original system. In our case, for example, a complementary channel

$$(\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+})^c: \mathcal{S}(\mathcal{H}_{\text{Alice}}) \rightarrow \mathcal{S}(\mathcal{H}_{\mathcal{H}^+}) \quad (12)$$

would map Alice’s initial state onto the resulting state of the *horizon* (equivalently, the state in Bob’s region). This

complementary channel transmits the entangling photons/gravitons produced by Alice’s experiment into Bob’s region of the spacetime, thereby providing the complementary which-path information that Bob must measure. The relevant formal definitions, as well as the substance of the proofs developed in this section, are presented in Appendix A. The goal of this section is to present the results in an intuitive and physically motivated context.

1. Definition of channel decoherence

An important measure of quantum coherence is the recoverability of a quantum state after that state has passed through a channel $\mathcal{N}$. Decoherence manifests as an inability to recover the full initial state, and we will provide a definition that formalizes this notion.

If a channel perfectly preserves the coherence of states, then there exists a recovery channel $\mathcal{R}$ (sometimes called a “decoding channel”), independent of Alice’s initial state such that

$$\mathcal{R} \circ \mathcal{N} = \text{Id}. \quad (13)$$

This means that the experiment is perfectly reversible and is the condition of *exact* quantum error correction. A trivial example of a quantum channel that is perfectly reversible is a unitary channel, such that the decoding channel is simply the adjoint unitary channel. Of course, practical experiments are not perfectly reversible, so we will use the language of approximate quantum error correction to define a general measure of channel decoherence that ranges between full coherence and total decoherence.

To do this, we need a more refined characterization of the similarity of quantum channels than strict equality. We seek a notion of “similarity” between the recovered channel $\mathcal{R} \circ \mathcal{N}$ and the identity channel Id that runs between 1 (equality/recoverability) and 0 (total distinguishability/unrecoverability). The “channel fidelity” gives the appropriate notion of similarity of quantum channels [37]. Here we build up its definition piece-by-piece. Along the way we explain each piece of the definition in the special case where the channels model the Stern-Gerlach experiment of Sec. II A.

Given any two quantum channels $\mathcal{N}$ and $\mathcal{M}$ and a chosen input state ω , the similarity of their outputs is captured by the Uhlmann fidelity as

$$f_{\text{U}}(\mathcal{N}(\omega), \mathcal{M}(\omega)). \quad (14)$$

We now seek a notion of the similarity between the channels themselves, independent of any particular initial state. Not all states are equally good when it comes to discriminating one channel from another. Consider, for example, a situation in which $\mathcal{N}$ corresponds to Alice running a spin through her Stern-Gerlach apparatus in flat spacetime, while we take $\mathcal{M}$ to be the identity channel. Then, if Alice prepares her spin in

a definite z -spin up state, the output of $\mathcal{N}$ and the output of Id will be identical (we assume she produces only negligible radiation). Should we then conclude that the Stern-Gerlach channel is identical to the identity channel? Of course, the answer is that we have merely chosen a poor state for discriminating between these two channels. A natural quantity, therefore, for the intrinsic distinguishability of channels, is to consider the fidelity when the *most discriminating* state is considered. In the case where $\mathcal{M}$ is the identity channel and $\mathcal{N}$ models Alice's z -aligned Stern-Gerlach experiment, the most discriminating state would be a state of definite x -spin of Alice's particle. In general, the formula becomes [38]

$$\inf_{\omega} f_{\text{U}}(\mathcal{N}(\omega), \mathcal{M}(\omega)). \quad (15)$$

This, however, has an undesirable property, peculiar to quantum entanglement: this quantity is unstable under the enlargement of the Hilbert space, even when the channels act trivially on the additional Hilbert space factors. That is to say, for some auxiliary Hilbert space dimension k ,

$$\inf_{\omega} f_{\text{U}}(\text{Id}_k \otimes \mathcal{N}(\omega), \text{Id}_k \otimes \mathcal{M}(\omega)) \leq \inf_{\omega} f_{\text{U}}(\mathcal{N}(\omega), \mathcal{M}(\omega)). \quad (16)$$

This might appear to be a rather formal property to concern ourselves with, but it has practical implications for Alice. Suppose, for example, that Alice does a more complicated experiment, in which she prepares her spin in a bell pair with another qubit in her laboratory. As before, suppose that $\mathcal{M}$ is the identity channel, so we are comparing the effect of $\mathcal{N}$ to the identity, where $\mathcal{N}$ acts trivially on the other qubit. If $\mathcal{N}$ destroys the entanglement of Alice's bell pair, we should consider this channel to be *farther* from the identity channel than one that preserves the entanglement. A better measure of the similarity of channels, therefore, would take this into account. A simple remedy presents itself, which is to consider the auxiliary Hilbert space (the other qubit in Alice's lab) that is most sensitive to such effects. By extremizing over this auxiliary Hilbert space dimension,⁵ we arrive at the definition of the channel fidelity [37],

$$F(\mathcal{N}, \mathcal{M}) := \inf_k \inf_{\omega} f_{\text{U}}((\text{Id}_k \otimes \mathcal{N})(\omega), (\text{Id}_k \otimes \mathcal{M})(\omega)), \quad (17)$$

which quantifies the similarity of two quantum channels.

Using this definition, the exact error correction condition of Eq. (13) can be relaxed into an approximate error correction condition. Given a recovery channel $\mathcal{R}$, this approximate recoverability condition is

⁵The Hilbert space dimension will never need to be larger than the Hilbert space dimension of the domain of the original channels, e.g., Alice's original spin Hilbert space [39].

$$F(\mathcal{R} \circ \mathcal{N}, \text{Id}) \simeq 1. \quad (18)$$

Conversely, however, if $F(\mathcal{R} \circ \mathcal{N}, \text{Id}) \ll 1$, we cannot conclude that $\mathcal{N}$ is a decohering channel. It may well be that $\mathcal{R}$ is simply a bad choice of recovery channel. The intrinsic decoherence of a channel should depend only on whether it is recoverable *in principle*, so we should take the supremum over recovery channels:

$$\sup_{\mathcal{R}} F(\mathcal{R} \circ \mathcal{N}, \text{Id}) \quad (19)$$

characterizes the supremum of the recoverable coherence of a quantum channel $\mathcal{N}$. This immediately implies a definition of the *irrecoverable decoherence* of a channel:

Definition 1. (Channel decoherence).

$$D(\mathcal{N}) := 1 - \frac{1}{1 - 1/d} [\sup_{\mathcal{R}} F(\mathcal{R} \circ \mathcal{N}, \text{Id}) - 1/d]_+. \quad (20)$$

For $d > 1$, the dependence on the Hilbert space dimension d of the domain of $\mathcal{N}$ is chosen so that

$$0 \leq D(\mathcal{N}) \leq 1, \quad (21)$$

where $D(\mathcal{N}) = 0$ indicates no decoherence and $D(\mathcal{N}) = 1$ indicates saturation of the decoherence measure.⁶ Here $[x]_+ := \max\{x, 0\}$.⁷ Intuitively, the channel decoherence (20) can be understood as the normalized loss of the optimal worst-case fidelity between an original state ω and what remains after acting with $\mathcal{N}$ and optimizing over recovery channels $\mathcal{R}$. As promised, it is a measure of irrecoverable decoherence inflicted by $\mathcal{N}$.⁸

⁶For an infinite-dimensional input Hilbert space, $1/d$ is understood as zero.

⁷Alternatively, one may omit the positive-part truncation and define $\tilde{D}(\mathcal{N}) := 1 - (\sup_{\mathcal{R}} F(\mathcal{R} \circ \mathcal{N}, \text{Id}) - 1/d)/(1 - 1/d)$, together with the analogous modification to the forthcoming definition of $\tilde{\mathcal{I}}$. Whereas the truncated definition treats complete dephasing as maximal decoherence, the untruncated definition continues to distinguish channels that also destroy the surviving classical information: $\tilde{D} = 1$ at complete dephasing, while $\tilde{D} > 1$ is possible for more destructive channels (with complete erasure giving $\tilde{D} = 1 + 1/d$ in finite dimensions). For the pure, equiprobable two-state case of Appendix A 2, however, the cutoff is inactive, so $\tilde{D} = D = 1 - |\langle \Psi_1 | \Psi_2 \rangle| = P_{\text{Bob}}^{\text{certain}}$; only the general probability-valued interpretation is lost once $\tilde{D} > 1$. The exact information-decoherence relation $\tilde{D}(\mathcal{N}) = \tilde{\mathcal{I}}(\mathcal{N}^c)$ remains intact.

⁸In the context of the complexity theory of quantum fields, Hollands and Ranallo [40] have introduced a closely related quantity as a measure of the *complexity* of a quantum channel. The exploration of the possible relationships between channel decoherence and channel complexity would seem to be worthy of further investigation, especially in light of the discussion of Sec. II D.

2. Definition of channel loss and channel information

Complementarity dictates that when Bob can distinguish between the states on his side of the horizon, i.e., he has “which-state information,” then Alice’s superposition must be decohered. To formalize this notion, we must formalize the notion of the information available to Bob. As a first step, we formalize the *loss* of information by a channel transmitting information from Alice to Bob.

As explained in Sec. II A, performing a Stern-Gerlach experiment in the vicinity of a horizon can provide a significant amount of which-path information to Bob. As shown in [13], however, the strength of the decoherence diminishes if Alice produces a smaller superposition—i.e., she reduces the spatial separation of the branches of her particle’s spatial wave function. Then the resulting long-range fields in Bob’s region are more difficult to distinguish and the amount of which-path information available to Bob is reduced.

The channel responsible for conveying this information from Alice to Bob is the complementary channel $(\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+})^c: \mathcal{S}(\mathcal{H}_{\text{Alice}}) \rightarrow \mathcal{S}(\mathcal{H}_{\mathcal{H}^+})$. If Alice produces an extremely small superposition, that channel will be very lossy: Bob will get very little information about Alice’s superposition, because the output state on the horizon will depend only very weakly on Alice’s superposition.

Generalizing from this example, we say that a channel $\mathcal{N}_{\text{lossy}}$ is lossy if it always outputs nearly the same, fixed state χ_0 , in which case

$$\inf_{\omega} f_U(\mathcal{N}_{\text{lossy}}(\omega), \chi_0) \simeq 1. \quad (22)$$

It is possible, however, to choose χ_0 so that it is distinguishable from all possible outputs of $\mathcal{N}_{\text{lossy}}$ (perhaps, for example, it lies outside the range of $\mathcal{N}_{\text{lossy}}$). Then a small fidelity would say nothing about the lossiness of $\mathcal{N}_{\text{lossy}}$ and rather more about the choice of χ_0 . When defining the “channel loss,” therefore, we should consider the fixed state χ that is most similar to the output of the channel, while still carrying no information about the actual, particular state that was input to $\mathcal{N}$:

$$\mathcal{L}(\mathcal{N}) := \sup_{\chi} F(\mathcal{N}, \mathcal{C}_{\chi}), \quad \mathcal{C}_{\chi}(\omega) := \text{Tr}(\omega)\chi. \quad (23)$$

Intuitively, the channel loss $\mathcal{L}(\mathcal{N})$ can be understood as a measure of whether, after acting $\mathcal{N}$ on ω , the result is similar to some state that carries no information about ω whatsoever.

If a channel is nearly lossless, it preserves a great deal of information about the input state. We define the information of a quantum channel $\mathcal{N}$ as follows:

Definition 2. (Channel information).

$$\mathcal{I}(\mathcal{N}) := 1 - \frac{1}{1 - 1/d} [\mathcal{L}(\mathcal{N}) - 1/d]_+. \quad (24)$$

Definition 2 measures the information *preserved* by $\mathcal{N}$. As with the channel decoherence,

$$0 \leq \mathcal{I}(\mathcal{N}) \leq 1, \quad (25)$$

where $\mathcal{I}(\mathcal{N}) = 1$ indicates saturation of the information measure and $\mathcal{I}(\mathcal{N}) = 0$ indicates a maximally lossy channel. We have already seen that, when Bob is presented with two equally probable states to discriminate between, the unambiguous state discrimination probability is controlled by the square root of the fidelity by Eq. (7). As discussed in Sec. II A, an unequal-prior bound on unambiguous state discrimination can be given in terms of the fidelity itself. Therefore, Eq. (23) can be understood as giving an operational bound on Bob’s ability to perform unambiguous state discrimination.

3. Information-decoherence trade-off

As anticipated, a tight relation holds between the amount of decoherence $D(\mathcal{N})$ inflicted by a quantum channel and the information shunted by the *complementary* channel $\mathcal{N}^c$ into the environment, $\mathcal{I}(\mathcal{N}^c)$. In fact, in Appendix A 1 we leverage an information-disturbance trade-off relation proved by Bény and Oreshkov [20] to prove these two operationally distinct quantities are precisely equal. In what might be called an information-decoherence trade-off, we prove the following equality:

$$\mathcal{I}(\mathcal{N}^c) = D(\mathcal{N}). \quad (26)$$

This equality holds for arbitrary finite-dimensional quantum channels, and for infinite-dimensional channels whenever the preceding fidelities and complementary-channel theorem apply,⁹ but we will concern ourselves with its application to the original gedanken experiment of Sec. II A. Recall that Alice used the spin degree of freedom of her particle to control her superposition, using a Stern-Gerlach apparatus, so that the final coherence of the paths of her particle is determined by its coherence in the spin basis. We rescale the affine generators of the horizon so that the cut $\mathcal{C}$ is labeled by a single affine time coordinate V_c . This is, in turn, related to the Killing time t in Alice’s laboratory by $V_c = e^{\kappa t}$, where κ is the surface gravity of the horizon. The channel $\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+}(t)$ evolves the initial state of Alice’s spin into the state of the spin as it evolves past region $\mathcal{A}_I$, accounting for any decoherence caused by the horizon. The complementary channel $(\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+})^c(t)$ maps that same initial spin onto a mixture of two distinct states of

⁹Although Bény and Oreshkov work in finite dimensions, an infinite-dimensional extension follows from adapting their Stinespring-Uhlmann argument to normal channels, replacing min and max by inf and sup, and using Sion’s theorem [41] to justify their minimax interchange. For regional algebras, fidelity must be formulated algebraically, and a “complementary channel” produces a state on the commutant algebra of the Stinespring dilation of the output domain.

coherent radiation on the horizon, which are responsible for Alice’s decoherence.

How does the channel decoherence inflicted by the horizon relate to the earlier notion of decoherence $D_{\text{Alice}|c}$? At finite times, the channel decoherence carries a distinct, but closely related, operational meaning. Whereas $D_{\text{Alice}|c}$ computed the fidelity between optimal purifications of Alice’s state given her final state at $V = V_c$, the channel decoherence $D(\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+}(t))$ takes an infimum over all states and is therefore an intrinsic decoherence due to the horizon itself, as required by D.i.

At late times, the two definitions become precisely equal. Appendix A 2 computes the exact value of $D(\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+})$ in the infinite-time limit, as $V_c \rightarrow \infty$ (correspondingly, no time constraints are placed on Bob). In that case, optimal recovery merely needs to rotate away the phases of the off-diagonal components of Alice’s reduced density matrix, and we find

$$D(\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+}) = D_{\text{Alice}} \quad \text{for } t \rightarrow \infty. \quad (27)$$

This, of course, also implies that

$$\mathcal{I}((\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+})^c) = P_{\text{Bob}}^{\text{certain}} \quad \text{for } t \rightarrow \infty, \quad (28)$$

so we have satisfied our requirements I.iii and D.iii. Looking back on Sec. II A with this more general perspective, we see that the complementary roles of the decoherence and distinguishability measures in (9) and (10) are consequences of the exact information-decoherence trade-off of Eq. (26). Operationally, and now in full generality, the complementary channel furnishes a channel-level bound on unambiguous state discrimination tasks involving quantum states in the exterior.

At finite times, the optimal recovery channel is more difficult to compute, but is undoubtedly still closely related to the optimal protocol established in [24], and therefore the Uhlmann fidelity between the two states sourced by Alice’s superposition. Despite the difficulty involved in evaluating the channel decoherence at finite times, Appendix C derives a bound between them:

$$D(\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+}(t)) \geq D_{\text{Alice}|c}. \quad (29)$$

This bound is saturated at late times.

The channel decoherence is an intrinsic definition of the decoherence due to a quantum channel, satisfying all of our requirements D.i–D.iii. It is, by definition, a measure of the unrecoverable decoherence inflicted by a quantum channel. Furthermore, it is precisely equal to the channel information of the complementary channel, which itself satisfies all of our requirements I.i–I.iii and provides an operational bound on one-shot unambiguous state discrimination.

Despite its operational clarity and suitability in satisfying the exact information-decoherence trade-off (IDT),

however, the channel decoherence and channel information are difficult to compute because of the extremization required. Computing the channel decoherence at finite times is likewise technically challenging; Eq. (29) gives the bound above. The simple Uhlmann-fidelity decoherence $D_{\text{Alice}|c}$ is, by contrast, much easier to compute in general than the rather difficult $D(\mathcal{N})$. Since we are interested in practical computations, later sections will be concerned with developing a framework for computing practical distinguishability measures, which can be used to bound the exact decoherence $D(\mathcal{N})$ and information $\mathcal{I}(\mathcal{N}^c)$. These distinguishability measures, such as the relative entropy and Rényi entropies, will obey similar bounds to the channel decoherence and channel information, albeit sometimes with a less clear or relevant operational meaning and without the exact IDT of Eq. (26). The computation of such distinguishability measures is, of course, also of independent interest to a broader audience of those interested in the quantum information theory of fields.

C. Monotone decoherence: A hallmark of horizons

The original gedanken experiment suggests that a horizon, as viewed from the outside, may be viewed as though it is made of highly entangling degrees of freedom. More abstractly, the apparent equivalence between “space-time horizons” and “maximally entangling degrees of freedom” raises the possibility that the presence of a horizon may be characterized in purely operational and quantum information-theoretic terms. Having developed a more general set of tools, we are now equipped to formulate such a characterization.

We have thus far focused on the effects of the horizon on a given subsystem of the exterior, namely Alice’s superposition. For a description of the effects of a horizon that is completely intrinsic to it, however, it is useful to note that evolution in the presence of the horizon can be understood as a channel acting on *any* subsystem of the exterior (or the entire exterior), including Alice and her experiment.

Let t label the Killing time coordinate corresponding to time in Alice’s laboratory. As the horizon cut $V_c = e^{\kappa t}$ advances in time along the horizon, Bob is able to measure a strictly larger region of spacetime on his side of the horizon. Let us now take

$$\mathcal{N}_{\text{ext}}^{\mathcal{H}^+}(t) : \mathcal{S}(\mathcal{U}_{\text{ext}}(0)) \rightarrow \mathcal{S}(\mathfrak{A}_{\text{ext}}(V_c)), \quad (30)$$

where $\mathcal{S}(\mathcal{U}_{\text{ext}}(0))$ denotes the state space of the full exterior (in which case $\mathcal{U}_{\text{ext}} = \mathfrak{A}_{\text{ext}}$) at $t = 0$ or of any fixed code subspace thereof. The channel maps the initial state of the exterior to the state at a later time under the influence of the black hole—namely, the state exterior to a horizon cut $V = V_c$. We label the channels by t because, as V_c advances, the flow gives a one-parameter family of channels

$$(\mathcal{N}_{\text{ext}}^{\mathcal{H}^+})^c(t) : \mathcal{S}(\mathcal{U}_{\text{ext}}(0)) \rightarrow \mathcal{S}(\mathfrak{A}_{\mathcal{H}^+|_c}(V_c)). \quad (31)$$

As Alice’s Killing time t advances outside the horizon, $(\mathcal{N}_{\text{ext}}^{\mathcal{H}^+})^c(t)$ targets states on one-parameter family of horizon algebra inclusions. This corresponds to a restriction of the algebra of observables in the exterior, where the flow corresponds to Killing time. Heuristically, this corresponds to the fact that, as time progresses, an exterior observer must “trace out” more degrees of freedom as they fall into the horizon. The composition of an algebra restriction with a recovery channel is itself a valid recovery channel, so the supremum over recovery channels in Definition (20) guarantees that the decoherence of the exterior is non-decreasing¹⁰:

$$\frac{d}{dt} D(\mathcal{N}_{\text{ext}}^{\mathcal{H}^+}(t)) \geq 0. \quad (32)$$

This is a succinct statement of the extraordinary efficiency of horizons, including black hole and de Sitter horizons, at decohering their environments. Semiclassically—i.e., the limit as $G_N \rightarrow 0$ —this is an exact equation. In existing models of quantum gravitational effects, however [43,44], it seems likely that the right-hand side of this equation will receive corrections suppressed in powers of G_N . The degree to which this equation holds at finite G_N , then, may prove useful in signaling the emergence of an effective horizon.

It is notable that Eq. (32) provides a consequence of a semiclassical horizon entirely in terms of quantum information-theoretic quantities. It is natural to wonder whether these conditions might be useful in giving a purely operational definition of an emergent “horizon” in the collective dynamics of even nongravitational bodies.

Whenever a quantum superposition is produced outside an ordinary material body, the long-range gravitational and electromagnetic fields sourced by the superposition become entangled to some degree with the internal degrees of freedom of the material body. However, most material bodies are incredibly inefficient in absorbing long-wavelength gravitational radiation. This renders it generally possible to recover the coherence of the superposition, by adiabatically recombining the branches of the superposition. Biggs and Maldacena [15] have, for example, investigated the extraordinary material properties that would be required for a material body to mimic the decoherence outside a black hole. The extraordinary efficiency of a black hole in decohering its environment is captured by Eq. (32). The information-decoherence trade-off of Eq. (26) reveals that this decoherence can also be understood as a statement of the fact that the distinguishability of the complementary, entangled states on the opposite side of the horizon is also nondecreasing:

$$\frac{d}{dt} \mathcal{I}((\mathcal{N}_{\text{ext}}^{\mathcal{H}^+})^c(t)) \geq 0. \quad (33)$$

The channel information bounds a one-shot state discrimination probability for an optimal observer in Bob’s region. Due to the inefficiency of most material bodies in absorbing long-wavelength gravitational radiation, the degrees of freedom of most material bodies will go into superposition while Alice’s superposition is held open, but will adiabatically recombine and, for the most part, follow the Newtonian field and disentangle themselves from Alice’s spin as she adiabatically closes her superposition. In that case one expects the state-discrimination information available to Bob at intermediate times to be much larger than at late times. Equation (33), along with (31) and (26), can be understood to show that a black hole decoheres its environment as though it were made of *maximally entangling* degrees of freedom.

D. Linear growth of distinguishability long after equilibration

The preceding discussion holds at a high level of generality, for arbitrary states outside black holes. By specializing to a particular case, we gain more explicit quantitative insight into the rate at which black holes decohere their environments—and their interiors harvest information. In the specific case of Alice’s superposition experiment outside a black hole, we may focus on $D(\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+})$ (rather than the entire exterior). Then, under the protocols of the gedanken experiment in which Alice attempts to *minimize* her decoherence, we recover $D(\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+}) = D_{\text{Alice}}$ at late times, and $-\ln(1 - D_{\text{Alice}})$ grows linearly with the time T over which she maintains a superposition [12,24].

Significant recent interest has focused on the interpretation of various quantities (e.g., “complexity,” “volume,” and “action”) relevant to black hole physics that exhibit linear growth long after the equilibration time of the black hole [45–49]. We point out that, in the context of the initial gedanken experiment involving Alice’s superposition, Alice decoheres due to a linearly growing number of soft photons radiated through the horizon [24]. As a result of the preceding discussion, the distinguishability measure

$$-\ln(1 - \mathcal{I}((\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+})^c)) \propto \kappa T \quad (34)$$

therefore exhibits secular, linear growth, where T denotes proper time over which Alice maintains her superposition in the black hole exterior and κ is the surface gravity. Here, $(\mathcal{N}_{\text{Alice}}^{\mathcal{H}^+})^c$ denotes the channel that carries out the joint evolution of Alice and the entire horizon before restricting to the interior. When Alice creates and recombines her superposition sufficiently slowly, an arbitrarily small amount of hard radiation falls into the horizon. The black hole rapidly equilibrates, as the field on the horizon settles

¹⁰The decoherence may not be C^1 , but because it is non-decreasing, it is almost-everywhere differentiable [42].

down to a new stationary state, but with a displaced value of the field on the horizon. The logarithmic distinguishability of interior states, however, continues to grow without bound at least until timescales comparable to the evaporation timescale, due to a linearly growing number of soft gravitons/photons produced due to the superposition in the exterior. From the exterior perspective, it is these soft, entangling gravitons/photons which are responsible for the corresponding decoherence.

From an interior point of view, it is quite natural that the distinguishability of states should increase long after the horizon equilibrates, because the size of the interior region (both in terms of volume and proper time) available for Bob to perform his measurement grows in direct analogy to Fig. 1.

We believe the fact that the decoherence of exterior states, and hence the distinguishability of interior states [by Eq. (26)], increases long after a black hole settles down may be of relevance to the understanding of interior dynamics of black holes in quantum gravity, which continues to be an area of active inquiry [45–50]. Naturally, it is possible—likely, even—that such qualitative statements may receive corrections when working nonperturbatively in G , but the semiclassical description applied here should nevertheless describe the most salient, emergent features available to an observer with bounded resources. It is anticipated that the nonperturbative description of a black hole in quantum gravity may, as viewed from the “exterior,” even more closely mirror the properties of a material body—albeit one with possibly extraordinary material properties [15]. In that case, the “degrees of freedom” of the black hole are expected to become identified with material degrees of freedom. The analogous “material description” for, e.g., a Rindler horizon or de Sitter horizon at finite G_N , however, would seem to be even more difficult to construct.

III. IR-COMPLETE HORIZON ALGEBRAS

While the previous section focused primarily on the operational interpretation of distinguishability measures and the exact trade-off between information and decoherence, this section and the following are concerned with the practical computation of distinguishability measures of the horizon radiation. For concreteness, we will restrict attention to the gedanken experiment considered in Sec. II, in which Alice creates a superposition of coherent states on the horizon.

In the phenomenon of decoherence due to horizons, it is the accumulation of soft radiation on the black hole horizon that is responsible for the loss of coherence of external quantum matter [12,31]. Indeed, in the limit in which Alice never recombines her superposition, the radiation field does not decay at asymptotically late times. Consequently, an infrared-divergent number of soft gravitons or photons fall through on the horizon, resulting in complete decoherence.

This infrared divergence is symptomatic of the fact that a proper treatment of soft quanta requires an enlargement of the theory on the horizon. In our gedanken experiment, however, Alice recombines her superposition at a finite time, and only a finite expected number of soft quanta are produced. Remarkably, the proper treatment of soft quanta will prove useful even in finite time computations that would otherwise be difficult, for reasons we will now explain.

In this section, for simplicity, we will work with the electromagnetic field on the horizon rather than the graviton field. Directly analogous considerations can be used to quantize the graviton field on the horizon, such as is done in [51]. We begin this section by giving a brief summary of the arguments we will use, to give an intuitive understanding of the methods to follow.

On the horizon, the distinguishability to the past of a cut $\mathcal{C}$ is a property of the corresponding algebra $\mathfrak{A}_{\mathcal{C}}$ to the past of that cut. As we will explain, we can get a simple description of the decoherence—as well as any other distinguishability measure that we consider—if the operator $U(A)$ that generates Alice’s radiation is an element of $\mathfrak{A}_{\mathcal{C}}$. This has been dubbed the easy case by Casini *et al.* [52], and such a radiation field is depicted in Fig. 2(a). However, at finite times during Alice’s experiment, $U(A)$ is *not* an element of $\mathfrak{A}_{\mathcal{C}}$ or its commutant, because the radiation “straddles” the cut $\mathcal{C}$ and the operator cannot be split into separate operators in the respective algebras.




FIG. 2. A spatial component A_μ , tangent to the horizon, of the electromagnetic potential is plotted as a function of affine time V parametrizing a horizon generator at angle x^B . (a) A field configuration in the “easy case” of Casini *et al.* (b) A field configuration in the “hard case.” (c) The result of applying a V -independent shift to (b), resulting in a field configuration in the easy case. In particular, the final field configuration can be written as a linear combination of two field configurations: one $\Pi_{\mathcal{C}}A$ supported only to the past of the horizon cut $\mathcal{C}$, and another $\Pi'_{\mathcal{C}}A$ supported only to the future. However, neither of the field configurations $\Pi_{\mathcal{C}}A$ and $\Pi'_{\mathcal{C}}A$ decay asymptotically in time—they exhibit electromagnetic memory [53], and thus do not lie in the standard Hilbert space of scattering theory due to their soft radiation content.

Casini *et al.* dub this the hard case, which is depicted in Fig. 2(b).

However, the radiation field on the horizon can be shifted by a smooth V -independent one-form, whose exact part is a familiar large-gauge transformation [28,54]

$$A_B \rightarrow A_B + h_B(x^A). \quad (35)$$

Using this shift, the radiation field A_A can be made to vanish at the cut $\mathcal{C}$. In this way we will be able to factorize the operator generating Alice's radiation

$$U(A) = U(\Pi_c A) U(\Pi'_c A) U(A|_c) e^{-i\Delta/2}, \quad (36)$$

where $\Pi_c A$ and $\Pi'_c A$ are radiation fields with support entirely to the past and future of $\mathcal{C}$, respectively, $U(A|_c)$ implements the corresponding shift. Here Δ is a c -number corresponding to the “memory” of A , the definition of which will be given shortly in Sec. III A. While this formally factorizes the operator, the radiation fields $\Pi_c A$ and $\Pi'_c A$ do not decay at asymptotically early and late times and contain an infinite number of radiative quanta. The corresponding operators are *not* elements of the standard von Neumann algebra of observables [52].

The purpose of this section is to enlarge the algebra to include such radiation fields and thereby achieve the factorization. This will result in a simple description of the distinguishability measures of Alice's radiation, which will be considered in Sec. IV.

A. Ordinary null quantization on the horizon

The algebra of local, gauge-invariant electromagnetic field observables in the bulk consists of the four-potential $A_a(x)$ smeared with divergence-free test functions $j^a(x)$ whose support decays to zero asymptotically on the spacetime manifold $\mathcal{M}$,

$$A(j) := \int_{\mathcal{M}} d^4x \sqrt{-g} A_a(x) j^a(x). \quad (37)$$

Such observables can be rewritten in terms of observables intrinsic to the horizon $\mathcal{H}^+$ by using the symplectic structure of the classical theory.

Let $\phi_{1,B}(V, x^A)$ and $\phi_{2,B}(V, x^A)$ be two classical four-potential solutions of the source-free Maxwell equations, restricted to the horizon. Working in a Coulomb-like gauge where $A_V = 0$, the symplectic form for the free, radiative degrees of freedom of Maxwell theory on $\mathcal{H}^+$ is

$$\mathcal{W}(\phi_1, \phi_2) = 2 \int_{\mathcal{H}^+} dV d^2x^A \phi_{[1B}(V, x^A) \partial_V \phi_{2]}^B(V, x^A). \quad (38)$$

Classically, this symplectic structure is associated with a Poisson algebra of electric field observables $E(\phi) := \mathcal{W}(A, \phi)$ on the horizon,¹¹ smeared with horizon-restricted four-potential test functions ϕ_B . Here, $E_B = -\partial_V A_B$. The classical Poisson relations are defined by

$$\{E(\phi_1), E(\phi_2)\} = \mathcal{W}(\phi_1, \phi_2) \mathbb{1}. \quad (39)$$

The bulk, quantum observables $A(j)$ can likewise be expressed in terms of an algebra on the horizon by the relation

$$A(j) = \mathcal{W}(\mathcal{E}j, A), \quad (40)$$

where $\mathcal{E}_{ab}$ is the advanced-minus-retarded Green's function and

$$(\mathcal{E}j)_a(x) := \int_{\mathcal{M}} d^4y \sqrt{-g} \mathcal{E}_{ab}(x, y) j^b(y) \quad (41)$$

is the solution corresponding to the source j^a . We use the symplectic form to rewrite the field algebra on the horizon purely in terms of the electric field $E_B = -\partial_V A_B$. To do this, we define the electric field observables,

$$E(s) := 2 \int_{\mathcal{H}^+} d^2x^A dV E_B(V, x^A) s^B(V, x^A), \quad (42)$$

which are related to the local algebra of the exterior via

$$A(j) = \mathcal{W}(\mathcal{E}j, A) = -E(\mathcal{E}j). \quad (43)$$

Given two solutions of the four-potential on the horizon $(s_1)_A$ and $(s_2)_A$ that fall off to zero asymptotically in V (and therefore do not describe field configurations with memory), the commutation relations follow from the symplectic form and are

$$[E(s_1), E(s_2)] = 2i \int_{\mathcal{H}^+} d^2x^A dV s_{[1B}(V, x^A) \partial_V s_{2]}^B(V, x^A) \mathbb{1}. \quad (44)$$

This defines the ordinary algebra of field observables on the horizon.

We will quantize with respect to the Poincaré invariant vacuum. In the case of a general Killing horizon a nearly identical procedure goes through for the Killing-symmetry invariant state [55]. The Poincaré-invariant state ω_Ω is a Gaussian state, fully determined by its vanishing one-point function and its two-point function,

¹¹When s_B decays asymptotically on the horizon, $E(s)$ is nothing more than twice the electric field smeared on the horizon with the test function s_B . For a nondecaying test function ϕ_B this is no longer true, but we retain the $E(\phi)$ notation for simplicity.

$$\omega_\Omega(\mathbf{E}(s_1)^*\mathbf{E}(s_2)) = \frac{-1}{\pi} \int_{\mathcal{H}^+} dV_1 dV_2 d^2x^A \frac{\bar{s}_{1,B}(V_1, x^A) s_2^B(V_2, x^A)}{(V_1 - V_2 - i0^+)^2}. \quad (45)$$

By specifying the quasifree state ω_Ω via the two-point function (45), we can also obtain a $*$ -representation of the horizon field algebra through the Gelfand–Naimark–Segal (GNS) construction.¹² Let $\tilde{\mathcal{A}}$ denote the $*$ -algebra generated by the smeared electric-field operators $\mathbf{E}(s)$ with $s_B(V, x^A)$ smooth and decaying as $V \rightarrow \pm\infty$, with commutator (44). Let $B(\mathcal{F}_0)$ be the set of bounded operators on $\mathcal{F}_0$. Then there exists a cyclic vector $|\Omega\rangle$ and a representation $\pi_\Omega: \tilde{\mathcal{A}} \rightarrow B(\mathcal{F}_0)$ such that

$$\omega_\Omega(\hat{a}) = \langle \Omega | \pi_\Omega(\hat{a}) | \Omega \rangle \quad (46)$$

for all $\hat{a} \in \tilde{\mathcal{A}}$. $\mathcal{F}_0$ is nothing more than the usual bosonic Fock space over the associated one-particle Hilbert space $\mathcal{H}_1$:

$$\mathcal{F}_0 := \mathbb{C} \oplus \bigoplus_{n=1}^{\infty} \mathcal{H}_n. \quad (47)$$

Here $\mathcal{H}_n$ is the n -fold symmetrized tensor product of the one-particle Hilbert space $\mathcal{H}_1$. The subscript “0” in $\mathcal{F}_0$ denotes the fact that *all* states in $\mathcal{F}_0$ must decay at early and late affine times [56]. As mentioned in the Introduction, the relevant infrared property is captured by the “horizon memory”¹³ of ϕ_B ,

$$(\Delta_\phi)_B(x^A) := \lim_{V \rightarrow +\infty} \phi_B(V, x^A) - \lim_{V \rightarrow -\infty} \phi_B(V, x^A) \neq 0, \quad (48)$$

and the decay of states simply is the statement that all states in $\mathcal{F}_0$ have $\Delta_A = 0$. The fact that the standard Fock space comprises only zero-memory configurations can be seen by considering the Fourier transform of the inner product (45). This yields

$$\omega_\Omega(\mathbf{E}(s_1)^*\mathbf{E}(s_2)) = \int_{\mathbb{R}^2} d^2x_A \int_0^\infty \frac{d\omega}{\pi} \omega \bar{s}_{1,B}(\omega, x^A) \tilde{s}_2^B(\omega, x^A), \quad (49)$$

where $\tilde{s}_A$ is the Fourier transform with respect to affine time

¹²In Sec. II we used $\mathcal{A}$ with roman numeral subscripts to denote regions of the spacetime accessible to Alice. From here onward we repurpose $\mathcal{A}$ (possibly with some decorations) to denote a $*$ -algebra and hope the reader will forgive this reuse of notation.

¹³By Maxwell’s equations, the exact part of the memory $\Delta_A(x^A)$ can be related to the difference in “large-gauge charges” associated to the radial component E_r of the electric field between early and late times (see., e.g., [54, 57]). The charges generate the exact shifts $h_B = \mathcal{D}_B \lambda$. These large-gauge charges can be defined on any Killing horizon [58] as well as at null infinity [59].

and $\bar{s}_A$ denotes complex conjugation. From (49) it follows that any one-particle state with memory has infinite norm in the standard Fock space. Indeed, the low frequency behavior of any solution with memory contains a soft pole,

$$\tilde{\phi}_B(\omega, x^A) = \frac{(\Delta_\phi)_B(x^A)}{i\omega} + O(1) \quad (\omega \rightarrow 0). \quad (50)$$

The corresponding norm of the vector $\mathbf{E}(s)|\Omega\rangle = |s\rangle$ for any s_A with memory logarithmically diverges at low frequencies. Thus, $|s\rangle$ is not a vector in $\mathcal{H}_1$. Similarly, the Weyl operator

$$\mathbf{W}(s) := e^{-i\mathbf{E}(s)} \quad (51)$$

satisfying

$$\mathbf{W}(s_1)\mathbf{W}(s_2)^* = \mathbf{W}(s_1 - s_2) e^{i\mathcal{W}(s_1, s_2)/2} \quad (52)$$

is a well-defined operator on $\mathcal{F}_0$ if and only if s_B has zero memory. We can construct a von Neumann algebra of the bounded operator $\mathbf{W}(s)$ on $\mathcal{F}_0$ by taking the weak closure of such operators with respect to $\mathcal{F}_0$. We denote the corresponding “zero memory” von Neumann algebra as $\tilde{\mathfrak{A}}$.

B. Infrared extension of the Weyl algebra

As we have argued, the key insight required to solve the hard case of Fig. 2 is that, in a gauge theory¹⁴ such as QED or general relativity, a coherent state that fails to factorize across the cut is related to one that *does* factorize by a V -independent shift, whose exact part is a large-gauge transformation, as shown in Fig. 2. To realize this, we will include the corresponding memory observables and memory-changing Weyl elements into the horizon algebra.

We seek to extend the algebra of the previous subsection to include objects charged under the memory observables. We face a technical difficulty in doing this, however, because a simple argument reveals that it is impossible to represent these operators on the ordinary Fock space. The argument is as follows.

The classical memory of a field configuration ϕ_B can be obtained by computing its Poisson bracket with so-called “memory observables,” which are in fact already present in our classical Poisson algebra: it consists of nothing more than the special case of electric field observables smeared with a V -independent one-form $h_B(x^A)$. As an observable on phase space,

$$\Delta(h) := -\mathcal{W}(A, h)/2, \quad (53)$$

and its Poisson relations are

¹⁴Similar considerations apply to scalar fields where the relevant transformation is not an asymptotic symmetry but is still a well-defined map on initial data.

$$\begin{aligned}\{E(\phi), \Delta(h)\} &= \frac{1}{2} \int_{\mathcal{H}^+} d^2x^A dV h_B(x^A) \partial_V \phi^B(x^A, V) \\ &= \frac{1}{2} \Delta_\phi(h) \mathbb{1}.\end{aligned}\quad (54)$$

Similar to Eq. (48), $\Delta_\phi(h)$ denotes the *value*, now smeared with h_B , of the memory of a classical solution ϕ_B . Note that $E(\phi)$ is a good classical observable even if ϕ does not decay to zero as $V \rightarrow \pm\infty$ on the horizon, because even in that case it has well-defined Poisson brackets with any other observable.

However, in the quantum theory, the ordinary field algebra, whose construction we reviewed in Sec. III A, excludes any test function that fails to decay to zero as $V \rightarrow \pm\infty$. Therefore, the corresponding quantum horizon memory observable $\Delta(h)$ is defined by

$$[E(s), \Delta(h)] := i\{E(s), \Delta(h)\} \mathbb{1} = \frac{i}{2} \Delta_s(h) \mathbb{1} = 0 \quad (55)$$

as well as the relation that Δ_B commutes with itself. Thus, if we extend the ordinary field algebra to $\tilde{\mathcal{A}}_\Delta$ which, in addition to $E(s)$ also includes memory observables $\Delta(h)$ for all h_B , then these memory observables are central elements of $\tilde{\mathcal{A}}_\Delta$.

For concreteness, we will refer to GNS Hilbert space $\mathcal{F}_0$ of the ordinary field algebra above the Poincaré-invariant vacuum as “the ordinary Fock space,” though the same arguments can be applied to the Hilbert space of the Hartle-Hawking or Unruh vacua of a black hole, or the Bunch-Davies vacuum of de Sitter spacetime, or any other state defined by invariance under the symmetries of a particular Killing horizon. The GNS representation of the field algebra above the Poincaré vacuum is irreducible, so we learn that $\Delta(h)$ is represented as a multiple of the identity operator in any irreducible representation. In particular, on $\mathcal{F}_0$, the memory operator is simply the zero operator—i.e., $\Delta_B = 0$ on $\mathcal{F}_0$ [25,56]. One can straightforwardly obtain memory Fock representations $\mathcal{F}_\Delta$ where all states have memory Δ_B (see, e.g., [25,56] for details on this construction). In other words, on these representations $\Delta(h) = \Delta(h) \mathbb{1}$ on $\mathcal{F}_\Delta$. However, in any definite memory representation, the Fock space $\mathcal{F}_\Delta$ cannot carry a representation of an operator “charged” under the memory observables. Such operators cannot commute with the memory and therefore cannot be represented on any definite-memory Fock space.

This might seem to be an insurmountable obstacle to our desired factorization (36), because the unitaries involved correspond to radiation fields with memory, and are therefore memory-changing. Such operators do not preserve the zero memory character of the standard Fock space. Nevertheless, we seek a representation in which the Poincaré-invariant vacuum state exists as a proper state; we seek to extend the vacuum onto this broader class of

operators. In fact, several authors have confronted this difficulty, and a variety of different methods have been considered to accommodate a description of electromagnetic (or gravitational) memory in Hilbert spaces. For example, a direct sum over memories might be considered [25,60]. If it is desirable that an infinitesimal generator of the unitary group relating memory sectors be represented as an observable, and if the existence of a Poincaré-invariant vacuum state were not required, then even a direct integral over memories might be appropriate [26,27]. Related challenges have arisen in other contexts where unitarily inequivalent sectors of a theory coexist, such as have been addressed by other means in the context of 1 + 1-dimensional conformal field theory by Hollands and Longo [61]. In this work we take an algebraic approach, in which we include precisely the algebra elements needed to answer the physical question at hand, and then see what Hilbert space representation naturally emerges. In Appendix E we elaborate on the different Hilbert space representations one might possibly consider.

Important to us, we are only concerned with Weyl unitaries that change the memory. We now show that the inclusion of these memory-changing coherent operators (i.e., those that produce a coherent state when acting on the vacuum) proceeds without issue.

To include memory-changing operators, we need to include operators which do not commute with the memory. Furthermore, on this enlarged algebra, we require that the vacuum be a well-defined state. Equation (54) suggests that we might try quantizing the classical observable $E(\phi)$ where ϕ has nonzero memory. Such observables do not classically commute with memory and would be a good candidate for our desired observables in the quantum theory. However, such observables are not well-defined operators on the vacuum. Indeed, by (49), the state “ $E(\phi)|\Omega$ ” would have infinite norm.¹⁵ While such operators may be of interest in other circumstances [26,27,51,62], this is not a candidate observable to address these issues that we are concerned with.

However, the Weyl algebra $W(s)$ given by (51) and (52) can be straightforwardly extended to include solutions with memory and, as we explain in the following section, these operators are well-defined on the vacuum. More precisely, we extend the Weyl algebra to include test functions $\phi_B(V, x^A)$ such that $\partial_V \phi_B(V, x^A)$ vanishes at asymptotically late and early affine times, but whose value in fact

¹⁵One might object that “ $E(\phi)$,” corresponding to an unbounded operator, need only be defined on a dense domain, and that perhaps this domain does not include the vacuum when ϕ has memory. In the full algebra, however, the memory-changing unitaries do not form a continuous group in the GNS norm—a fact that we will revisit when constructing unitaries charged under the memory. This failure of continuity rules out the possibility of any such infinitesimal generator if a proper Poincaré-invariant vacuum state is to be retained.

approaches some nontrivial function on the horizon; i.e., $\lim_{V \rightarrow \pm\infty} \phi_B(V, x^A)$ is a nonzero function of x^A along the asymptotic horizon cross section. The commutation relations follow immediately from the ordinary Weyl relations, so that

$$U(\phi_1)U(\phi_2)^* = U(\phi_1 - \phi_2)e^{i\mathcal{W}(\phi_1, \phi_2)/2}. \quad (56)$$

Before we can obtain the Tomita operator it is useful to further extend the algebra to accommodate the decomposition of classical solutions into “components” that factorize across the cut $\mathcal{C}$. This will allow us to factorize coherent operators that do *not* factorize across a given cut $\mathcal{C}$ of the horizon in the original, ordinary Weyl algebra.

Now we will make use of the projectors Π_c and Π'_c introduced earlier in the section, beginning with their proper definition. Given a smooth function $f_B(V, x^A)$ (valued in vectors on the horizon cross section), we now define a projector (i.e., an idempotent map on test functions) Π_c that shifts the function so that it evaluates to zero at V_c and forces the function to zero thereafter. We also define a “complementary” projector Π'_c whose action is reflected about the cut,

$$\begin{aligned} \Pi_c f_B(V, x^A) &:= \Theta(V - V_c) f_B(V_c, x^A) \\ &\quad + \Theta(V_c - V) f_B(V, x^A) - f_B(V_c, x^A), \\ \Pi'_c f_B(V, x^A) &:= \Theta(V_c - V) f_B(V_c, x^A) \\ &\quad + \Theta(V - V_c) f_B(V, x^A) - f_B(V_c, x^A), \end{aligned} \quad (57)$$

where Θ is the Heaviside step.

These projectors are chosen so as to set the function to zero at the cut and to the future (or past), respectively, as required to achieve the factorization of Eq. (36) and Fig. 2. It should be noted that even if f_B is smooth, then the projection $\Pi_c f_B$ is in general not smooth. One might be concerned that this would obstruct the desired extension to the Weyl algebra (56), which depends on the symplectic product between the solutions. The singularity of the projected functions does not spoil the classical Poisson algebra, however, because $\Pi_c \phi_B$ and $\Pi'_c \phi_B$ are both once-differentiable (in the Sobolev sense)

$$\begin{aligned} \partial_V(\Pi_c f_B(V)) &= (\partial_V f_B(V))\Theta(V_c - V), \\ \partial_V(\Pi'_c f_B(V)) &= (\partial_V f_B(V))\Theta(V - V_c), \end{aligned} \quad (58)$$

due to the fact that all intermediate δ -functions cancel identically. This has the result of guaranteeing that the symplectic product of any two projected functions $\mathcal{W}(\Pi_c f_B, \Pi_c g_B)$ is well-defined, as is the symplectic product of a projected function with a smooth one.

With these tools in hand, a straightforward application of the extended Weyl relations of Eq. (56) gives an identity that realizes, at the level of the algebra, the factorization

property we sought in Eq. (36). This identity is the key technical advance of this section, and constitutes the mathematical realization of the observation illustrated in Fig. 2.

Letting $(\phi_c)_B(V, x^A) := \phi_B(V_c, x^A)$ we obtain a factorization identity for general coherent algebra elements

$$U(\phi) = U(\Pi_c \phi) U(\Pi'_c \phi) e^{i2\Delta(\phi_c)} e^{-i\Delta_\phi(\phi_c)/2}. \quad (59)$$

Note that $\Delta_\phi(\phi_c) = \mathcal{W}(\phi_c, \phi)$ where we trivially extend $(\Delta_\phi)_B$ to act on functions on $\mathcal{H}^+$ that are constant in V . As promised at the start of this section, we have in Eq. (59) rewritten a unitary¹⁶ that does not, strictly speaking, factorize across the cut in terms of a product of unitaries belonging to the Weyl algebras on either side of the cut, at the expense of an additional unitary corresponding to a V -independent shift (generated by the memory Δ_B). In doing so we have, in effect, transformed the hard case of Casini *et al.* into the easy case as illustrated in Sec. III.

We also note that

$$[U(\Pi_c f), U(\Pi'_c g)] = 0 \quad (60)$$

for any test functions f and g , justifying our claim that they belong to the commuting algebras on either side of the cut $\mathcal{C}$, respectively.

In addition to the ordinary $*$ -algebra of fields $E(s)$ smeared into smooth, asymptotically decaying functions, and the memory observable, the global horizon algebra includes all $U(\phi)$ for ϕ_B in the extended class of test functions. In addition to smooth test functions ϕ_B with asymptotically-decaying first V -derivatives this includes all projections $\Pi_c \phi_B$ or $\Pi'_c \phi_B$, for each cut of the horizon $\mathcal{C}$ and its associated projectors Π_c and Π'_c .

C. Vacuum state and Hilbert space

An abstract algebra is of little use without a physical state within which to compute expectation values and moments of probability distributions for the observables of that algebra. To make use of our algebra, and to achieve a Hilbert space representation of it via the GNS construction, we extend the Killing-symmetry invariant state of [63]. In the present context of a Rindler horizon, as well as at null infinity, this constitutes the Poincaré-invariant vacuum state.

On the original Weyl algebra,

$$\omega_\Omega(W(s_1)^* W(s_2)) = e^{-\frac{1}{2}\|s_1 - s_2\|^2} e^{i\mathcal{W}(s_1, s_2)/2}. \quad (61)$$

¹⁶Note that this is *not* a unitary operator on the ordinary Fock space. Here we use the word “unitary” to refer to either a unitary element of a $*$ -algebra, or its representative as a unitary operator in a GNS representation of that algebra.

Here, the norm is the Klein-Gordon-like norm on the positive frequency part of the test function

$$\|s\|^2 := \omega_\Omega(E(s)^2). \quad (62)$$

The vacuum state on the extended Weyl algebra is

$$\omega_\Omega(U(\phi_1)^*U(\phi_2)) = \begin{cases} e^{\frac{i}{2}\Delta_{\phi_1}(\phi_1-\phi_2)} & \partial_V(\phi_{2,B} - \phi_{1,B}) = 0 \\ e^{-\frac{1}{2}\|\phi_1-\phi_2\|^2} e^{\frac{i}{2}\Delta(\phi_1,\phi_2)} & \Delta_{\phi_1} = \Delta_{\phi_2} \\ 0 & \Delta_{\phi_1} \neq \Delta_{\phi_2}, \end{cases} \quad (63)$$

where in the first case, the condition $\partial_V(\phi_{2,B} - \phi_{1,B}) = 0$ implies that $\phi_{1,B}$ and $\phi_{2,B}$ differ by a V -independent shift. This state extends onto the memory observable $\Delta(h)$ by the relation

$$[U(\phi), \Delta(h)] = U(\phi)\Delta_\phi(h)/2 \quad (64)$$

and the fact that ω_Ω is an eigenstate of $\Delta(h)$ with eigenvalue zero. This implies that because $\Delta_\phi(h)$ is no longer required to vanish, $\Delta(h)$ is no longer central in the algebra.

In fact, the Weyl element corresponding to the memory observable was already implicit in the extended Weyl algebra. To see this, consider any two solutions $(f_1)_B(V, x^A)$ and $(f_2)_B(V, x^A)$ such that $\partial_V((f_1)_B - (f_2)_B) = 0$, in which case $h_B := (f_1)_B - (f_2)_B$ only depends on the spatial cross sections of the horizon, and $\Delta_{f_1} = \Delta_{f_2}$. Then we find that

$$U(f_1)U(f_2)^* = e^{i\Delta_{f_1}(h)/2} e^{i2\Delta(h)}, \quad (65)$$

so the (exponentiated) memory observable $\Delta(h)$ was, in fact, already present in the extended operator algebra.

At this point one may ask whether the same strategy might be employed to extend the electric field algebra itself “ $E(\phi)$ ” to include classical solutions ϕ_B with memory. Indeed, there is no obstruction defining an *abstract* algebra corresponding to the classical field algebra with memory, in correspondence with the Poisson relations of Eq. (39), to obtain an abstract field algebra that includes test functions ϕ that do not decay asymptotically on the horizon. However, the vacuum state ω_Ω does not extend onto this algebra. This can be seen from the fact that Eq. (45) suffers from an infrared divergence when $s_{1,B}$ and $s_{2,B}$ are both replaced by the same function ϕ_B with memory, i.e., a function such that $\lim_{V \rightarrow \infty} \phi_B(V, x^A) \neq \lim_{V \rightarrow -\infty} \phi_B(V, x^A)$. Thus, although the vacuum ω_Ω has a natural extension to the Weyl algebra given in Eq. (63), it does not extend onto an enlarged field algebra “ $E(\phi)$ ” of classical solutions with memory. Therefore, along with the extended Weyl algebra via Eq. (63), the resulting GNS Hilbert space will carry a representation of the *ordinary*, zero-memory field algebra via Eq. (45).

Let $\mathcal{H}$ be the GNS Hilbert space of the fully extended Horizon algebra $\mathcal{A}$ with respect to ω_Ω where the one-particle inner product is given by

$$\langle s_1 | s_2 \rangle := \omega_\Omega(E(s_1)^*E(s_2)) \quad (66)$$

for any s_1 and s_2 that decay asymptotically, and the inner product on coherent states is

$$\langle U(\phi_1)\Omega | U(\phi_2)\Omega \rangle := \omega_\Omega(U(\phi_1)^*U(\phi_2)). \quad (67)$$

Note that $\Delta(h)$ is no longer central in the algebra of observables, so it is no longer represented by a multiple of the identity. Regarding those elements that do commute with the memory, as elements of the von Neumann algebra $\mathfrak{A}$ of bounded operators on $\mathcal{H}$ one has, for *asymptotically decaying* test functions s ,

$$e^{-iE(s)} = W(s) = U(s). \quad (68)$$

No such relation holds, however, for the elements of the Weyl algebra of operators $U(\phi)$ when ϕ_B has memory. This is because, for any ϕ_B with memory, the one-parameter unitary group $t \mapsto U(t\phi)$ is neither strongly nor weakly continuous in the GNS representation, as can be seen from Eq. (63). This is consistent with the fact that the algebra of observables does not contain a quantum analog of the classical observable $E(\phi)$ for ϕ_B with memory, which, if included in the algebra, would constitute the generator of that unitary group. This topological obstruction can be understood as the reason why the ordinary Fock space of scattering theory fails to incorporate physical states with memory.

The fact that the formal “exponential of a field observable with memory” should be better behaved than the field observable itself has been anticipated in the literature [64]. The question of the existence of the infinitesimal generator of such unitaries is intimately tied to the existence (or lack thereof) of a Poincaré-invariant vacuum state for the resulting algebra. In our context, a Poincaré-invariant vacuum is physically reasonable, and indeed the GNS constructions of Secs. III A and III C guarantee the existence of this vacuum by construction. Other physical contexts, however, may favor other representations [26,27,51,62].

In fact, the GNS Hilbert space we obtain in Sec. III C is the direct sum Hilbert space over memory representations [26,65], and its properties are detailed in Appendix E.¹⁷ The cautious reader may be concerned that this Hilbert space is not separable, and it is sometimes argued that nonseparable

¹⁷While there exist uncountably infinite memory “sectors” labeled by memory, unlike in superselection theory, one cannot restrict to any chosen sector because there exist physical memory-changing operators. This is discussed further in Sec. V A.

Hilbert spaces are not useful. However, as we have already seen, this Hilbert space is the GNS representation of a physically relevant algebra. Furthermore, it will prove very useful insofar as it provides a setting for Tomita-Takesaki theory, which remains well-defined regardless of separability [66].

IV. TOMITA-TAKESAKI THEORY FOR GENERAL COHERENT STATES AND SOFT MODES

A. Standard Tomita-Takesaki theory

We first review standard Tomita-Takesaki theory of the free, radiative fields on the horizon. We recall that the standard von Neumann algebra $\tilde{\mathfrak{A}}_c$ consists of the (weak closure of) the Weyl operators $W(s)$ on the zero memory Fock space $\mathcal{F}_0$ where the operators have support to the past of a cut $\mathcal{C}$. On $\tilde{\mathfrak{A}}_c$ the vacuum state $|\Omega\rangle$ is a cyclic and separating state.¹⁸ Given two cyclic and separating states $|\Psi\rangle, |\Phi\rangle \in \mathcal{F}_0$, one can (densely) define a relative Tomita operator $\hat{S}_{\Psi|\Phi}$ which satisfies

$$\hat{S}_{\Psi \rightarrow \Phi} \hat{a} |\Psi\rangle = \hat{a}^\dagger |\Phi\rangle. \quad (69)$$

Given the relative Tomita operator, one can define a “relative modular operator” $\hat{\Delta}_{\Psi \rightarrow \Phi}$ as

$$\hat{\Delta}_{\Psi \rightarrow \Phi} = \hat{S}_{\Psi \rightarrow \Phi}^\dagger \hat{S}_{\Psi \rightarrow \Phi}. \quad (70)$$

The *modular operator* is defined as the case where both entries of the relative modular operator refer to the same state—i.e., for $|\Psi\rangle$ the modular operator is $\hat{\Delta}_\Psi = \hat{\Delta}_{\Psi \rightarrow \Psi}$. This operator generates a strongly continuous, one-parameter group of automorphisms of any von Neumann algebra $\mathfrak{A}$

$$\hat{a}(t) = \hat{\Delta}_\Psi^{-it} \hat{a} \hat{\Delta}_\Psi^{it} \in \mathfrak{A} \quad \forall \hat{a} \in \mathfrak{A}, t \in \mathbb{R}, \quad (71)$$

where $\hat{a}(t)$ is the one-parameter family of algebra elements which constitute the “modular flow” of $\hat{a}$. The modular operator satisfies the direct analog of the Kubo–Martin–Schwinger (KMS) conditions now with respect to modular flow:

$$\hat{\Delta}_\Psi |\Psi\rangle = |\Psi\rangle, \quad \langle \Psi | \hat{a}(t) \hat{b} | \Psi \rangle = \langle \Psi | \hat{b} \hat{a}(t+i) | \Psi \rangle \quad (72)$$

for any $\hat{a}, \hat{b} \in \mathfrak{A}$. The right-hand side is well-defined because the correlation function $\langle \Psi | \hat{b} \hat{a}(t) | \Psi \rangle$ is analytic on the strip $0 < \Im t < 1$ and is bounded and continuous on the boundary of the strip. As we will review in the following subsections, the modular operator plays a

significant role in the definition of many quantum-information-theoretic quantities. However, an explicit construction of the modular operator for a general state is rarely available in practice.

A notable exception is the vacuum state $|\Omega\rangle$ where on the algebra $\mathfrak{A}_c$, the corresponding operator $\hat{\Delta}_\Omega^{it}$ yields a boost about V_c with rapidity $\xi = 2\pi t$. We will henceforth refer to t as the “boost parameter.” The action of such a boost about V_c on the horizon is

$$\Lambda^{2\pi t}(V, x^A) = (V_c + e^{-2\pi t}(V - V_c), x^A). \quad (73)$$

The action of the vacuum modular operator on the one-particle Hilbert space is determined by its action on the corresponding classical solutions. The classical field $E_A(V, x^B) = -\partial_V A_A(V, x^B)$ transforms under the (pullback Λ_*^ξ of a) boost as

$$\Lambda_*^\xi[E_A(V, x^B)] = e^\xi E_A(V_c + e^{-\xi}(V - V_c), x^B). \quad (74)$$

Another case in which the modular operator is known exactly is due to Casini *et al.* [52] and Longo [67], in the easy case of Sec. III. In that case, the relative modular operator between two coherent states that factorize across the cut as

$$|\Psi\rangle = W(s_\Psi) W(s'_\Psi) |\Omega\rangle, \quad |\Phi\rangle = W(s_\Phi) W(s'_\Phi) |\Omega\rangle \quad (75)$$

is given to be

$$\hat{\Delta}_{\Psi \rightarrow \Phi} = W(s'_\Psi) W(s_\Phi) \hat{\Delta}_\Omega W(s_\Phi)^\dagger W(s'_\Psi)^\dagger. \quad (76)$$

In the hard case, however, no such formula has appeared in the literature, nor has such a formula been derived in the presence of field configurations with memory. In what follows, we provide a formula for the relative Tomita operator between general coherent states.

B. Relative Tomita operator for general coherent states

From here on we will work within the GNS representation of the full algebra¹⁹ defined in Secs. III A and III B. This horizon algebra includes the ordinary zero-memory field algebra $E(s)$ and the corresponding Weyl algebra $W(s)$, as well as the Weyl elements $U(\phi)$ corresponding to coherent radiation with memory, and the memory observable itself $\Delta(h)$. Using this algebra, we will provide the relative Tomita operator for the algebra to the past of the cut $\mathfrak{A}_c$. The full GNS Hilbert space of $\mathfrak{A}$ has several possibly unfamiliar properties relative to the standard Tomita-Takesaki theory we just presented. First, the Hilbert space $\mathcal{H}$ is nonseparable (i.e., it does not admit a countable basis).

¹⁸We say $|\Omega\rangle$ is *cyclic* for $\mathfrak{A}$ if a dense subset of vectors in $\mathcal{H}$ are produced by acting all elements of $\mathfrak{A}$ on $|\Omega\rangle$. We say $|\Omega\rangle$ is *separating* for $\mathfrak{A}$ if $\hat{a}|\Omega\rangle = 0$ implies $\hat{a} = 0$ for all $\hat{a} \in \mathfrak{A}$.

¹⁹The construction also applies to free massive and massless scalar fields, even though the large-gauge interpretation does not apply.

This is no obstruction, however, to the formulation of Tomita-Takesaki theory [66].

To apply the Tomita-Takesaki theory, all that is typically required is the existence of a cyclic and separating state, the definitions of which are given in footnote 18. The algebra $\mathfrak{A}_c$ to the past of the cut will include the standard Weyl algebra $W(s)$ as well as elements of the extended $U(\phi)$ that are supported strictly to the past of the cut. Notably, this does not include $\exp(i\Delta(h))$. With the inclusion of the enlarged Weyl algebra including operators with memory $U(\phi)$, it is straightforward to show that Ω is again cyclic and separating on the algebra $\mathfrak{A}_c$ to the past of the cut.²⁰ Since the state (63) is boost-KMS, the modular operator for

the Poincaré-invariant vacuum is unchanged from its definition on the Poincaré-invariant vacuum of the standard Fock space. Thus, despite the unusual setting, we are equipped with the necessary tools to proceed with the Tomita-Takesaki theory.

We now construct the relative Tomita operator for two general coherent states $|\Psi\rangle = U(\phi^\Psi)|\Omega\rangle$ and $|\Phi\rangle = U(\phi^\Phi)|\Omega\rangle$, where we do not assume that ϕ_B^Ψ and ϕ_B^Φ vanish at the cut, or decay asymptotically. Then, restricting these states to the past of $\mathcal{C}$, we will show that the relative Tomita operator of the state $|\Psi\rangle_c$ relative to $|\Phi\rangle_c$ is

$$\hat{S}_{\Psi \rightarrow \Phi} = e^{-i\Delta_\phi(\phi_c^\Phi)/2} U(\Pi_c \phi^\Psi) U(\Pi'_c \phi^\Phi) \hat{S}_\Omega U(\Pi_c \phi^\Phi)^\dagger U(\Pi'_c \phi^\Psi)^\dagger e^{i\Delta_\phi(\phi_c^\Psi)/2}, \quad (77)$$

where S_Ω is the Tomita operator for the vacuum. Here $\Delta_\phi(h) := \int d^2x^A h^B (\Delta_\phi)_B$.

To justify this, let us verify that $\hat{S}_{\Psi \rightarrow \Phi}$ satisfies the defining property of the relative Tomita operator: $\hat{S}_{\Psi \rightarrow \Phi} \hat{a} |\Psi\rangle = \hat{a}^\dagger |\Phi\rangle$ for any $\hat{a} \in \mathfrak{A}_c$ (including the infrared sector) to the past of the horizon cut $\mathcal{C}$. For any element $\hat{a}$ of $\mathfrak{A}_c$, we have

$$\begin{aligned} \hat{S}_{\Psi \rightarrow \Phi} \hat{a} |\Psi\rangle &= e^{-i\Delta_\phi(\phi_c^\Phi)/2} U(\Pi_c \phi^\Psi) U(\Pi'_c \phi^\Phi) \hat{S}_\Omega U(\Pi_c \phi^\Phi)^\dagger U(\Pi'_c \phi^\Psi)^\dagger e^{i\Delta_\phi(\phi_c^\Psi)/2} \hat{a} |\Psi\rangle \\ &= e^{-i\Delta_\phi(\phi_c^\Phi)/2} U(\Pi_c \phi^\Psi) U(\Pi'_c \phi^\Phi) \hat{S}_\Omega U(\Pi_c \phi^\Phi)^\dagger U(\Pi'_c \phi^\Psi)^\dagger e^{i\Delta_\phi(\phi_c^\Psi)/2} \hat{a} U(\phi^\Psi) |\Omega\rangle. \end{aligned} \quad (78)$$

Next we apply the observation of Sec. III that the easy case differs from the hard case (the case at hand) by a V -independent shift, as illustrated in Fig. 2. We factor the final unitary using the resulting factorization identity (59),

$$= e^{-i\Delta_\phi(\phi_c^\Phi)/2} U(\Pi_c \phi^\Psi) U(\Pi'_c \phi^\Phi) \hat{S}_\Omega U(\Pi_c \phi^\Phi)^\dagger U(\Pi'_c \phi^\Psi)^\dagger e^{i\Delta_\phi(\phi_c^\Psi)/2} \hat{a} U(\Pi_c \phi^\Psi) U(\Pi'_c \phi^\Psi) e^{i2\Delta(\phi_c^\Psi)} e^{-i\Delta_\phi(\phi_c^\Psi)/2} |\Omega\rangle. \quad (79)$$

Then because Ω belongs to the null space of Δ ,

$$= e^{-i\Delta_\phi(\phi_c^\Phi)/2} U(\Pi_c \phi^\Psi) U(\Pi'_c \phi^\Phi) \hat{S}_\Omega U(\Pi_c \phi^\Phi)^\dagger U(\Pi'_c \phi^\Psi)^\dagger \hat{a} U(\Pi_c \phi^\Psi) U(\Pi'_c \phi^\Psi) |\Omega\rangle. \quad (80)$$

Using the fact that coherent operators projected onto opposite sides of the cut commute (60),

$$= e^{-i\Delta_\phi(\phi_c^\Phi)/2} U(\Pi_c \phi^\Psi) U(\Pi'_c \phi^\Phi) \hat{S}_\Omega U(\Pi_c \phi^\Phi)^\dagger \hat{a} U(\Pi_c \phi^\Psi) |\Omega\rangle, \quad (81)$$

and the defining property of $\hat{S}_\Omega$ gives

$$\begin{aligned} &= e^{-i\Delta_\phi(\phi_c^\Phi)/2} U(\Pi_c \phi^\Psi) U(\Pi'_c \phi^\Phi) U(\Pi_c \phi^\Psi)^\dagger \hat{a}^\dagger U(\Pi_c \phi^\Phi) |\Omega\rangle \\ &= e^{-i\Delta_\phi(\phi_c^\Phi)/2} U(\Pi'_c \phi^\Phi) \hat{a}^\dagger U(\Pi_c \phi^\Phi) |\Omega\rangle. \end{aligned} \quad (82)$$

We again use the fact that Ω is in the null space of Δ to write

$$= e^{-i\Delta_\phi(\phi_c^\Phi)/2} U(\Pi'_c \phi^\Phi) \hat{a}^\dagger U(\Pi_c \phi^\Phi) e^{i2\Delta(\phi_c^\Phi)} |\Omega\rangle. \quad (83)$$

Finally, we apply the fact that $U(\Pi'_c \phi^\Phi)$ is in the commutant of $\mathfrak{A}_c$ and again use the factorization identity (59) to obtain

$$\begin{aligned} &= \hat{a}^\dagger U(\phi^\Phi) |\Omega\rangle \\ &= \hat{a}^\dagger |\Phi\rangle, \end{aligned} \quad (84)$$

thereby satisfying the defining property of the relative Tomita operator. Thus, we see that Eq. (77) does, indeed, specify the relative Tomita operator.

With this, we can compute various quantum information theoretic quantities, such as the Holevo fidelity and relative entropy between general coherent states, in both the “easy” and “hard” cases of Casini *et al.*, and also for coherent states of soft radiation.

²⁰ Acting with $U(\phi_\Delta)W(s) \in \mathfrak{A}_c$, with s ranging densely over the zero-memory one-particle space and Δ_B over all memory profiles, produces a dense set in every $\mathcal{F}_\Delta$ and hence in $\mathcal{H}$. The analogous future algebra is cyclic on Ω , hence Ω is separating for $\mathfrak{A}_c$.

The so-called relative modular operator is the modulus of the relative Tomita operator $\hat{\Delta}_{\Psi \rightarrow \Phi} = \hat{S}_{\Psi \rightarrow \Phi}^\dagger \hat{S}_{\Psi \rightarrow \Phi}$, and it is useful in the computation of many information-theoretic quantities. In our case it evaluates to

$$\hat{\Delta}_{\Psi \rightarrow \Phi} = U(\Pi'_c \phi^\Psi + \Pi_c \phi^\Phi) \hat{\Delta}_\Omega U(\Pi_c \phi^\Phi + \Pi'_c \phi^\Psi)^*. \quad (85)$$

In the case that Φ is the vacuum $\Omega|_c$ and Ψ has the vanishing one-point function at $\mathcal{C}$, the projection (including the V -independent shift) implemented by Π'_c acts trivially, and we recover Casini *et al.*'s expression for the relative modular operator in the easy case. We are, however, interested in the fully general case.

Formerly, in the hard case of Casini *et al.*, the lack of a general formulation of the corresponding Tomita-Takesaki theory meant that each new quantum-information theoretic quantity of interest required a new method of computation. Thus, each quantity of interest (e.g., Rényi entropies, etc.) would need to be treated on a case-by-case basis. Now, having suitably extended the algebra and the framework of Tomita-Takesaki theory, a variety of information-theoretic applications can be developed with relative ease. We first return to the Holevo fidelity, which was the original motivation for our foray into Tomita-Takesaki theory, in Sec. II.

With the relative modular operator in hand, we can explicitly evaluate the distinguishability of the states on Bob's side of the horizon sourced by Alice's superposition. We will express this in terms of the Holevo fidelity, the Uhlmann fidelity, the relative entropy, and the Rényi entropies.

C. The Holevo fidelity between coherent states

The (square root of the) Holevo fidelity between two coherent states Ψ and Φ restricted to the past of V_c is given by

$$\sqrt{f_H}(\Psi, \Phi) = \langle \Psi | \hat{\Delta}_{\Psi \rightarrow \Phi}^{1/2} | \Psi \rangle, \quad (86)$$

where the square root of the relative modular operator is

$$\hat{\Delta}_{\Psi \rightarrow \Phi}^{1/2} = U(\Pi'_c \phi^\Psi + \Pi_c \phi^\Phi) \hat{\Delta}_\Omega^{1/2} U(\Pi_c \phi^\Phi + \Pi'_c \phi^\Psi)^*. \quad (87)$$

We will ultimately be interested in the case where $\Phi = \Omega$, so we restrict our attention to that case for simplicity. The calculation follows by exact analogy for a general coherent state Φ .

Because the expectation value $\langle \Psi | \Delta^{it} | \Psi \rangle$ has an analytic continuation to strip $-1 < \Im t < 0$, the conjugation by $\hat{\Delta}_\Omega^{1/2}$ corresponds to a boost with the boost parameter analytically continued to $\xi/(2\pi) = t = -i/2$, which gives the element $\Lambda_{RT}(V, x^A) = (2V_c - V, x^A)$ of the Lorentz group.

The fidelity is symmetric in its arguments, so for convenience we calculate

$$\begin{aligned} \sqrt{f_H}(\Omega, \Psi) &= \langle U(\Pi_c \phi^\Psi)^* \hat{\Delta}_\Omega^{1/2} U(\Pi_c \phi^\Psi) \rangle_\Omega \\ &= \langle U(\Pi_c \phi^\Psi)^* \hat{\Delta}_\Omega^{1/2} U(\Pi_c \phi^\Psi) \hat{\Delta}_\Omega^{-1/2} \rangle_\Omega \\ &= \langle U(\Pi_c \phi^\Psi)^* U(-\Lambda_*^{RT} \Pi_c \phi^\Psi) \rangle_\Omega, \end{aligned} \quad (88)$$

where Λ_*^{RT} is the action of $\Lambda_{-i\pi} = \Lambda_{RT}$ and reflects a test function h^A on the horizon about V_c :

$$\Lambda_*^{RT} h^A(V, x^B) = h^A(2V_c - V, x^B). \quad (89)$$

In the second line of Eq. (88) we have used the fact that Ω is an eigenstate of $\hat{\Delta}_\Omega$ with eigenvalue 1.

Thus, we obtain a simple expression for the fidelity between two coherent states restricted to the past of a cut $\mathcal{C}$,

$$\sqrt{f_H}(\Psi, \Omega) = e^{-\frac{1}{2} \|\Pi_c \phi^\Psi + \Lambda_*^{RT} \Pi_c \phi^\Psi\|^2}. \quad (90)$$

As promised, this formula applies regardless of whether the function ϕ^Ψ is supported at V_c or not. That is, this gives the Holevo fidelity in both the hard and easy cases of Casini *et al.*, and therefore determines the distinguishability of interior states even at intermediate times, while Alice's superposition is held open as discussed in Sec. II.

We previously analyzed an intimately related question [24], in which we determined the optimal protocol for Alice to follow so that she minimizes the decoherence she finds in the exterior of a black hole. There, it was shown that to minimize her decoherence, Alice should not merely close her superposition, but also subsequently reopen it and reclose it according to a specific pattern. By doing this she, in effect, produces a radiation field on the horizon that more closely resembles a vacuum fluctuation and thus reduces the probability of Bob achieving certainty about the state of her superposition. In that paper, it was shown that a reasonably good procedure (though not the optimal one) would be for the reopening and reclosing process to proceed as the *CRT* conjugate of her original process, giving rise to what was there referred to as the “*CRT* purification” on the horizon.

The operation *CRT* on the Hilbert space of a quantum field theory corresponds to the element Λ_{RT} of the Lorentz group about $\mathcal{C}$. The squared magnitude of the inner product of the *CRT* purification with the vacuum is precisely the Holevo fidelity, which we have computed here by much simpler means. As shown in Appendix B the Holevo fidelity approximates the Uhlmann fidelity by

$$1 - \sqrt{f_H}(\rho, \sigma) \leq \sqrt{1 - f_U(\rho, \sigma)} \leq \sqrt{1 - f_H(\rho, \sigma)}, \quad (91)$$

which bounds Bob's probability of success by Eq. (7) and corresponds to the *optimal* purification of [24], and will be the focus of the next section (Sec. IV D).

It is instructive to consider the special case that the coherent state Ψ corresponds to a classical solution s_B that decays asymptotically (zero memory) and also vanishes at the cut. As was shown in [68] it is possible to decompose the exponent of Eq. (90) as

$$\begin{aligned} & -\frac{1}{2} \|\Pi_{c,s} + \Lambda_*^{RT} \Pi_{c,s}\|^2 \\ &= -\frac{1}{2} \|\Pi_{c,s}\|^2 - \frac{1}{2} \|\Lambda_*^{RT} \Pi_{c,s}\|^2 - \Re \omega(E(\Pi_{c,s}), E(\Lambda_*^{RT} \Pi_{c,s})) \\ &= -\|\Pi_{c,s}\|^2 - \Re \omega(E(\Pi_{c,s}), E(\Lambda_*^{RT} \Pi_{c,s})). \end{aligned} \quad (92)$$

The first term counts the number of particles in the coherent state corresponding to $s_B(V, x^A)$ while the final term is the symmetric part of the two-point function, namely [55],

$$\begin{aligned} & \Re \omega(E(\Pi_{c,s}), E(\Lambda_*^{RT} \Pi_{c,s})) \\ &= \frac{-1}{\pi} \int_{\mathbb{R}^2 \times \mathbb{R}^2} dV_1 dV_2 d^2 x^B \\ & \quad \times \text{P.V.} \left[\frac{\Pi_{c,s}^A(V_1, x^B) \Lambda_*^{RT} \Pi_{c,s}^A(V_2, x^B)}{(V_1 - V_2)^2} \right], \end{aligned} \quad (93)$$

where the principal value acts trivially because the test functions have disjoint support. Following [68], then, in the easy case the fidelity cleanly splits into two factors for finite V_c ,

$$\sqrt{f_H}(\Psi|_c, \Omega|_c) = e^{-\|\Pi_{c,s}\|^2} e^{-\Re \omega(E(\Pi_{c,s}), E(\Lambda_*^{RT} \Pi_{c,s}))}. \quad (94)$$

In the case that the radiation remains zero forever after the cut, e.g., in the case that Alice has finished her experiment, the first factor is the fidelity due to the expected number of particles in the state $\langle N \rangle_\Psi = \|s\|^2$, while the second piece is caused by the uncertainty in Bob's ability to measure these particles, due to vacuum fluctuations. Only the latter factor depends on V_c , and its exponent monotonically decays toward zero as V_c increases as can be seen from evaluating (93) and its first V_c derivative.

This implies that as Bob is allowed more time in the interior, he can better distinguish the configuration of the interior field from vacuum fluctuations. This is complementary to the derivation of [24] showing Alice can minimize her decoherence in the exterior by performing her experiment so as to “mimic” a vacuum fluctuation (although the actual improvement in her coherence is minuscule if her superposition has been open for an appreciable time).

In the case in which no time constraints are placed on Bob, then the states on the Rindler horizon would become pure states, in which case there is guaranteed to be some measurement available to Bob that will saturate the bound (7) on his one-shot probability $P_{\text{Bob}}^{\text{certain}}(\mathcal{B})$ to unambiguously distinguish the state.

The infinite time case corresponds to the fidelity along the entire horizon,²¹ which is

$$\sqrt{f_U}(\Psi, \Omega) = e^{-\frac{1}{2} \|\phi^\Psi\|^2} = e^{-\frac{1}{2} \langle \hat{N} \rangle_\Psi}, \quad (95)$$

where $\langle \hat{N} \rangle_\Psi$ is the expected number of particles in the coherent state Ψ . As one would expect, we have recovered the formula of [12–14,31] for the final decoherence $1 - \sqrt{f_U}$ of Alice in terms of the number of entangling particles $\langle \hat{N} \rangle_\Psi$ in the limit that we account for the distinguishability of all entangling degrees of freedom on the horizon. The discrepancy between this formula and the distinguishability of the states according to Bob at finite times as given in Eq. (94) is notable. Evidently, even at very large V_c and thus very small $\Re \omega(E(\Pi_{c,s}), E(\Lambda_*^{RT} \Pi_{c,s})) \simeq 0$, the uncertainty due to vacuum fluctuations poses a significant obstacle to Bob's ability to perform discerning measurements of the field in his region. If, even at very late times, it is possible for Bob to infer with certainty the full future time evolution on the horizon (by, for example, learning by some other means that Alice is finished with her experiment), he gains a significant amount of additional which-path information. This is demonstrated by the exponent of the fidelity passing from $\langle \hat{N} \rangle_\Psi$ to $\langle \hat{N} \rangle_\Psi/2$ in the pure-state case.

D. The Uhlmann fidelity between general coherent states

Recall the definition (6) of the Uhlmann fidelity $f_U(\rho, \sigma)$,

$$f_U(\rho, \sigma) := \sup_{\pi} |\langle \rho | \sigma \rangle_{\pi}|^2. \quad (96)$$

As before, the supremum runs over all Hilbert space representations π of the $*$ -algebra such that $|\rho\rangle$ and $|\sigma\rangle$ are vectors in the corresponding Hilbert space and are purifications of the states ρ and σ , respectively. This supremum over all purifications makes the Uhlmann fidelity very difficult to compute in general. For practical purposes, therefore, the Holevo fidelity is often more useful, because the latter bounds the former from above and below by

$$1 - \sqrt{f_H}(\rho, \sigma) \leq \sqrt{1 - f_U(\rho, \sigma)} \leq \sqrt{1 - f_H(\rho, \sigma)} \quad (97)$$

as shown in Appendix B.

If one restricts to the space of coherent states, then one can achieve a better estimate than the Holevo fidelity. In a related article [24] we, together with J. Kudler-Flam and R. M. Wald, obtained a formula for the optimal purification of the states in $\mathcal{G}$. For simplicity we consider the case of a

²¹Perhaps surprisingly, this expression is *not* identical to the $V_c \rightarrow \infty$ limit of the finite-time, mixed state fidelity of Eq. (94). This discrepancy is due to the fact that $\lim_{V_c \rightarrow \infty} \|\Lambda_*^{RT} \Pi_{c,s}\| \neq \|\lim_{V_c \rightarrow \infty} \Lambda_*^{RT} \Pi_{c,s}\|$.

single coherent state Ψ along with the vacuum Ω . Let s_B be the classical solution corresponding to Ψ . Restricted to a subalgebra, Ref. [24] provides an explicit formula for the optimal purification of Ψ into the space of pure coherent states. That is to say, it was shown that when the purification $|\Psi\rangle$ of Ψ extremizes

$$\max_{|\Psi\rangle \in \mathcal{G}} |\langle \Psi | \Omega \rangle| \quad (98)$$

in the subspace $\mathcal{G}$ of *coherent* states within the Hilbert space, then the optimal coherent state is

$$|\Psi\rangle_{\max} = U(\Pi_c s + R_c \Pi_c s + s_c) |\Omega\rangle, \quad (99)$$

where the optimal recovery R_c acting on some solution $f_B(V, x^A)$ is

$$R_c f_B(V, x^A) := \frac{1}{\pi} \int_{V_c}^{2V_c} dV' \frac{f_B(2V_c - V', x^A)}{V - 2V_c + V'} \sqrt{\frac{V - V_c}{V' - V_c}}. \quad (100)$$

It is natural to expect that $\max_{|\Psi\rangle \in \mathcal{G}} |\langle \Psi | \Omega \rangle| = \sup_{\pi} |\langle \Psi | \Omega \rangle_{\pi}|$, i.e., this is the optimal purification in absolute terms, not just in the space of coherent states above the vacuum. In that case the Uhlmann fidelity becomes

$$f_U(\Psi, \Omega) = e^{-\|\Pi_c s + R_c \Pi_c s\|^2}. \quad (101)$$

Even if it should turn out that an alternative recovery protocol (leveraging purifications outside coherent states) can be found that achieves a larger overlap between the states than the coherent-state protocol of [24], this nevertheless dramatically outperforms the *CRT* purification of the Holevo fidelity for many states. Therefore, this gives an improved bound on the ultimate coherence of Alice's superposition even if it turns out not to be strictly equal to the Uhlmann fidelity.

In the case of field configurations without memory that vanish at the cut, a calculation directly analogous to that of the preceding section again finds that the distinguishability according to Bob at late finite times is controlled by $\sqrt{f_U} \rightarrow \exp(-\langle \hat{N} \rangle_{\Psi})$, whereas in the case that he can infer the entire future state on the horizon we have $\sqrt{f_U} = \exp(-\frac{1}{2} \langle \hat{N} \rangle_{\Psi})$.

E. Relative entropy of general coherent states

The calculation of the relative entropy between a coherent state and the vacuum forms the primary result of Casini *et al.* In the easy case in which the coherent state factorizes across the cut of the initial data surface, their calculation mirrors ours exactly. In the hard case, however, their calculation requires much greater effort. In the

following calculation, the two cases are treated at once, with the hard case being reduced to the level of simple calculation that could previously be employed in the easy case. This highlights the usefulness of the extended Tomita-Takesaki theory for analyzing local quantum information theoretic questions such as the relative entropy to the past of a horizon cut $\mathcal{C}$.

Let $|\Omega\rangle$ and $|\Psi\rangle$ be *CRT* (i.e., “natural cone”) purifications of the states $\Omega|_c$ and $\Psi|_c$, related by $U(\phi)|\Omega\rangle = |\Psi\rangle$. Let $K\phi$ and $\bar{K}\phi$ be the positive and negative frequency parts of ϕ , respectively. Then,

$$\begin{aligned} S_{\Psi|\Omega} &= -\langle \Psi | \ln \Delta_{\Psi \rightarrow \Omega} | \Psi \rangle \\ &= i \frac{d}{dt} \langle \Psi | U(\Pi_c \phi) \Delta_{\Omega}^{it} U(\Pi_c \phi)^* | \Psi \rangle |_{t=0} \\ &= i \frac{d}{dt} \langle \Omega | U(\Pi_c \phi)^* \Delta_{\Omega}^{it} U(\Pi_c \phi) | \Omega \rangle |_{t=0} \\ &= i \frac{d}{dt} \exp(-\|\Pi_c \phi - e^{2\pi t} \Lambda_*^{2\pi t} \Pi_c \phi\|^2 / 2 \\ &\quad + i \mathcal{W}(\Pi_c \phi, e^{2\pi t} \Lambda_*^{2\pi t} \Pi_c \phi) / 2) |_{t=0} \\ &= \frac{-1}{2} \left[i \frac{d}{dt} \|\Pi_c \phi - e^{2\pi t} \Lambda_*^{2\pi t} \Pi_c \phi\|^2 \right. \\ &\quad \left. + \frac{d}{dt} \mathcal{W}(\Pi_c \phi, e^{2\pi t} \Lambda_*^{2\pi t} \Pi_c \phi) \right]_{t=0} \\ &= \frac{-1}{2} \frac{d}{dt} [\mathcal{W}(K \Pi_c \phi - e^{2\pi t} K \Lambda_*^{2\pi t} \Pi_c \phi, \\ &\quad \bar{K} \Pi_c \phi - e^{2\pi t} \bar{K} \Lambda_*^{2\pi t} \Pi_c \phi) \\ &\quad + \mathcal{W}(\Pi_c \phi, e^{2\pi t} \Lambda_*^{2\pi t} \Pi_c \phi)]_{t=0}, \end{aligned} \quad (102)$$

where in going from the second line to the third we have used the fact that $U(\phi)|\Omega\rangle = |\Psi\rangle$ and the identity of Eq. (59). Using the chain rule for the symplectic product

$$\frac{d}{dt} \mathcal{W}(g, h) = \mathcal{W}\left(\frac{d}{dt} g, h\right) + \mathcal{W}\left(g, \frac{d}{dt} h\right), \quad (103)$$

one finds that the contribution from the Klein-Gordon norm vanishes, giving

$$\begin{aligned} S_{\Psi|\Omega} &= -\frac{1}{2} \frac{d}{dt} \mathcal{W}(\Pi_c \phi, e^{2\pi t} \Lambda_*^{2\pi t} \Pi_c \phi) |_{t=0} \\ &= -\frac{1}{2} \mathcal{W}\left(\Pi_c \phi, 2\pi \Lambda_*^{2\pi t} \Pi_c \phi + e^{2\pi t} \frac{d}{dt} \Lambda_*^{2\pi t} \Pi_c \phi\right) \Big|_{t=0} \\ &= -\frac{1}{2} \mathcal{W}\left(\Pi_c \phi, \frac{d}{dt} \Lambda_*^{2\pi t} \Pi_c \phi\right) \Big|_{t=0} \\ &= 2\pi \int_{-\infty}^{V_c} dV d^2 x^A (V_c - V) (\partial_V \phi)^2 + \pi \lim_{V \rightarrow -\infty} (V_c - V) \\ &\quad \times \int_{\mathbb{R}^2} d^2 x^A \partial_V \phi_B(V, x^A) (\phi^B(V, x^A) - \phi^B(V_c, x^A)). \end{aligned} \quad (104)$$

At present we are considering a Hilbert space of functions whose first derivatives decay so that the final term is zero. This leaves²²

$$S_{\Psi|\Omega} = 2\pi \int_{-\infty}^{V_c} dV d^2x^A (V_c - V)(\partial_V \phi_B)^2. \quad (105)$$

This is in agreement with the formula of [52], but its validity is now established even for coherent states corresponding to classical solutions that do not decay at infinity, such as arise in the presence of memory effects.

As Eq. (105) makes clear, the relative entropy $S_{\Psi|\Omega}$ is directly related to the square of the affine derivative of the potential, $(\partial_V \phi_B)^2$. The relative entropy would therefore appear to be much more closely related to the energy of the potential than it is to any “soft” quantity, such as the horizon memory Δ_B .

At the same time, we have seen that fidelity between states that differ by a substantial number of soft photons or gravitons can be made arbitrarily small. Such “soft radiation” therefore carries an enormous amount of which-state information as measured by the fidelity. Evidently the relative entropy (105) and the fidelity capture different aspects of the quantum information content of the radiation.

The gedanken experiment of II helps to illustrate the difference between the two, and in particular, the relative sensitivities of each quantity to the soft radiation produced by Alice. To illustrate the point, we consider two scenarios: (I) Alice opens her superposition, and then holds it open forever, and (II) Alice recombines her superposition shortly after opening it. In both scenarios we analyze the relative entropy and the Uhlmann fidelity of the resulting coherent state restricted to the past of an arbitrary cut V_c .

In scenario (I), suppose that Alice opens her superposition over an affine time interval $V \in (V_0, V_1)$ by keeping one branch of the superposition stationary for all times, and displacing the other branch to a new position where it is then left stationary. Then the relevant, entangling radiation content on the horizon will be determined by the radiation field ϕ_B sourced by the displaced trajectory, relative to the vacuum. Because the integrand of Eq. (105) is only supported when the derivative of the field is nonzero, the relative entropy becomes

$$S_{\Psi|\Omega} = 2\pi \int_{V_0}^{V_1} dV d^2x^A (V_c - V)(\partial_V \phi_B)^2. \quad (106)$$

At late cut times $V_c \gg V_1$ this simplifies further. We write the time at the cut in terms of Alice’s proper time t_c , which is related to V_c by $V_c = e^{\kappa t_c}$, and obtain

²²During the preparation of this manuscript, we became aware that Stefan Hollands and Akihiro Ishibashi have independently derived directly parallel expressions, also allowing for the presence of soft radiation (private communication).

$$S_{\Psi|\Omega} \sim 2\pi e^{\kappa t_c} \int_{V_0}^{V_1} dV d^2x^A (\partial_V \phi_B)^2. \quad (107)$$

From this we learn two things: (i) at late times, the relative entropy grows exponentially in Alice’s proper time t_c , and (ii) the relative entropy is determined by the hard part of the field configuration, in terms of its affine time derivative. In fact, one could replace Alice’s entangling field configuration with a different one, ϕ_{hard} , in which Alice opens and then immediately closes her superposition, and her protocol could be chosen such that

$$\int_{V_0}^{V_1} dV d^2x^A (\partial_V \phi_{\text{hard},B})^2 = \int_{V_0}^{V_1} dV d^2x^A (\partial_V \phi_B)^2. \quad (108)$$

In that case, the relative entropies would be essentially identical at late times, despite the fact that $\phi_{\text{hard},B}$ contains no soft radiation whatsoever, because Alice has closed her superposition. The growth of the relative entropy with time illustrates that even an arbitrarily small amount of energy deposited at early times can give rise to a very large relative entropy as measured at sufficiently late times. More to the point, however, this example illustrates that there are situations in which the relative entropy is not directly sensitive to the soft radiation content of field configurations.

We contrast this with the Uhlmann fidelity, in the same setup in which Alice holds her superposition open. At late times the Uhlmann fidelity

$$f_U(\Psi, \Omega) = e^{-\|\Pi_c s + R_c \Pi_c s\|^2} \quad (109)$$

decays exponentially [24] in t_c as

$$\ln f_U(\Psi, \Omega) \sim -\kappa t_c \quad (\text{I}) \quad (110)$$

for $V_c \gg V_1$. Thus, we see that, while Alice’s superposition is open, the fidelity decays exponentially, due to the linearly increasing number of soft photons in the optimal purification of her superposition [24].

In scenario (II), Alice, instead, closes her superposition shortly after opening it and sources $\phi_{\text{hard},B}$ on the horizon. The resulting fidelity would become essentially t_c independent at late times [24]

$$\ln f_U(\Psi, \Omega) \sim \text{const} \quad (\text{II}) \quad (111)$$

for $V_c \gg V_1$. This can be understood as the result of the fact that the “optimal recovery protocol” for Alice is, in the case that she has already closed her superposition long ago, simply to leave it closed. However, from the discussion above, the relative entropy is still well-approximated by (107) in scenario (II).

Thus, we learn that, in this example, although the relative entropy is not sensitive to the soft radiation content of the

field ϕ_B , the fidelity is sensitive to its soft photon content (or, more precisely, the soft photon content of its optimal purification). This is a reflection of the fact that the fidelity obtains its operational interpretation in terms of one-shot unambiguous state discrimination, and that whether or not Alice holds her superposition open has an enormous influence on the ease with which Bob may distinguish the resulting change in the field from a vacuum fluctuation.²³

For general coherent states Ψ and Φ , the fidelity bounds the relative entropy by [70]

$$e^{S_{\Psi|\Phi}} \geq 1/f_U(\Psi, \Phi). \quad (112)$$

Again letting $\Psi = \Omega|_c$, we see that, in the case under consideration at late times, this is a very weak bound for $\phi_{\text{hard},B}$: the relative entropy grows exponentially in t_c , while the fidelity becomes t_c -independent. When ϕ_B contains soft radiation, however, this bound can become nontrivial, because in that case both quantities grow with t_c . The relative entropy grows because of the time-dependent coefficient in front of the purely hard integral of Eq. (107). The fidelity, on the other hand, is sensitive to the soft photon content of ϕ_B . Since we are considering the case in which Alice never recombines her superposition, in the limit that $t_c \rightarrow \infty$, the corresponding states of soft radiation are precisely orthogonal and have zero fidelity. Likewise, their relative entropy is infinite. If, on the other hand, Alice had opened and then immediately closed her superposition, we would replace ϕ_B with $\phi_{\text{hard},B}$. The fidelity would stay finite, while the relative entropy would nonetheless diverge. Evidently, it is the fidelity that reflects the quantum information content of soft radiation most faithfully.

V. DISCUSSION

In this paper, we have argued that the framework of approximate quantum error correction is the natural language for understanding the gedanken experiments of [10–12,31,71]. From this perspective, the trade-off—indeed, the equivalence of Eq. (26)—between decoherence due to entangling radiation and decoherence due to an eavesdropper learning which-path information is clear. We were then able to explicitly compute the information in soft radiation accessible to an observer behind the horizon. We conclude with a few remarks that may deserve further investigation.

²³The relative entropy, on the other hand, obtains its operational interpretation in a multishot setting [69], where vacuum fluctuations average out over many trials.

A. Connections to scattering theory

Much of the technical content of this article has, for the sake of simplicity, focused on applications to a Rindler horizon. The results, however, apply to horizons more generally, as well as at null infinity of asymptotically flat spacetimes. In fact, much of the related literature until this point has focused *primarily* on soft radiation at null infinity [28,72,73], i.e., in the context of scattering theory. In scattering theory, questions of the coherence of scattering states [74,75] and the quantum information content of soft radiation have been complicated by the ubiquitous production of radiation states with memory in scattering processes. This generic production of radiation with memory is signaled by the ubiquitous appearance of “infrared divergences” in conventional scattering calculations.

The fact that soft modes, in the infinite time limit, become totally orthogonal is the origin of the observation of Carney *et al.* that, at future infinity, an electron is totally decohered by soft radiation [74,75]. This strong degree of entanglement between hard and soft radiation is essential to the proper treatment of quantum information-theoretic questions in scattering theory [74,75]. The present article presents a framework with which to analyze questions such as the distinguishability of soft radiation, and the resulting decoherence of the entangled matter, at finite cuts of null infinity, as this entanglement is still accumulating. These techniques can be applied without appealing to the so-called “inclusive” quantities typically leveraged to remedy infrared divergences. If one is interested, for example, in the amount of entanglement between hard and soft quanta, an inclusive approach gives an incomplete account, because these techniques discard information about the soft radiation present in a given state. The present work sets the stage for more detailed analysis of the quantum information content of soft radiation at null infinity, and the related phenomena of decoherence in scattering theory, at both finite and asymptotic times. Indeed, the latter topic is the focus of upcoming work [76].

B. Decoherence by causal horizons

While we have made a substantial effort in this paper to generalize our previous work [12–14,24], our results still fall short of fully justifying the conclusions of the gedanken experiment presented in the Introduction. In particular, our analysis here has focused on the decoherence induced by Killing horizons in the vicinity of Alice’s laboratory. Nevertheless, the structure of the argument suggests a more general conclusion: namely, a universal bound on the coherence of Alice’s quantum superposition in the presence of any causal horizon.

To make this precise, let Γ be a future- and past-inextendible timelike curve in a spacetime $\mathcal{M}$ representing the worldline of Alice’s laboratory. We assume that the proper time along Γ extends to $-\infty$ as $\tau \rightarrow -\infty$ ensuring

that Alice can, in principle, carry out her experiment for an arbitrarily long duration of proper time T . We do not, however, assume that her proper time similarly extends to $+\infty$ since Γ may terminate at a singularity.

If $D(\Gamma)$ is the domain of communication of Γ (i.e., the causal diamond associated with Γ) then Alice’s worldline possesses a causal horizon if the complement $\mathcal{M} \setminus D(\Gamma)$ is nonempty. The future causal horizon $\mathcal{H}^+$ is the future boundary of $D(\Gamma)$. Suppose that the degrees of freedom beyond the horizon constitute the complete complementary output of Alice’s channel and carry nonzero information, $\mathcal{I}(\mathcal{N}^c) > 0$, about her experiment. The inaccessible region could then, in principle, contain “Bob(s)” who access that information. If, moreover, the complementary outputs form a nested cumulative family as T increases, these hypothetical Bob(s) can accumulate a nondecreasing amount of information. By the arguments presented in the Introduction, Alice is commensurately decohered. Under this nesting assumption, her decoherence is therefore nondecreasing with T , and its effect is *as if* she is becoming entangled with degrees of freedom on the other side of the horizon.

In all previous work [15,77–82] as well as the present paper thus far, $D(\Gamma)$ is the exterior of a stationary black hole, or the de Sitter static patch, and $\mathcal{H}^+$ is a Killing horizon. However, from the above arguments, we see that an analogous decoherence effect can occur in any space-time with a causal horizon—e.g., Alice falling toward a black hole, or a big crunch cosmology. Indeed, the arguments and conclusions of Sec. II can be straightforwardly generalized to this case. Assuming Bob can obtain any information whatsoever about Alice’s state on his side of the horizon, so that $\mathcal{I}(\mathcal{N}^c) > 0$, one again obtains an exact IDT for Alice’s quantum channel $\mathcal{N}$ given by

$$D(\mathcal{N}) = \mathcal{I}(\mathcal{N}^c) > 0, \quad (113)$$

where the positivity follows in any situation in which Bob, on the other side of the horizon, can gain any information whatsoever about Alice’s state. Additionally, under the same nesting assumption, Bob’s accessible information, and hence her decoherence, cannot decrease in time. However, a nonstationary spacetime need not admit preferred notions of “photon number” or global thermality, which played a crucial role in previous analyses of decoherence. Nevertheless, the causal and information-theoretic framework introduced here remains well-defined in this more general setting. We leave a detailed investigation of these issues to future work [83].

ACKNOWLEDGMENTS

We thank Jonah Kudler-Flam for crucial insights and initial collaboration. We also thank Stefan Hollands, Nicholas LaRacuenta, Jonathan Sorce, Robert Wald, and Edward Witten for stimulating discussions. D. L. D.

acknowledges support as a Black Hole Initiative Fellow, as a Fannie and John Hertz Foundation Fellow holding the Barbara Ann Canavan Fellowship, as an Eckhardt Graduate Scholar at the University of Chicago, and as a Long-Term TURIS-IQOQI Fellow of the Austrian Academy of Sciences and the University of Vienna. D. L. D. acknowledges the Julian Schwinger Foundation for financial support at the 2025 Peyresq Spacetime Meeting, where useful discussions influenced this work. This work was performed in part at Aspen Center for Physics, which is supported by National Science Foundation (NSF) Grant No. PHY-2210452. This research was supported in part by NSF Grant No. PHY-2403584 and John Templeton Foundation Grant No. 62845 to the University of Chicago. This publication is funded in part by Gordon and Betty Moore Foundation Grant No. 13526, and John Templeton Foundation Grant No. 63445.

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

APPENDIX A: CHANNEL INFORMATION AND CHANNEL DECOHERENCE

The primary purpose of this Appendix is to elaborate on the definitions, properties, and utility of the channel decoherence and channel information of Sec. II. In Sec. A 1 of this appendix we establish the result of Eq. (26): the exact information-decoherence trade-off of Sec. II. In Sec. A 2 we establish that the channel decoherence recovers a familiar notion of decoherence when applied to a simple experiment, and is therefore a generalization of the familiar notion of “decoherence.” Section A 3 highlights three desirable properties of the definition of channel decoherence: A 3 a. its elimination of the problem of “false decoherence” [32], A 3 b. its independence of a “choice of state,” which guarantees it reflects the intrinsic decoherence of the channel, and A 3 c. its stability under the introduction of an entangled, auxiliary system, due to extremization over all such systems. The validity and utility of the results of Sec. II, and the conclusions drawn from there, rest on the content of this Appendix. Due to its rather technical nature, however, we have deferred this material until this point so as to preserve the clarity of the primary physics points of the paper.

1. Proof of the equivalence between information capacity and complementary decoherence capacity

Quantum channels are the most general maps describing the physical evolution of a system. Here we assume, for the sake of pedagogy, that one has already obtained a Hilbert space representation. For the corresponding definitions

working directly with algebraic states on the C^* algebra of bounded operators, see [36]. One such map is the channel $\mathcal{N}$ taking Alice's subsystem between its initial and final states. If Alice's subsystem comprises operators on the Hilbert space $\mathcal{H}_A$, then

$$\mathcal{N}: \mathcal{S}(\mathcal{H}_A) \rightarrow \mathcal{S}(\mathcal{H}_A)$$

$$\omega_A \mapsto \text{Tr}_{\mathcal{H}^+}[\hat{U}^\dagger(\omega_A \otimes \omega_{\text{QFT}})\hat{U}]. \quad (\text{A1})$$

Here, $\mathcal{S}(\mathcal{H})$ is the space of quantum states on the Hilbert space $\mathcal{H}$, and $\hat{U}$ is the unitary evolution operator on the combined system of Alice's particle and the quantum field. The state ω_{QFT} is fixed, independent of ω_A , and may be taken pure without loss of generality by including any purifying degrees of freedom in $\mathcal{H}_{\mathcal{H}^+}$; the partial trace is over this environment Hilbert space. Given a suitable Hilbert space representation, all quantum channels can be represented by this form, namely tensoring in an auxiliary state, global unitary evolution, and partial trace [36].

There is a unique (up to isometric equivalence) *complementary channel*, $\mathcal{N}^c$, that performs the same evolution, but instead traces over the quantum state of Alice's system,

$$\mathcal{N}^c: \mathcal{S}(\mathcal{H}_A) \rightarrow \mathcal{S}(\mathcal{H}_{\mathcal{H}^+})$$

$$\omega_A \mapsto \text{Tr}_A[\hat{U}^\dagger(\omega_A \otimes \omega_{\text{QFT}})\hat{U}]. \quad (\text{A2})$$

This is the quantum state available to Bob. Unitary evolution constrains the information jointly available through $\mathcal{N}$ and $\mathcal{N}^c$.

We recall the definition, introduced in Sec. II, of the intrinsic quantum decoherence of a quantum channel $\mathcal{N}$.

Definition. (Channel decoherence).

$$D(\mathcal{N}) := 1 - \frac{[\sup_{\mathcal{R}} F(\mathcal{R} \circ \mathcal{N}, \text{Id}) - 1/d]_+}{1 - 1/d}. \quad (\text{A3})$$

Here we have used the channel fidelity F between quantum channels, which is defined in terms of the fidelity f_U as

$$F(\mathcal{N}, \mathcal{M}) := \inf_k \inf_{\omega} f_U((\text{Id}_k \otimes \mathcal{N})(\omega), (\text{Id}_k \otimes \mathcal{M})(\omega)), \quad (\text{A4})$$

where k is the dimension of an ancillary Hilbert space into which the state ω is purified. Throughout this Appendix, the channel fidelity F is defined using the “squared” convention for the Uhlmann fidelity f_U of Eq. (6).

The infimum over k enforces that a perfect recovery channel must not only restore the initial reduced state, but must also restore its correlations with any external system, as implied by Eq. (13).

As a preliminary sign of the usefulness of the channel decoherence, note that when $D(\mathcal{N})$ is very small (say $\epsilon \ll 1$), this can be understood as an indicator of quantum coherence, because one may take ω to be *any* initial state for Alice and the quantum channel $\mathcal{N}$ will be almost perfectly reversible via the decoding channel, the errors being of size ϵ .

In Appendix Sec. A2 we show that the definition of channel decoherence is consistent with the ordinary notion of decoherence of Eq. (8). In Appendix Sec. A3 we remark on various desirable properties of the definition (20) such as its ability to distinguish false decoherence [32] from physical decoherence.

We recall the definition, also introduced in Sec. II, of the information capacity of a channel $\mathcal{N}$.

Definition. (Channel information).

$$\mathcal{I}(\mathcal{N}) := 1 - \frac{[\mathcal{L}(\mathcal{N}) - 1/d]_+}{1 - 1/d}, \quad (\text{A5})$$

where $\mathcal{L}(\mathcal{N})$ is defined as

$$\mathcal{L}(\mathcal{N}) := \sup_{\chi} F(\mathcal{N}, \mathcal{C}_{\chi}), \quad \mathcal{C}_{\chi}(\omega) := \text{Tr}(\omega)\chi. \quad (\text{A6})$$

We will show that this measure of information capacity is tightly related (in fact, equal to) to the complementary channel decoherence. First, we must establish an intermediate result, that the channel loss $\mathcal{L}(\mathcal{N})$ is equivalent to the channel fidelity between $\mathcal{N}$ and a “discard-and-recover” channel:

$$\mathcal{L}(\mathcal{N}) = \sup_{\mathcal{R}'} F(\mathcal{N}, \mathcal{R}' \circ \text{Id}^c). \quad (\text{A7})$$

To see this, note that an example of a channel complementary to Id is Tr , whose target is one-dimensional. In this case the domain of $\mathcal{R}'$ will be one-dimensional, so it must map all inputs onto a fixed state, χ . Thus,

$$\sup_{\mathcal{R}'} F(\mathcal{N}, \mathcal{R}' \circ \text{Id}^c) = \sup_{\chi} F(\mathcal{N}, \mathcal{C}_{\chi}) = \mathcal{L}(\mathcal{N}). \quad (\text{A8})$$

a. Exact information-disturbance trade-off. Having given sensible measures of the decoherence capacity $D(\mathcal{N})$ of a quantum channel as well as its information content $\mathcal{I}(\mathcal{N})$ we apply Bény and Oreshkov's Theorem 1 [20], which is a version of what is known as an information-disturbance trade-off between the information conveyed by the channel $\mathcal{N}^c$ to the environment (in our case, the horizon) and the disturbance the original channel $\mathcal{N}$ exerts on the system under consideration (in our case, Alice's particle). In general, we can also allow for another channel $\mathcal{M}$ to act on the system; since squaring is order preserving on $[0, 1]$, the theorem remains valid for the squared fidelity used here. The *information-disturbance trade-off* is

$$\sup_{\mathcal{R}'} F(\mathcal{N}^c, \mathcal{R}' \circ \mathcal{M}^c) = \sup_{\mathcal{R}^F} F(\mathcal{R} \circ \mathcal{N}, \mathcal{M}). \quad (\text{A9})$$

Specializing to the case where $\mathcal{M} = \text{Id}$, our desired result follows immediately from Eq. (A7):

$$D(\mathcal{N}) = \mathcal{I}(\mathcal{N}^c). \quad (\text{A10})$$

Therefore, the channel decoherence is *equal* to the channel information, as claimed in Eq. (26).

2. The channel decoherence generalizes an ordinary notion of decoherence

In the specific case of Alice's experiment, the generalized definition of decoherence $D(\mathcal{N})$ recovers the familiar notion of decoherence, as we will see here.

By joint concavity of the square root Uhlmann fidelity, the infimum in Eq. (A4) may be taken over pure states on $\mathcal{H}_A \otimes \mathcal{H}_R$.

For simplicity, let us assume that Alice's system consists of a spin. Since any state can be purified using an ancilla of at most the same dimension, the ancillary Hilbert space dimension k need not be any larger than that of Alice's spin $\mathcal{H}_A$, which is the domain of $\mathcal{N}$ in this example. If $\mathcal{H}_R$ denotes the ancillary, or reservoir, Hilbert space, then a general pure state in $\mathcal{H}_A \otimes \mathcal{H}_R$ can be expressed in the z -spin basis as

$$|\rho_{AR}\rangle = \sum_{ij} c_{ij} |i\rangle_A |j\rangle_R \quad \text{with} \quad \sum_{ij} |c_{ij}|^2 = 1, \quad (\text{A11})$$

where $|i\rangle \in \{|\psi_1, \downarrow\rangle, |\psi_2, \uparrow\rangle\}$. If $\mathcal{N}$ implements the effects of decoherence on Alice's superposition at asymptotically late times, then

$$\begin{aligned} & (\text{Id}_R \otimes \mathcal{N})(|\rho_{AR}\rangle\langle\rho_{AR}|) \\ &= \sum_{ijkl} c_{ij} c_{kl}^* |i\rangle\langle k|_A \otimes |j\rangle\langle l|_R \langle\Psi_k|\Psi_i\rangle_{\mathcal{H}^+}. \end{aligned} \quad (\text{A12})$$

The state of Alice's spin and ancilla immediately after Alice completes her experiment is well-approximated by this asymptotic state.

If $\hat{K}_s$ are Kraus operators implementing the recovery channel $\mathcal{R}$, then

$$\begin{aligned} & (\text{Id}_R \otimes \mathcal{R} \circ \mathcal{N})(|\rho_{AR}\rangle\langle\rho_{AR}|) \\ &= \sum_{ijkl} \sum_s c_{ij} c_{kl}^* \hat{K}_s |i\rangle\langle k|_A \hat{K}_s^\dagger \otimes |j\rangle\langle l|_R \langle\Psi_k|\Psi_i\rangle_{\mathcal{H}^+}. \end{aligned} \quad (\text{A13})$$

For an input Schmidt diagonal in the z -spin basis, one obtains

$$\begin{aligned} & (\text{Id}_R \otimes \mathcal{R} \circ \mathcal{N})(|\rho_{AR}\rangle\langle\rho_{AR}|) \\ &= \sum_{ik} \sum_s c_{ii} c_{kk}^* \hat{K}_s |i\rangle\langle k|_A \hat{K}_s^\dagger \\ & \otimes |i\rangle\langle k|_R \langle\Psi_k|\Psi_i\rangle_{\mathcal{H}^+}. \end{aligned} \quad (\text{A14})$$

Here $c_{ii} \in \mathbb{R}$.

$\mathcal{R}$ maps Alice's Hilbert space to itself, so the Kraus matrices $\hat{K}_s$ are *square* matrices. The corresponding Uhlmann fidelity is

$$\begin{aligned} & f_U((\text{Id}_R \otimes \mathcal{R} \circ \mathcal{N})(|\rho_{AR}\rangle\langle\rho_{AR}|), |\rho_{AR}\rangle\langle\rho_{AR}|) \\ &= \langle\rho_{AR}|(\text{Id}_R \otimes \mathcal{R} \circ \mathcal{N})(|\rho_{AR}\rangle\langle\rho_{AR}|)|\rho_{AR}\rangle \\ &= \sum_{i,k,s} c_{ii}^2 c_{kk}^2 \langle i|\hat{K}_s|i\rangle\langle k|\hat{K}_s^\dagger|k\rangle_A \langle\Psi_k|\Psi_i\rangle_{\mathcal{H}^+}. \end{aligned} \quad (\text{A15})$$

Let $p := |\mathcal{P}|$ and absorb the phase of $\mathcal{P}$ into a unitary recovery. Let X, Y, Z denote the Pauli operators in the ordered basis $\{|\psi_1, \downarrow\rangle, |\psi_2, \uparrow\rangle\}$. The induced channel is the qubit dephasing channel

$$\mathcal{N}_p(\rho) = q\rho + (1-q)Z\rho Z, \quad q := \frac{1+p}{2}. \quad (\text{A16})$$

For the identity recovery and any purification $|\rho\rangle$ of an input state ρ ,

$$\begin{aligned} & f_U((\text{Id}_R \otimes \mathcal{N}_p)(|\rho\rangle\langle\rho|), |\rho\rangle\langle\rho|) \\ &= q + (1-q)|\text{Tr}(\rho Z)|^2 \geq q, \end{aligned} \quad (\text{A17})$$

with equality at $\rho = \mathbf{1}/2$. Conversely, if $\{\hat{R}_s\}$ are Kraus operators for an arbitrary recovery and $\hat{R}_s = a_s \mathbf{1} + b_s Z + c_s X + d_s Y$, then the maximally entangled input Φ gives

$$\begin{aligned} & f_U((\text{Id}_R \otimes \mathcal{R} \circ \mathcal{N}_p)(\Phi), \Phi) \\ &= q \sum_s |a_s|^2 + (1-q) \sum_s |b_s|^2 \leq q, \end{aligned} \quad (\text{A18})$$

where trace preservation implies $\sum_s (|a_s|^2 + |b_s|^2 + |c_s|^2 + |d_s|^2) = 1$. Since the channel fidelity minimizes over inputs, the identity recovery is optimal and

$$\begin{aligned} \max_{\mathcal{R}} F(\mathcal{R} \circ \mathcal{N}, \text{Id}) &= \frac{1+|\mathcal{P}|}{2}, \\ D(\mathcal{N}) &= 1 - 2 \left[\frac{1+|\mathcal{P}|}{2} - \frac{1}{2} \right]_+ = 1 - |\mathcal{P}|, \end{aligned} \quad (\text{A19})$$

where $\mathcal{P} := \langle\Psi_1|\Psi_2\rangle_{\mathcal{H}^+}$. This agrees with the ordinary notion of decoherence $D_{\text{Alice}}|_c = 1 - |\langle\Psi_1|\Psi_2\rangle_{\mathcal{H}^+}|$ and of Eq. (8) that is applied in, e.g., [12,31]. Thus, we see that $D(\mathcal{N})$ is a generalization of the ordinary notion of decoherence to general quantum channels. This also agrees

with the long-time limit of $\mathcal{I}_{\text{Bob}} = 1 - \sqrt{f_U}(\Psi_1, \Psi_2) = 1 - |\langle \Psi_1 | \Psi_2 \rangle|$, as anticipated by Eq. (A10).

3. Remarks on channel decoherence

In Sec. II we treat Alice’s exterior experiment (and its horizon generalization) as a quantum channel $\mathcal{N}$ on her laboratory degrees of freedom, with the radiative data on $\mathcal{H}^+ \cup \mathcal{I}^+$ playing the role of an environment. The channel decoherence $D(\mathcal{N})$, defined in Eqs. (A1)–(A4), is the operational measure of *irreversible* loss of coherence that enters the exact information-decoherence trade-off (26) and, through it, the horizon monotonicity statements of Sec. II C. Since this notion is deliberately more refined than the naïve “trace over the field and read off off-diagonals” criterion—especially at finite times, where one encounters Unruh’s false decoherence—we record here a few remarks clarifying the roles of the decoding optimization, the extremization over input states, and the ancillary reference system that stabilizes the definition.

a. Good decoding channels eliminate false decoherence. The definition of channel decoherence quantifies the irreversible decoherence incurred by a quantum channel. Any disturbance of the quantum state which can be reversed purely by manipulating the output of the quantum channel (e.g., by operations performable entirely in Alice’s own laboratory) is automatically corrected by the application of the recovery channel $\mathcal{R}$. In the preceding example, the decoding channel had the effect of correcting the phases of Alice’s superposition.

The inclusion of a decoding channel in the definition of channel decoherence has the effect of eliminating the problem of false decoherence, as discussed in [32]. Suppose, for instance, Alice were to create a superposition and hold it open for some time, but *not* recombine its spatial branches. Formally, the joint state of her particle and the sourced field are

$$\frac{1}{\sqrt{2}}(|\psi_1, \downarrow\rangle \otimes |\Psi_1\rangle + |\psi_2, \uparrow\rangle \otimes |\Psi_2\rangle), \quad (\text{A20})$$

where $\psi_1 \neq \psi_2$. Due to the distinct Coulomb fields of the left and right branches of the superposition, this state is totally decohered.²⁴ This is, however, false decoherence, because it can be trivially reversed by Alice. The optimization over decoding channels in Definition A 1 eliminates such false decoherence: in this case, the optimal decoding channel would involve Alice adiabatically recombining the branches of her spatial superposition, thus recombining the branches of the Coulomb field.

²⁴As given, the state is merely a formal expression, because the “Coulomb subtraction” procedure of [13] cannot be defined for a charge in superposition. However, following [32], the field can be given a small mass to make the analysis precise.

b. Infimum over initial states. A key change in perspective from [12,13] is that we are characterizing decoherence as a property of the quantum channel, independent of the actual state of Alice. Since we will quantify coherence in terms of the recoverability of Alice’s initial state, we must be careful to avoid any definition that trivializes the recovery process. The following example illustrates the type of “trivial case” any sensible definition must discount.

If Alice’s objective is merely to recover Alice’s initial state ω_A after performing her experiment, then regardless of the experiment that she does, there will always exist a decoding channel that achieves this:

$$\begin{aligned} \mathcal{R}: \mathcal{S}(\mathcal{H}_A) &\rightarrow \mathcal{S}(\mathcal{H}_A) \\ \omega'_A &\mapsto \text{Tr}[\omega'_A] \omega_A. \end{aligned} \quad (\text{A21})$$

It is straightforward to check that this linear map is both completely positive and trace preserving. The result is that $f_U(\mathcal{R} \circ \mathcal{N}(\omega_A), \omega_A) = 1$. Of course, we should not conclude that $\mathcal{N}$ cannot decohere Alice’s experiment, because this decoding channel requires Alice to know in advance the state of the system that she is to recover. Had Alice begun with a different state, $\chi_A \neq \omega_A$, then the above decoding channel would be catastrophic. It is for this reason that the channel fidelity of Eq. (A4) takes an infimum over initial states: to ensure that it defines a fidelity intrinsic to the channel, and not merely a characterization of its performance on one highly specialized state. Thus, our decoherence measure provides a measure of unavoidable decoherence due to the channel itself, and not merely due to a poor choice of initial state.

c. The need for the ancillary reference system. What motivates the inclusion of an ancillary identity channel in Eq. (A4)? This serves to stabilize the norm, as the usual one-norm generally may increase even when acting trivially on the k ancillary degrees of freedom—a peculiar consequence of entanglement. To properly characterize the channel decoherence when acting on mixed states, we must consider how the channel acts on the purifications of these states. To illustrate this point, we consider a simple gedanken experiment involving an Einstein–Podolsky–Rosen (EPR) pair.

Suppose that Alice begins with one half of an EPR pair, the other half being held by Bob in his respective laboratory. EPR pairs are valuable resources, allowing for useful protocols such as quantum teleportation, so Alice wants to maintain the coherence of the joint state of her particle and Bob’s so that she retains this valuable resource at later times. Alice performs her Stern–Gerlach experiment while Bob does nothing at all. If the quantum channel irreversibly ruins the state of the EPR pair, even when acting trivially on Bob, we must consider this to be a form of decoherence. Only when an ancillary reference system is included, can the channel decoherence properly account for this form of decoherence, in which

entanglement is degraded even while the reduced density matrix of Alice's subsystem is preserved.

APPENDIX B: THE HOLEVO FIDELITY

The purpose of this Appendix is to establish bounds, from above and below, between the Uhlmann fidelity and the Holevo fidelity between two quantum states. Since both fidelities are used in the article, it is important to know how they relate to one another. The most important result of this Appendix is to show that the Holevo fidelity can be used as an approximant for the Uhlmann fidelity. This is useful in part because the Uhlmann fidelity has a clear operational meaning in terms of unambiguous state discrimination, as reviewed in Sec. II, while the Holevo fidelity is easily computed, using the methods of Tomita-Takesaki theory. The Uhlmann fidelity, on the other hand, is generally quite a difficult quantity to evaluate.

The Holevo fidelity is defined in Tomita-Takesaki theory as

$$\sqrt{f_H}(\rho, \sigma) := \langle \rho | \hat{\Delta}_{\rho \rightarrow \sigma}^{1/2} | \rho \rangle, \quad (\text{B1})$$

where $|\rho\rangle$ is the CRT (i.e., natural cone) purification of the algebraic state ρ . When ρ and σ can be represented as density matrices this reduces to

$$\sqrt{f_H}(\rho, \sigma) = \text{Tr}[\sqrt{\rho}\sqrt{\sigma}]. \quad (\text{B2})$$

The Holevo fidelity is bounded by the Uhlmann fidelity,

$$f_H(\rho, \sigma) \leq f_U(\rho, \sigma), \quad (\text{B3})$$

where equality holds, in particular, when the two states commute.

Both the Uhlmann and Holevo fidelities satisfy respective Fuchs–Van deGraaf-type inequalities [84,85],

$$1 - \sqrt{f}(\rho, \sigma) \leq \frac{1}{2} \|\rho - \sigma\|_1 \leq \sqrt{1 - f(\rho, \sigma)}. \quad (\text{B4})$$

Taken together, Eqs. (B3) and (B4) imply

$$\begin{aligned} 1 - \sqrt{f_U}(\rho, \sigma) &\leq 1 - \sqrt{f_H}(\rho, \sigma) \leq \frac{1}{2} \|\rho - \sigma\|_1 \\ &\leq \sqrt{1 - f_U(\rho, \sigma)} \leq \sqrt{1 - f_H(\rho, \sigma)}, \end{aligned} \quad (\text{B5})$$

which includes a bound between the Uhlmann and Holevo fidelities,

$$1 - \sqrt{f_H}(\rho, \sigma) \leq \sqrt{1 - f_U(\rho, \sigma)} \leq \sqrt{1 - f_H(\rho, \sigma)}. \quad (\text{B6})$$

APPENDIX C: APPROXIMATION OF THE CHANNEL DECOHERENCE

Here we compute bounds on the channel decoherence $D(\mathcal{N})$, or equivalently the channel information, which are more readily computable in general situations than the channel decoherence itself. In particular, given any two states in the range of the complementary channel, such as might arise in Alice's experiment from Bob's point of view, we give a relatively simple lower bound on the decoherence.

The computation relies on a number of bounds and technical arguments, which are not essential to the main arguments of the paper, so we have deferred their presentation until now. Nonetheless, the final bound is important for establishing the computational utility of the channel decoherence, beyond its utility as an operationally useful quantity and its satisfaction of the exact information-decoherence trade-off.

Choosing an idempotent complement Id^c of Id , Bény and Oreshkov prove [20] that

$$\begin{aligned} \frac{1}{2} \sqrt{1 - \sqrt{F(\mathcal{N}^c, \mathcal{N}^c \circ \text{Id}^c)}} &\leq \sqrt{1 - \sqrt{\sup_{\mathcal{R}} F(\mathcal{R} \circ \mathcal{N}, \text{Id})}} \\ &\leq \sqrt{1 - \sqrt{F(\mathcal{N}^c, \mathcal{N}^c \circ \text{Id}^c)}}. \end{aligned} \quad (\text{C1})$$

Because $F(\mathcal{N}^c, \mathcal{N}^c \circ \text{Id}^c) \leq \sup_{\mathcal{R}} F(\mathcal{R} \circ \mathcal{N}, \text{Id})$ and $F \leq \sqrt{F} \leq 1$, this gives rise to a weaker form of information-disturbance trade-off (IDT) reminiscent of the diamond norm IDT of [86],²⁵

$$\begin{aligned} \frac{1}{2} \sqrt{1 - \sqrt{F(\mathcal{N}^c, \mathcal{N}^c \circ \text{Id}^c)}} &\leq \sqrt{1 - \sup_{\mathcal{R}} F(\mathcal{R} \circ \mathcal{N}, \text{Id})} \\ &\leq \sqrt{1 - F(\mathcal{N}^c, \mathcal{N}^c \circ \text{Id}^c)}. \end{aligned} \quad (\text{C2})$$

This can be reexpressed in terms of the channel decoherence,

$$\begin{aligned} &\frac{1}{4} \left[1 - \sqrt{F(\mathcal{N}^c, \mathcal{N}^c \circ \text{Id}^c)} \right] \\ &\leq (1 - 1/d) D(\mathcal{N}) \\ &\leq \min\{1 - 1/d, 1 - F(\mathcal{N}^c, \mathcal{N}^c \circ \text{Id}^c)\}, \end{aligned} \quad (\text{C3})$$

²⁵Explicitly, the diamond-norm IDT is $\frac{1}{4} \inf_{\mathcal{R}} \|\mathcal{R} \circ \mathcal{N} - \text{Id}\|_{\diamond}^2 \leq \|\mathcal{N}^c - \Phi\|_{\diamond} \leq 2 \inf_{\mathcal{R}} \|\mathcal{R} \circ \mathcal{N} - \text{Id}\|_{\diamond}^{1/2}$, where $\|\cdot\|_{\diamond}$ denotes the diamond norm of [86] and the channel Φ takes any input state and replaces it with some fixed state.

where d is the dimension of the input Hilbert space on which $\mathcal{N}$ acts (with $1/d := 0$ for an infinite-dimensional input space). Thus, we see that the channel decoherence can be approximated without solving the supremum-infimum optimization problem required to evaluate $D(\mathcal{N})$ directly.

The preceding inequalities are extremely general in that they hold for arbitrary quantum channels. We now obtain a direct lower bound whenever two states in the range of the complementary channel are known. Let $\omega_i = \mathcal{N}^c(\rho_i)$, $i = 1, 2$, for two input states ρ_1 and ρ_2 . The Bures angle between two states is

$$\theta_B(\rho, \sigma) := \arccos \sqrt{f_U(\rho, \sigma)}. \quad (\text{C4})$$

It obeys the triangle inequality. Thus, for any state χ ,

$$\begin{aligned} \theta_B(\omega_1, \omega_2) &\leq \theta_B(\omega_1, \chi) + \theta_B(\chi, \omega_2), \\ \max_{i=1,2} \theta_B(\omega_i, \chi) &\geq \frac{1}{2} \theta_B(\omega_1, \omega_2). \end{aligned} \quad (\text{C5})$$

Since $\cos \theta$ is decreasing for $0 \leq \theta \leq \pi/2$,

$$\begin{aligned} \min_{i=1,2} f_U(\omega_i, \chi) &\leq \cos^2 \left[\frac{1}{2} \theta_B(\omega_1, \omega_2) \right] \\ &= \frac{1 + \sqrt{f_U(\omega_1, \omega_2)}}{2}. \end{aligned} \quad (\text{C6})$$

The defining infimum in the channel fidelity may be evaluated on either input state ρ_i , so

$$\begin{aligned} \mathcal{L}(\mathcal{N}^c) &\leq \sup_{\chi} \min_{i=1,2} f_U(\omega_i, \chi) \\ &\leq \frac{1 + \sqrt{f_U(\omega_1, \omega_2)}}{2}. \end{aligned} \quad (\text{C7})$$

Using Eq. (26) and the definition of the channel decoherence, for $d > 1$ we obtain

$$\begin{aligned} (1 - 1/d)D(\mathcal{N}) &= 1 - 1/d - [\mathcal{L}(\mathcal{N}^c) - 1/d]_+ \\ &\geq 1 - 1/d - \left[\frac{1 + \sqrt{f_U(\omega_1, \omega_2)}}{2} - 1/d \right]_+ \\ &= \frac{1}{2} [1 - \sqrt{f_U(\omega_1, \omega_2)}]. \end{aligned} \quad (\text{C8})$$

Applying this to the particular case of Sec. II B 3 in which the channel $\mathcal{N}_{\text{Alice}}^+(t)$ models the influence of the horizon on Alice's quantum particle, taking $\omega_1 = \Psi_1|_c$ and $\omega_2 = \Psi_2|_c$ gives

$$\begin{aligned} D(\mathcal{N}) &\geq 1 - \sqrt{f_U(\Psi_1|_c, \Psi_2|_c)} \\ &= D_{\text{Alice}}|_c. \end{aligned} \quad (\text{C9})$$

Because the Holevo fidelity is often easier to compute explicitly than the Uhlmann fidelity, it is useful to apply

Raggio's inequality [87], which in the squared conventions used here is

$$f_H(\rho, \sigma) \leq f_U(\rho, \sigma) \leq \sqrt{f_H(\rho, \sigma)}. \quad (\text{C10})$$

Therefore, for Alice's experiment we obtain the less tight but more readily computable bound

$$D(\mathcal{N}) \geq 1 - f_H(\Psi_1|_c, \Psi_2|_c)^{1/4}. \quad (\text{C11})$$

More generally, for any two states ω_1, ω_2 in the range of $\mathcal{N}^c$,

$$\frac{1}{2} [1 - f_H(\omega_1, \omega_2)^{1/4}] \leq (1 - 1/d)D(\mathcal{N}). \quad (\text{C12})$$

APPENDIX D: RÉNYI ENTROPIES FOR GENERAL COHERENT STATES

In Sec. IV E we apply the infrared extension of Tomita-Takesaki theory to compute the relative entropy of general coherent states. We make no assumption about whether the field configurations vanish at the cut demarcating the restriction of the state, nor do we assume that the field configurations are zero-memory configurations that decay asymptotically. We include this for completeness, but defer it until this point because it does not have direct bearing on the main physics content of the article, which pertains more to the fidelities and the relative entropy.

We proceed by considering any two coherent states, expressed as states on the horizon $\mathcal{H}^+$, $|\Psi\rangle = U(\phi^\Psi)|\Omega\rangle$ and $|\Phi\rangle = U(\phi^\Phi)|\Omega\rangle$, and their restrictions to the past of the cut $\mathcal{C}$, namely $\Psi|_c$ and $\Phi|_c$. Define

$$\phi_B^{\Psi \rightarrow \Phi} := \Pi_c \phi_B^\Psi - \Pi_c \phi_B^\Phi. \quad (\text{D1})$$

Then, for $0 < \alpha < 1$, the order α Petz Rényi entropy of Ψ relative to Φ is

$$\begin{aligned} S_{\Psi|\Phi}^\alpha &= \frac{1}{\alpha - 1} \ln \langle \Psi | \hat{\Delta}_{\Psi \rightarrow \Phi}^{1-\alpha} | \Psi \rangle \\ &= \frac{1}{\alpha - 1} \ln \langle U(\phi^{\Psi \rightarrow \Phi})^* \hat{\Delta}_\Omega^{1-\alpha} U(\phi^{\Psi \rightarrow \Phi}) \rangle_\Omega \\ &= \frac{1}{\alpha - 1} \left[\frac{-1}{2} \| \phi^{\Psi \rightarrow \Phi} - e^{i2\pi\alpha} \Lambda_*^{i2\pi\alpha} \phi^{\Psi \rightarrow \Phi} \|^2 \right. \\ &\quad \left. + \frac{i}{2} e^{i2\pi\alpha} \mathcal{W}(\phi^{\Psi \rightarrow \Phi}, \Lambda_*^{i2\pi\alpha} \phi^{\Psi \rightarrow \Phi}) \right]. \end{aligned} \quad (\text{D2})$$

As was the case for the relative entropy, this formula applies regardless of whether the coherent states lie in the easy or hard cases of Casini *et al.* This formula is consistent with that of [68], with its validity now established in the case that the coherent states contain soft radiation, so the classical solutions do not decay asymptotically.

APPENDIX E: HILBERT SPACE REPRESENTATION

The discussion of Sec. III focused on an extension of the algebra and an accompanying extension of the Poincaré-invariant vacuum state, as an algebraic state. That gave a Hilbert space by the GNS construction, on which our Tomita-Takesaki theory acts. We arrived at this construction in the process of answering some practical questions in quantum field theory, but an equivalent Hilbert space has, in fact, arisen under various guises in the prior literature [25,65]. Here we discuss the features of this Hilbert space representation. The body of the article took an algebraic viewpoint, with the Hilbert space emerging as an immediate by-product (via GNS) after the Poincaré-invariant vacuum extends onto the full horizon algebra. Many of the results in the article can also be understood by understanding the properties of the resulting Hilbert space, so we discuss the definition and properties of that Hilbert space here, to provide a complementary perspective on the construction of Sec. III.

The Hilbert space $\mathcal{H}$ is nothing more than the direct sum over definite-memory Fock spaces. Given a Fock space $\mathcal{F}_0$ of radiation on the horizon, one obtains a unitarily inequivalent Fock space $\mathcal{F}_\Delta$ by the pullback of a state in $\mathcal{F}_0$ along the algebra automorphism [56] $E(\phi) \mapsto E(\phi') + \phi_\Delta(\phi')\mathbf{1}$ (and the corresponding GNS construction of a new Hilbert space), where $(\phi_\Delta)_A(V, x^A)$ is a classical solution with memory, and ϕ'_A is a test function on the horizon.

One can then consider the direct sum over all such so-called “memory representations,”

$$\mathcal{H} := \bigoplus_{\Delta} \mathcal{F}_{\Delta}, \quad (\text{E1})$$

where the sum runs over all allowed memory profiles. Our enlarged algebra is naturally represented on this Hilbert space, and the state ω_Ω corresponds to the Poincaré invariant vacuum state in the $\Delta = 0$ term in the direct sum.

Clearly this Hilbert space is not separable: for each memory profile Δ_B there is a normalized vacuum vector $\Omega_\Delta \in \mathcal{F}_\Delta$, and vectors from different memory sectors are mutually orthogonal. Since the set of allowed profiles $\Delta_B(x^A)$ (e.g., smooth vector fields on $\mathbb{S}^2$) is uncountable, this gives an uncountable orthonormal family in $\mathcal{H}$, which is impossible in a separable Hilbert space. Separability aside, as we have seen, this space is very useful for information-theoretic computations. Due to the direct sum structure, the group action that permutes the memory representations is not continuous in the Hilbert space topology. Thus, the Hilbert space does not carry a representation of the infinitesimal generators of the memory-changing unitaries. This is another way to understand the final discussion of Sec. III C. In some physically realistic scenarios it may be useful to include such infinitesimal generators, and in that case, one may wish to consider a direct integral representation such as those constructed in [26,27,51,60]. In Sec. III C, the difficulty in representing such generators arose due to the inability to define a suitable action of the Poincaré-invariant vacuum on such operators. In a direct integral representation, one sacrifices the existence of a normalizable Poincaré-invariant vacuum, in exchange for the representation of such operators.

-
- [1] S. W. Hawking, Particle creation by black holes, *Commun. Math. Phys.* **43**, 199 (1975).
 - [2] J. D. Bekenstein, Black holes and entropy, *Phys. Rev. D* **7**, 2333 (1973).
 - [3] A. Strominger and C. Vafa, Microscopic origin of the Bekenstein-Hawking entropy, *Phys. Lett. B* **379**, 99 (1996).
 - [4] E. Witten, Gravity and the crossed product, *J. High Energy Phys.* **10** (2022) 008.
 - [5] J. Kudler-Flam, S. Leutheusser, and G. Satishchandran, Generalized black hole entropy is von Neumann entropy, *Phys. Rev. D* **111**, 025013 (2025).
 - [6] D. Harlow, Jerusalem lectures on black holes and quantum information, *Rev. Mod. Phys.* **88**, 015002 (2016).
 - [7] N. Engelhardt and A. C. Wall, Quantum extremal surfaces: Holographic entanglement entropy beyond the classical regime, *J. High Energy Phys.* **01** (2015) 073.
 - [8] R. Bousso, Z. Fisher, S. Leichenauer, and A. C. Wall, A quantum focussing conjecture, *Phys. Rev. D* **93**, 064044 (2016).
 - [9] T. Faulkner, A. Lewkowycz, and J. Maldacena, Quantum corrections to holographic entanglement entropy, *J. High Energy Phys.* **11** (2013) 074.
 - [10] G. Satishchandran, Black Holes, Entanglement and Decoherence, *J. Phys. Conf. Ser.* **3177**, 012002 (2026).
 - [11] D. L. Danielson, Gravitationally Mediated Entanglement and Decoherence, Ph.D. thesis, The University of Chicago, 2025.
 - [12] D. L. Danielson, G. Satishchandran, and R. M. Wald, Black holes decohere quantum superpositions, *Int. J. Mod. Phys. D* **31**, 2241003 (2022).
 - [13] D. L. Danielson, G. Satishchandran, and R. M. Wald, Killing horizons decohere quantum superpositions, *Phys. Rev. D* **108**, 025007 (2023).

- [14] D. L. Danielson, G. Satishchandran, and R. M. Wald, Local description of decoherence of quantum superpositions by black holes and other bodies, *Phys. Rev. D* **111**, 025014 (2025).
- [15] A. Biggs and J. Maldacena, Comparing the decoherence effects due to black holes versus ordinary matter, [arXiv:2405.02227](#).
- [16] A. Uhlmann, The “transition probability” in the state space of a $*$ -algebra, *Rep. Math. Phys.* **9**, 273 (1976).
- [17] J. Kudler-Flam and G. Penington, It costs nothing to teleport information into a black hole, *Int. J. Mod. Phys. D* **34**, 2543002 (2025).
- [18] C. A. Fuchs, Information gain vs state disturbance in quantum theory, [arXiv:quant-ph/9611010](#).
- [19] C. A. Fuchs, Distinguishability and accessible information in quantum theory, Ph.D. thesis, New Mexico University, 1995.
- [20] C. Bény and O. Oreshkov, General conditions for approximate quantum error correction and near-optimal recovery channels, *Phys. Rev. Lett.* **104**, 120501 (2010).
- [21] A. Lewkowycz and J. Maldacena, Generalized gravitational entropy, *J. High Energy Phys.* **08** (2013) 090.
- [22] A. Almheiri, T. Hartman, J. Maldacena, E. Shaghoulian, and A. Tajdini, Replica wormholes and the entropy of hawking radiation, *J. High Energy Phys.* **05** (2020) 013.
- [23] H. Casini, Relative entropy and the Bekenstein bound, *Classical Quantum Gravity* **25**, 205021 (2008).
- [24] D. L. Danielson, J. Kudler-Flam, G. Satishchandran, and R. M. Wald, How to minimize the decoherence caused by black holes, *Phys. Rev. D* **112**, 025012 (2025).
- [25] K. Prabhu, G. Satishchandran, and R. M. Wald, Infrared finite scattering theory in quantum field theory and quantum gravity, *Phys. Rev. D* **106**, 066005 (2022).
- [26] K. Prabhu and G. Satishchandran, Infrared finite scattering theory: Amplitudes and soft theorems, *Phys. Rev. D* **110**, 085022 (2024).
- [27] K. Prabhu and G. Satishchandran, Infrared finite scattering theory: Scattering states and representations of the BMS group, *J. High Energy Phys.* **08** (2024) 055.
- [28] A. Strominger, Lectures on the infrared structure of gravity and gauge theory, [arXiv:1703.05448](#).
- [29] A. Ashtekar, M. Campiglia, and A. Laddha, Null infinity, the BMS group and infrared issues, *Gen. Relativ. Gravit.* **50**, 140 (2018).
- [30] M. Fierz, Über die künstliche umwandlung des protons in ein neutron, *Helv. Phys. Acta* **9**, 245 (1936).
- [31] D. L. Danielson, G. Satishchandran, and R. M. Wald, Gravitationally mediated entanglement: Newtonian field versus gravitons, *Phys. Rev. D* **105**, 086001 (2022).
- [32] W. G. Unruh, False loss of coherence, in *Relativistic Quantum Measurement and Decoherence*, edited by H.-P. Breuer and F. Petruccione (Springer, Berlin, Heidelberg, 2000), pp. 125–140.
- [33] C. W. Helstrom, Quantum detection and estimation theory, *J. Stat. Phys.* **1**, 231 (1969).
- [34] Y. Eldar, A semidefinite programming approach to optimal unambiguous discrimination of quantum states, *IEEE Trans. Inf. Theory* **49**, 446 (2003).
- [35] T. Rudolph, R. W. Spekkens, and P. S. Turner, Unambiguous discrimination of mixed states, *Phys. Rev. A* **68**, 010301 (2003).
- [36] W. F. Stinespring, Positive functions on c^* -algebras, *Proc. Am. Math. Soc.* **6**, 211 (1955).
- [37] D. Kretschmann, D. Schlingemann, and R. F. Werner, A continuity theorem for Stinespring’s dilation, [arXiv:0710.2495](#).
- [38] M. Raginsky, A fidelity measure for quantum channels, *Phys. Lett. A* **290**, 11 (2001).
- [39] D. Puzzuoli and J. Watrous, Ancilla dimension in quantum channel discrimination, *Ann. Henri Poincaré* **18**, 1153 (2017).
- [40] S. Hollands and A. Ranallo, Channel divergences and complexity in algebraic QFT, *Commun. Math. Phys.* **404**, 927 (2023).
- [41] M. Sion, On general minimax theorems, *Pac. J. Math.* **8**, 171 (1958).
- [42] H. Lebesgue, *Leçons Sur L’intégration et la Recherche des Fonctions Primitives* (Gauthier-Villars, Paris, 1904).
- [43] A. Almheiri, N. Engelhardt, D. Marolf, and H. Maxfield, The entropy of bulk quantum fields and the entanglement wedge of an evaporating black hole, *J. High Energy Phys.* **12** (2019) 063.
- [44] G. Penington, Entanglement wedge reconstruction and the information paradox, *J. High Energy Phys.* **09** (2020) 002.
- [45] D. Stanford and L. Susskind, Complexity and shock wave geometries, *Phys. Rev. D* **90**, 126007 (2014).
- [46] A. R. Brown and L. Susskind, Second law of quantum complexity, *Phys. Rev. D* **97**, 086015 (2018).
- [47] A. R. Brown, D. A. Roberts, L. Susskind, B. Swingle, and Y. Zhao, Complexity, action, and black holes, *Phys. Rev. D* **93**, 086006 (2016).
- [48] A. Belin, R. C. Myers, S.-M. Ruan, G. Sárosi, and A. J. Speranza, Does complexity equal anything?, *Phys. Rev. Lett.* **128**, 081602 (2022).
- [49] A. Belin, R. C. Myers, S.-M. Ruan, G. Sárosi, and A. J. Speranza, Complexity equals anything II, *J. High Energy Phys.* **01** (2023) 154 (2023).
- [50] S. Hollands, R. M. Wald, and J. Zahn, Quantum instability of the Cauchy horizon in Reissner–Nordström–deSitter spacetime, *Classical Quantum Gravity* **37**, 115009 (2020).
- [51] J. Kudler-Flam, K. Prabhu, and G. Satishchandran, Vacua and infrared radiation in de Sitter quantum field theory, [arXiv:2503.19957](#).
- [52] H. Casini, S. Grillo, and D. Pontello, Relative entropy for coherent states from Araki formula, *Phys. Rev. D* **99**, 125020 (2019).
- [53] L. Bieri and D. Garfinkle, An electromagnetic analog of gravitational wave memory, *Classical Quantum Gravity* **30**, 195009 (2013).
- [54] S. Pasterski, Asymptotic symmetries and electromagnetic memory, *J. High Energy Phys.* **09** (2017) 154.
- [55] B. S. Kay and R. M. Wald, Theorems on the uniqueness and thermal properties of stationary, nonsingular, quasifree states on space-times with a bifurcate killing horizon, *Phys. Rep.* **207**, 49 (1991).
- [56] A. Ashtekar, *Asymptotic Quantization: Based on 1984 Naples Lectures*, Vol. 2 (Bibliopolis, Naples, 1987).
- [57] T. He, P. Mitra, A. P. Porfyriadis, and A. Strominger, New symmetries of massless QED, *J. High Energy Phys.* **10** (2014) 112.

- [58] V. Iyer and R. M. Wald, Some properties of Noether charge and a proposal for dynamical black hole entropy, *Phys. Rev. D* **50**, 846 (1994).
- [59] D. Kapec, V. Lysov, and A. Strominger, Asymptotic symmetries of massless QED in even dimensions, *Adv. Theor. Math. Phys.* **21**, 1747 (2017).
- [60] A. Herdegen, Semidirect product of CCR and CAR algebras and asymptotic states in quantum electrodynamics, *J. Math. Phys. (N.Y.)* **39**, 1788 (1998).
- [61] S. Hollands and R. Longo, Non-equilibrium thermodynamics and conformal field theory, *Commun. Math. Phys.* **357**, 43 (2017).
- [62] É. É. Flanagan and I. Z. Shehzad, The classical dynamics of gauge theories in the deep infrared, *J. High Energy Phys.* **05** (2023) 185.
- [63] B. S. Kay and R. M. Wald, Theorems on the uniqueness and thermal properties of stationary, nonsingular, quasifree states on space-times with a bifurcate killing horizon, *Phys. Rep.* **207**, 49 (1991).
- [64] E. Himwich, S. A. Narayanan, M. Pate, N. Paul, and A. Strominger, The soft S-matrix in gravity, *J. High Energy Phys.* **09** (2020) 129.
- [65] T. W. B. Kibble, Coherent soft-photon states and infrared divergences. I. Classical currents, *J. Math. Phys. (N.Y.)* **9**, 315 (1968).
- [66] M. Takesaki, *Tomita's Theory of Modular Hilbert Algebras and its Applications*, Lecture Notes in Mathematics (Springer-Verlag, Berlin, 1970).
- [67] R. Longo, Entropy of coherent excitations, *Lett. Math. Phys.* **109**, 2587 (2019).
- [68] M. B. Fröb and L. Sangaletti, Petz–Rényi relative entropy in QFT from modular theory, *Lett. Math. Phys.* **115**, 30 (2025).
- [69] F. Hiai and D. Petz, The proper formula for relative entropy and its asymptotics in quantum probability, *Commun. Math. Phys.* **143**, 99 (1991).
- [70] S. Gudder, J.-P. Marchand, and W. Wyss, Bures distance and relative entropy, *J. Math. Phys. (N.Y.)* **20**, 1963 (1979).
- [71] A. Mari, G. De Palma, and V. Giovannetti, Experiments testing macroscopic quantum superpositions must be slow, *Sci. Rep.* **6**, 22777 (2016).
- [72] H. Bondi, M. J. G. van der Burg, and A. W. K. Metzner, Gravitational waves in general relativity. VII. Waves from axi-symmetric isolated systems, *Proc. R. Soc. A* **269**, 21 (1962).
- [73] R. K. Sachs, Gravitational waves in general relativity. VIII. Waves in asymptotically flat space-time, *Proc. R. Soc. A* **270**, 103 (1962).
- [74] D. Carney, L. Chaurette, D. Neuenfeld, and G. W. Semenoff, Dressed infrared quantum information, *Phys. Rev. D* **97**, 025007 (2018).
- [75] D. Carney, L. Chaurette, D. Neuenfeld, and G. W. Semenoff, Infrared quantum information, *Phys. Rev. Lett.* **119**, 180502 (2017).
- [76] D. L. Danielson, G. Satishchandran, and R. M. Wald, Bulk locality and asymptotic states in quantum field theory and quantum gravity (to be published).
- [77] S. E. Gralla and H. Wei, Decoherence from horizons: General formulation and rotating black holes, *Phys. Rev. D* **109**, 065031 (2024).
- [78] J. Wilson-Gerow, A. Dugad, and Y. Chen, Decoherence by warm horizons, *Phys. Rev. D* **110**, 045002 (2024).
- [79] R. Li, Decoherence of quantum superpositions by Reissner–Nordström black holes, *Phys. Rev. D* **111**, 024040 (2025).
- [80] R. Li, Note on the local calculation of decoherence of quantum superpositions in de Sitter spacetime, *Phys. Rev. D* **111**, 044022 (2025).
- [81] R. Li, Z.-X. Man, and J. Wang, Decoherence of quantum superpositions in near-extremal Reissner–Nordström black holes with quantum gravity corrections, *J. High Energy Phys.* **08** (2025) 079.
- [82] S. Kawamoto, D.-S. Lee, and C.-P. Yeh, Decoherence by black holes via holography, *J. High Energy Phys.* **01** (2026) 154.
- [83] N. Rao and G. Satishchandran, Causal horizons decohere quantum superpositions (to be published).
- [84] C. A. Fuchs and J. van de Graaf, Cryptographic distinguishability measures for quantum-mechanical states, *IEEE Trans. Inf. Theory* **45**, 1216 (1999).
- [85] A. Kholevo, On quasiequivalence of locally normal states, *Theor. Math. Phys.* **13**, 1071 (1972).
- [86] D. Kretschmann, D. Schlingemann, and R. F. Werner, The information-disturbance tradeoff and the continuity of Stinespring's representation, *IEEE Trans. Inf. Theory* **54**, 1708 (2008).
- [87] G. A. Raggio, Comparison of Uhlmann's transition probability with the one induced by the natural positive cone of von neumann algebras in standard form, *Lett. Math. Phys.* **6**, 233 (1982).