

MIT Open Access Articles

Cryptographic Censorship

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Engelhardt, N., Folkestad, Å., Levine, A. et al. Cryptographic Censorship. J. High Energ. Phys. 2025, 122 (2025).

As Published: [https://doi.org/10.1007/JHEP01\(2025\)122](https://doi.org/10.1007/JHEP01(2025)122)

Publisher: Springer Berlin Heidelberg

Persistent URL: <https://hdl.handle.net/1721.1/158276>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution



Cryptographic Censorship

Netta Engelhardt ^a, Åsmund Folkestad ^a, Adam Levine ^a, Evita Verheijden ^{a,b}
and Lisa Yang ^{a,c}

^aCenter for Theoretical Physics, Massachusetts Institute of Technology,
Cambridge, MA 02139, U.S.A.

^bBlack Hole Initiative at Harvard University,
20 Garden Street, Cambridge, MA 02138, U.S.A.

^cComputer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology,
Cambridge, MA 02139, U.S.A.

E-mail: engeln@mit.edu, afolkest@mit.edu, arlevine@mit.edu,
evitamhv@mit.edu, lisayang@mit.edu

ABSTRACT: We formulate and take two large strides towards proving a quantum version of the weak cosmic censorship conjecture. We first prove “Cryptographic Censorship”: a theorem showing that when the time evolution operator of a holographic CFT is approximately pseudorandom (or Haar random) on some code subspace, then there must be an event horizon in the corresponding bulk dual. This result provides a general condition that guarantees (in finite time) event horizon formation, with minimal assumptions about the global spacetime structure. Our theorem relies on an extension of a recent quantum learning no-go theorem and is proved using new techniques of pseudorandom measure concentration. To apply this result to cosmic censorship, we separate singularities into classical, semi-Planckian, and Planckian types. We illustrate that classical and semi-Planckian singularities are compatible with approximately pseudorandom CFT time evolution; thus, if such singularities are indeed approximately pseudorandom, by Cryptographic Censorship, they cannot exist in the absence of event horizons. This result provides a sufficient condition guaranteeing that seminal holographic results on quantum chaos and thermalization, whose general applicability relies on typicality of horizons, will not be invalidated by the formation of naked singularities in AdS/CFT.

KEYWORDS: AdS-CFT Correspondence, Black Holes, Spacetime Singularities

ARXIV EPRINT: [2402.03425](https://arxiv.org/abs/2402.03425)

Contents

1	Introduction	1
1.1	Preliminaries	10
2	Warmup: QESs lie behind horizons	13
3	Cryptographic Censorship	16
3.1	Review: pseudorandomness and complexity of learning	18
3.2	Complexity of learning in holography	20
3.3	Pseudorandomness implies horizons	23
3.4	Possible subspaces with pseudorandom time evolution	28
4	From Cryptographic to cosmic censorship	30
4.1	The size of singularities	31
4.2	Accommodating gravitational pseudorandomness	35
4.3	Quantum cosmic censorship	39
5	Discussion	40
A	Proof of Theorem 3	44
B	Proof of Corollary 3.1	51

1 Introduction

Curvature singularities are ubiquitous in General Relativity: the evolution of smooth initial data generically results in such singularities [1–3]. Observationally, however, evidence for singularity formation from generic gravitational processes (such as matter collapse) remains at best indirect. This apparent absence of directly-observable singularities is often explained by the general expectation that the formation of horizons is likewise generic: so singularities are usually cloaked behind event horizons. This hypothesis — that singularities resulting from initial data evolution should be generically hidden behind horizons — is known as the Weak Cosmic Censorship Conjecture (WCCC) [1]. It is a cornerstone assumption of almost all results in modern gravitational physics, from classical theorems such as the Hawking area law [4], to key results from the general “it from qubit” program. In particular, the expectation that CFT thermalization corresponds to black hole equilibration, and more generally that black hole statistics are reflected in CFT thermal states, crucially depends on the assumption that states with horizons are typical in some measure.¹ If the WCCC were false, naked singularities could be as typical as black holes; our understanding of quantum chaos, eigenstate thermalization, and black hole information in holography would be called into question.

¹This is not identical to the firewall sense of typicality of [5]. See footnote 6 on page 5.

Surprisingly, it turns out that naked singularities do exist in solutions to the Einstein equation satisfying various energy conditions and regular initial data, and these states do *not* constitute a measure zero set. There is by now a large repertoire of counterexamples to the WCCC: in more than four spacetime dimensions, the studies of Gregory-Laflamme type instabilities of [6] have illustrated the existence of “pinch-off” singularities as the end-stage of certain black holes (e.g., the black string [7]) — see [8] for a review; more recently, [9–12] constructed extended naked singularities in asymptotically AdS₄ spacetimes with various matter profiles; see also [13]. A similar mechanism was used by [14] to construct a likely counterexample in the four-dimensional asymptotically flat setting.

It thus appears that horizons and singularities are not related in any trivial or obvious way: curvature singularities may form outside (or in the absence) of any horizons, and horizons may form without singularities behind them [15]. There is no extant formulation of WCCC that is (1) not violated and (2) sufficient to guarantee the consistency of seminal results in AdS/CFT such as the universality of the saturation of the chaos bound and the bulk manifestation of the CFT thermal partition function.

Some progress towards understanding cosmic censorship in quantum gravity was made in [9–11], where it was noted that many (though not quite all) of the counterexamples also violate the Weak Gravity Conjecture [16], suggesting the possibility that the WCCC in four dimensions is valid in the landscape but violated in the swampland. Further evidence in favor of this conclusion was presented in [13], which constructed large sets of AdS₄ initial data satisfying the gravitational constraints and the null energy condition whose future development must include naked singularities, but these theories violated properties required of theories admitting a UV completion via holography [17]. This hypothesis would suggest that while classical General Relativity may violate the WCCC, the semiclassical limit of quantum gravity does not.

However, this naïve “swampland censorship” scenario is clearly incomplete. In higher dimensions, the Gregory-Laflamme instability is not excluded by any known swampland constraints; moreover the inclusion of quantum corrections is expected to lead to black hole evaporation, which also violates the WCCC. However, the absence of naked singularities in the observable universe together with the top down arguments from quantum gravity suggest that *some* version of WCCC remains valid: intuitively, a commonality between the Gregory-Laflamme instability and the evaporating black hole is that there is a heuristic sense, which in this article we will make precise, that the naked part of the singularities is “small”. See figure 1. That is, naked singularities can exist, but they must only be visible to asymptotic observers for a “short” period of time. Furthermore, they can only exist in spacetimes that have horizons — so that the bulk of the singularity can remain hidden.

A formulation of such a “quantum gravitational” version of cosmic censorship — which we propose later in this article — and its subsequent proof would require as *sine qua non* both (1) a condition that can guarantee the formation of horizons in broad generality, and (2) a definition of a “small” naked singularity.

Point (1) is a notoriously difficult open problem. Event horizons are by definition teleological: an event horizon is defined with reference to future infinity. Consequently, very few extant results can guarantee the existence of such a global construct even under very




Figure 1. The evaporation point of a black hole is, in some sense that we make precise in section 4.1, “small”.

restrictive assumptions (and none guarantee that a given event horizon will successfully hide the full extent of a curvature singularity). See [18–20] for a series of failed attempts to construct criteria to this effect. Addressing point (1) would also have the potential to shed light on a number of other research programs, such as thermalization in AdS/CFT (see e.g. [21] for a review) and the firewall problem [5, 22, 23]. Point (2) has seen various approaches (see [24]), but none of those have successfully made contact with the generic existence of horizons, since (1) has thus far remained unresolved.

Here we accomplish both (1) and (2) in the context of AdS/CFT: we give a general guarantee for the existence of event horizons at finite time, and criteria for the size of a singularity that make contact with the guarantee for event horizons. In particular, we classify singularities into three categories: classical, semi-Planckian and Planck-sized, and we motivate an assumption on the behavior of singularities that would guarantee the existence of horizons at finite time — via point (1) — in spacetimes containing such classical or semi-Planckian singularities. In the rest of this section, we give a rough outline of the central ingredients of (1) and (2).

A general guarantee for horizon formation.

Under the standard assumption of global hyperbolicity,² which guarantees various consistency results such as unitary invariance of the von Neumann entropy and entanglement wedge nesting, there are various sufficient conditions that may guarantee existence of an event horizon, such as exponential reconstruction complexity or chaos. However, global hyperbolicity *implies* the WCCC as a consequence: almost all results ensuring consistency of the AdS/CFT dictionary assume cosmic censorship — even though it is known that cosmic censorship can be violated in perfectly reasonable AdS/CFT setups (e.g., the evaporating AdS black hole). This strongly suggests that while the WCCC is sufficient for consistency, it is not in fact necessary. We will argue for a more minimalistic form of censorship, which *is* necessary;³ it remains to be seen whether this version is also sufficient.

²We really mean AdS hyperbolicity here, which is essentially global hyperbolicity when the asymptotic boundary is included. See [25] for the precise definition.

³We hope that this sentence is not taken out of context.

Our goal is to find a cosmic censorship conjecture that follows directly from consistency of the AdS/CFT dictionary. Clearly, an essential part of a holographic formulation of cosmic censorship is the holographic dual of an event horizon: we are looking for the boundary dual to the process of bulk horizon formation. This specifically means that we are interested in the existence of a region outside of the causal wedge of the entire asymptotic boundary. The separation between the causal (or simple [26–28]) wedge and the entanglement wedge has been previously linked in [28–30] to the separation between polynomial and exponential quantum complexity classes in the dual theory (see also [31] and footnote 9). Our approach to finding a guarantee for the formation of event horizons will thus focus on reconstruction restrictions by complexity class.

More concretely, we ask when there is a guarantee for the existence of some operator Q in the large- N limit that cannot be reconstructed from knowledge of the causal wedge alone, given that the latter can be reconstructed efficiently, e.g., by jumping into the bulk, making a measurement, and sailing back towards the boundary.⁴

In [32], two of us provided such a guarantee at large but finite N under certain assumptions. The full statement of the relevant theorem will be reviewed at length in section 3.1; here we provide a heuristic statement only, starting with the most significant assumption: that the action of the fundamental time evolution operator U of the system is well modeled on our code subspace by a *pseudorandom* unitary [33] (we use “fundamental” to refer to the full quantum description of the gravitational system, as in [30]). In AdS/CFT, this assumption is valid whenever the time evolution of the CFT on a particular choice of code subspace is well approximated by a pseudorandom unitary at any finite value of N ;⁵ this is not dissimilar from the assumption of e.g. [34] that the fundamental time evolution of a black hole is well modeled by a two-design or as in [30, 35] by a (pseudo)random unitary. Pseudorandom unitaries are operators drawn from an ensemble $\{U_k\}$ that are indistinguishable from Haar random unitaries by any efficient quantum algorithm. Our results below also hold for Haar random unitaries, but the advantage of pseudorandom unitaries over Haar random unitaries is that by definition they can be constructed by an efficient process. While existence of pseudorandom unitaries has not been explicitly demonstrated yet, candidate constructions have appeared in, e.g., [33]. We will not make use of any such explicit constructions.

Under this first assumption of pseudorandom unitary (PRU) dynamics, the theorem of [32] considers any efficient quantum algorithm \mathcal{A} that:

- (a) Prepares polynomially complex states $|\varphi\rangle$ and has access to the action of the time evolution operator U on the expectation values of polynomially complex observables in such states;
- (b) Uses the information in (a) to attempt to learn U (or a circuit representation of U), the pseudorandom time evolution operator, thereby producing some operator \hat{U} ;

⁴By an efficient reconstruction, we mean one that can be implemented by a quantum circuit whose size is polynomially bounded in the logarithm of the dimension of the code subspace under consideration.

⁵We will eventually work at strictly infinite N , but then we will work only with time-evolved states rather than the unitary operator that generates that time evolution, since the aforementioned unitary operator is not well-defined in the strict infinite N limit.

- (c) After completing (b), is handed the actual state of the system $|\psi\rangle$ and is asked to predict U acting on $|\psi\rangle$ using the outcome of (b).

The theorem states that on average — and as we will show in section 3.2, for typical pseudorandom states and unitaries⁶ — any efficient algorithm will fail to produce an accurate guess for U acting on $|\psi\rangle$.⁷ That is, an efficient algorithm’s best guess \hat{U} for U will satisfy:

$$|\langle\psi|\hat{U}^\dagger U|\psi\rangle|^2 \leq 1 - \alpha, \tag{1.1}$$

where α is an $O(1)$ constant, independent of the size of the Hilbert space in which $|\psi\rangle$ lives. This result in particular guarantees the existence of a distinguishing operator Q whose expectation value differs between $\hat{U}|\psi\rangle$ and $U|\psi\rangle$ by at least 2α .

We now apply this result to our gravitational setup using the fact that reconstruction is a particular type of learning. We will consider code subspaces $\mathcal{H}_{\text{code}}$ of dimension $e^{O(G_N^a)}$, where $-1 \leq a < 0$, which do not necessarily have horizons in typical states. For example, a microcanonical window of width $O(N)$; another example is a set of states where the leading order contribution to the geometry is the same for all of the states at least for a time of $O(G_N^0)$. We will show that it is easy to prepare candidates for such code subspaces without obviously incurring horizons. We next assume that the time evolution is *gravitationally* pseudorandom: we will define this term more precisely in section 3.2, but roughly speaking it denotes an $O(G_N^a)$ amount of pseudorandomness; we find that this amount of pseudorandomness is generally compatible with an $\mathcal{H}_{\text{code}}$ of size $e^{O(G_N^a)}$, which are the code subspaces that we consider. Next, we argue for the existence of an efficient algorithm that, as part of the learning phase, can reconstruct the causal wedge of any efficiently-preparable state in this code subspace. Under certain technical assumptions, there is an algorithm that, given access to some CFT data, is able to output a guess \hat{U} for the time evolution of the initial state that accurately reproduces the expectation values of all operators in the causal wedge. However, the theorem of [32] guarantees that the fidelity between the predicted time evolution and the actual time evolved state is low, which implies that there exists a distinguishing operator Q whose expectation values differ at $O(1)$ between the actual time evolved state and the predicted time evolution of the state.

Since there exists a learning algorithm that can accurately reconstruct the time evolution of all operators in the causal wedge but cannot accurately reconstruct the expectation value of some operator Q , which we show exists in the large- N limit, we immediately find that the dual to Q lies outside of the causal wedge. Thus, the boundary of the causal wedge is nontrivial and an event horizon must exist.

⁶Caution: this is not the typical use of ‘typical’. Here we will use the qualifier typical specifically for states in the code subspace to which measure concentration applies; this depends on the choice of distribution of states under consideration. For example, we may consider “typical Haar random states” or “typical pseudorandom states”. We will always note the distribution under discussion. This nomenclature is not to be confused with the previously used ‘generic’, which refers to an open subset of the space of solutions to the Einstein equation.

⁷Note that *prima facie*, the fact that we extend Theorem 2 of [32] to any fixed, typical U and ψ appears to contradict some of the statements in the introduction of [32], i.e., that randomness is a *sine qua non* for a quantum learning no-go theorem. This appears to be the case because if $U = U_0$ is fixed, there exists a trivial \mathcal{A}^U that just exactly outputs U_0 . The point is that U_0 is here ‘atypical’, and the algorithm will fail for all other $U \neq U_0$.

We call this result *Cryptographic Censorship*:

*(Pseudo)random dynamics guarantee event horizon formation in typical
(pseudo)random states.*

This is the first general condition that guarantees the formation of horizons without requiring asymptotic time evolution, thus addressing the first task above: the prescription of a broadly applicable quantum complexity theoretic diagnostic of event horizon formation.⁸ This result can be thought of as a quantum instability of generic asymptotically AdS spacetimes, assuming that typical pseudorandom dynamics are a good model for (sufficiently long) time evolution of generic spacetimes. We may apply this result to our two examples of code subspaces: microcanonical windows of width $O(N)$ and states with an identical leading order geometry at some fixed time. In the case of the former, we find that pseudorandom (respectively, random) states in a given microcanonical window are exponentially (respectively, doubly exponentially) likely to have horizons. In the case of the latter, we find that whenever it is possible to squeeze a gravitational amount of pseudorandomness into an $O(1)$ time band on the boundary, an event horizon must form.

Before we address the second task mentioned above — classifying the size of a singularity — let us motivate why pseudorandom time evolution is natural in strongly gravitating systems. There is a general expectation that some amount of (approximate) randomness is crucial for understanding black hole dynamics: due to scrambling, many aspects of black holes appear to be well modeled by Haar random unitary dynamics or sometimes by the weaker notion of k -designs [34]. The chaotic nature of black hole physics [36–38] and the connection between randomness and chaos [39] have strengthened this expectation. However, we also expect that time evolution in quantum gravity can be efficiently implemented. Haar random unitaries and (exact) k -designs do not have this property. We thus find pseudorandom unitaries to be a well-motivated model of black hole dynamics. On the boundary side of the AdS/CFT duality, many holographic CFTs are known to exhibit chaotic dynamics in typical (high energy) states. Here, we will show that sufficiently pseudorandom unitary (boundary) time evolution in fact *guarantees* the existence of an event horizon in typical states (and thus, by definition, of a black hole).⁹

Thus far we have argued that pseudorandomness is a valid model for time evolution of strongly gravitating systems using established facts about black holes. We shall postpone motivating pseudorandomness as a model for time evolution in the context of naked singularities until we have defined what we mean by “large” and “small” naked singularities. We therefore proceed to give an overview of this second task.

⁸Here we are defining an event horizon as $\partial J^-[\mathcal{S}]$, so this includes holographic cosmologies in AdS.

⁹At this point it is worth noting that [31] also discussed pseudorandomness in the context of reconstruction complexity: they showed that an easily prepared set of black hole microstates — specifically black holes formed by collapse together with the application of some randomly chosen shocks — should form a pseudorandom ensemble, and that as a consequence, boundary reconstruction of the volume of the black hole interior (for a code subspace containing all these states) must be exponentially complex. Our aim here is notably different: crucially, we do not start with an ensemble of black hole microstates, but instead show that gravitational pseudorandom boundary dynamics *implies* the existence of an event horizon.




Figure 2. A spacetime with a Python’s Lunch. In the Penrose diagram on the left, the lunch, causal wedge (\mathcal{W}_C) and the wedge of the outermost QES (\mathcal{W}_O) are indicated. The left and right green dots are the minimal and outermost QES, respectively. The right panel depicts a timeslice.

Size of a naked singularity.

Intuitively, the naked singularities that we would like quantum gravity to exclude are classically extended with divergent curvatures, and emitting large amounts of observable, causally propagating radiation as prescribed by some fundamental quantum gravity evolution.

In many cases, such large curvatures create (quantum) trapped surfaces and (quantum) extremal surfaces. In [40], unitary invariance of S_{vN} in the dual CFT was used to prove that classical extremal (and trapped) surfaces resulting from classical singularities that do not “evaporate” must lie behind horizons in AdS/CFT. As a warmup, we will prove that even upon inclusion of quantum corrections and under no restrictions on the behavior of the singularity, quantum extremal surfaces (QESs) generically lie behind horizons. The argument makes essential use of the geometrization of regions of high- and low-reconstruction complexity: the causal wedge, and more generally the wedge of the outermost QES, admits an efficient reconstruction — i.e., subexponentially complex in $\log \dim \mathcal{H}_{\text{code}}$, where $\mathcal{H}_{\text{code}}$ is as before the code subspace on which reconstruction is taking place [29], while the complementary wedge — the so-called Python’s Lunch — does not [41]. See figure 2. In particular, acting on the Python’s Lunch is an exponentially difficult (in $\log \dim \mathcal{H}_{\text{code}}$) task in the boundary theory. If any QES lies in causal contact of the boundary, it is possible to act on the lunch by throwing in some timelike probe, as illustrated in figure 3. This is a simple task, and thus cannot happen: QESs — minimal or otherwise — must therefore lie outside of the causal wedge¹⁰ of the relevant region.¹¹

¹⁰For those readers who have previously seen the Quantum Focusing Conjecture (QFC) used to justify why QESs lie behind horizons [42, 43], note that all such proofs to date have assumed cosmic censorship (in the form of global hyperbolicity). We prove this result using the holographic dictionary instead of using the QFC.

¹¹Any readers confused about the islands outside the horizon phenomenon [44, 45] should recall that the horizons of [44, 45] are not the horizons of the system for which the island is quantum extremal but are rather the horizons of a smaller system.

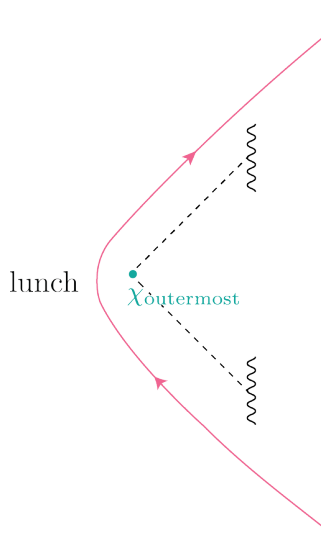


Figure 3. A forbidden situation: a QES in causal contact with the boundary.

However, singularities are not always associated with (quantum) trapped or (quantum) extremal surfaces. In these cases, the aforementioned result showing that extremal surfaces imply horizons will not be useful. However, if a singularity could be shown to be associated to a gravitational amount of pseudorandomness, then Cryptographic Censorship ensures that these singularities induce horizons. But how can we know which singularities are “severe enough” to accommodate gravitational pseudorandomness? As a first step towards this question, we give a precise definition of the size of a naked singularity Γ , denoted $|\Gamma|$. This quantity measures the minimal amount of CFT time evolution required to prescribe (or equivalently, obtain) all information about the singularity. For a timelike singularity, we furthermore show that this quantity also measures how many modes within the semiclassical EFT can interact with the singularity. Of course, the dynamics of the singularity itself would not be semiclassical, but we argue that in the $\ell_{\text{Pl}} \rightarrow 0$ limit, from the perspective of weakly backreacting probe fields, a timelike singularity might reasonably be conjectured to act primarily as a complicated UV-sensitive boundary condition for their evolution. Then we show that if $|\Gamma|$ does not go to zero as $\ell_{\text{Pl}} \rightarrow 0$, or if it goes to zero sufficiently slowly, then the singularity is compatible with inducing gravitational pseudorandomness in the dynamics. With this in mind, we classify singularities into three groups: classical, semi-Planckian, and Planckian. We speculate that classical naked singularities, and possibly semi-Planckian ones, are forbidden by Cryptographic Censorship. Compatibility with gravitational pseudorandomness is of course necessary but not sufficient to guarantee actual gravitational pseudorandomness in the dynamics. We of course cannot give a top-down derivation that the unitary dynamics of a holographic CFT are approximately pseudorandom on such code subspaces, but we shall motivate this as a reasonable assumption for generic classical timelike singularities. Cryptographic Censorship, together with this expectation, is the basis of our subsequent conjecture of Quantum Cosmic Censorship: that classical (and perhaps semi-Planckian) singularities must be hidden behind horizons; only Planckian singularities may be visible to the asymptotic observer.

Let us now motivate the assumption of pseudorandomness as a model for singularity dynamics (which we have previously motivated for black holes; see also [32, 35]). A natural source of pseudorandom dynamics associated to a singularity is due to the Planckian physics at the singularity itself — what we described above as a UV-sensitive boundary condition. We expect such a boundary condition cannot be predicted by effective field theory. We will argue in more detail in section 4.2 that this UV-sensitive boundary condition at the singularity could plausibly be the effect of pseudorandomness in the fundamental dynamics. We expect that in many cases, pseudorandomness is an accurate model for the boundary dynamics describing a singularity. When this criterion is satisfied, the theorem of [32] together with the argument outlined above guarantees that the spacetime has an event horizon.¹² It immediately follows that *whenever the fundamental dynamics of a singularity are well approximated by a pseudorandom unitary, the singularity incurs an event horizon.*

It is worth mentioning another potential source of pseudorandom dynamics associated to a singularity: spacetime dynamics in the vicinity of the singularity might induce chaos in the probe fields. The paradigmatic example of chaos in classical gravity — other than event horizons via the butterfly effect — is the BKL phenomenon [46]. This phenomenon, which involves chaotically oscillating metric components, is expected to be generic in the descent to a spacelike singularity [47]. BKL behavior is however *only* established for spacelike singularities; it is possible that certain timelike singularities could induce similar chaos upon approach (see, e.g., [48]), though we remain agnostic on the plausibility of this potential source of pseudorandomness.

We emphasize here that throughout this paper we will be considering pseudorandomness in the *boundary* dynamics. Readers interested in understanding what happens in the situation of bulk pseudorandomness should refer to section 5.

We now briefly summarize the paper. We begin in the next subsection by setting up notations and conventions. In section 2, we prove that nonminimal QESs must lie behind a horizon. In section 3 we prove that when the time evolution in the fundamental description is sufficiently pseudorandom, the bulk spacetime contains an event horizon. We then apply this to the microcanonical window and a code subspace representing states with the same (leading order) geometry for times of $O(G_N^0)$. In section 4, we provide a definition that quantifies the size of (the naked part of) a singularity, and we show that classical (and semi-Planckian) singularities can accommodate gravitational pseudorandomness. By Cryptographic Censorship, this would then imply that such singularities must incur horizons. In section 4, we also speculate on how Cryptographic Censorship can be used to prove a version of cosmic censorship. We end with a conjecture for a new, inherently quantum gravitational form of cosmic censorship that may be consistent with known counterexamples to the standard WCCC, including the evaporating black hole. In section 5, we end with a short discussion of our aspirations and a few possible future directions. In appendices A and B, we provide proof details for the statements of section 3.

¹²Note that our argument does not guarantee that the singularity lies behind the event horizon, nor should it: we expect that small singularities can exist outside of horizons. Our techniques do suggest that it may be possible to limit the extent to which a singularity can lie outside of a horizon, but we leave a detailed investigation of this question to future work.

1.1 Preliminaries

Assumptions. Unlike much of the AdS/CFT literature, we will not be imposing the assumption of global hyperbolicity on the bulk spacetimes under consideration, as this is tantamount to assuming weak (and strong) cosmic censorship. This means that, even in the $G_N \rightarrow 0$ limit, QFT in curved spacetime (including linearized gravitons) cannot be utilized for time evolution. However, we will still assume that QFT in curved spacetime is valid within causal diamonds that contain no singularities. We will further assume that any spacetime under consideration admits an initial data slice Σ with an everywhere defined and nowhere vanishing continuous timelike unit normal; in particular, we require that $D[\Sigma]$ contains an open set containing Σ . Any singularity in the spacetime will be assumed to act as a UV-sensitive boundary condition that determines the relation between the algebras of different causal diamonds.¹³ Finally, we will take a conformal frame which is taken to be fixed at any value of N , including in the large- N limit.

Assumptions on the semiclassical limit. We will also often implicitly use the framework of *asymptotically isometric codes* for discussing the notion of “geometric” states in the boundary theory. This framework was inspired by the work of [49, 50]. By “a geometric state,” we really have in mind that our bulk state is described by a sequence of boundary states, $|\psi_N\rangle \in \mathcal{H}_N$, labeled by an index N , which controls $1/G_N$ in the bulk. The Hilbert space \mathcal{H}_N is that of the boundary theory at finite N . In the case where the boundary theory is $\mathcal{N} = 4$ SYM, we have in mind that N is the rank of the gauge group, but more generally it is just some parameter in the boundary theory.

We now state what we mean by an asymptotically isometric code for our work. These assumptions are only a slight modification of those listed in [51].

Assumption 1. *The bulk Hilbert space is given by a direct sum $\mathcal{H} = \mathcal{H}_{g_1} \oplus \mathcal{H}_{g_2} \oplus \dots$, where \mathcal{H}_{g_i} is the Hilbert space of QFT states defined perturbatively about a fixed background manifold M_{g_i} . Furthermore, the algebra of bounded operators on this Hilbert space, $\mathcal{B}(\mathcal{H})$, breaks up into a direct sum over bounded operators on each Hilbert space individually, $\mathcal{B}(\mathcal{H}) = \bigoplus_i \mathcal{B}(\mathcal{H}_{g_i})$.*

Assumption 2. *There exists a sequence of bounded linear maps $V_N : \mathcal{H} \rightarrow \mathcal{H}_N$ such that for all $|\psi\rangle, |\phi\rangle \in \mathcal{H}$*

$$\lim_{N \rightarrow \infty} \langle \psi | V_N^\dagger V_N | \phi \rangle = \langle \psi | \phi \rangle, \tag{1.2}$$

where \mathcal{H}_N is the full boundary Hilbert space. From the results of [49], we will just need that for every bulk state ρ on $\mathcal{B}(\mathcal{H})$ and every operator $\mathcal{O} \in \mathcal{B}(\mathcal{H})$, there is a sequence of states ρ_N and boundary operators $\mathcal{O}_N \in \mathcal{B}(\mathcal{H}_N)$ such that

$$\lim_{N \rightarrow \infty} \rho_N(\mathcal{O}_N) = \rho(\mathcal{O}), \tag{1.3}$$

and where ρ_N and \mathcal{O}_N are related by conjugation by V_N to their bulk duals, $\rho = V_N^\dagger \rho_N V_N$ and $\mathcal{O}_N = V_N \mathcal{O} V_N^\dagger$.

¹³The singularity only has an effect on the relation between two diamonds D_1 and D_2 when the set $[J^+(D_1) \cup J^-(D_1)] \cap [J^+(D_2) \cup J^-(D_2)]$ contains part of the singularity. If this set does not contain a part of the singularity, the EFT can be used to relate the algebra of observables in D_1 and D_2 .

Given these assumptions, we can more carefully define what we mean by a boundary “geometric” state in the context of an asymptotically isometric code. Namely, we say that a sequence of states $\{\psi_N : \psi_N \in \mathcal{H}_N\}$ is *geometric* if there exists $|\psi\rangle \in \mathcal{H}$ such that $\|V_N |\psi\rangle - |\psi_N\rangle\|_1 = 1/\exp(N)$. In such a situation, we will sometimes say that the states $|\psi_N\rangle$ labeled by N are “represented” by the asymptotically isometric code.

In what follows, we will be concerned with whether or not bulk observables lie in the causal wedge. We will denote the algebra associated to the causal wedge as \mathcal{N}_{cw} . This algebra is a subalgebra of $\mathcal{B}(\mathcal{H}) \supseteq \mathcal{N}_{\text{cw}}$. Whether this inclusion is strict or not is the essential point of this work. If $\mathcal{B}(\mathcal{H}) \supsetneq \mathcal{N}_{\text{cw}}$, then there exists a horizon by definition.

In what follows, we will make various assumptions about the time evolution of the boundary theory. To conform with the standard expectations of AdS/CFT, we will assume that whenever we have an asymptotically isometric code, then bulk time evolution is boundary time evolution on the code subspace. We can encapsulate this in the following assumption.

Assumption 3. *Let H_{bulk} be the bulk Hamiltonian acting on \mathcal{H} . Furthermore, let H_N be the boundary Hamiltonian at finite N acting on \mathcal{H}_N . Then*

$$V_N^\dagger e^{-iH_N t} V_N = e^{-iH_{\text{bulk}} t}. \tag{1.4}$$

Notation and conventions. We now proceed to introduce and review some language and notation from the quantum computing literature, which will be important for the formulation of Cryptographic Censorship. The rest of this section contains notation that will mostly be used in section 3. We will use \mathcal{H} to denote Hilbert spaces and for a pure state $|\psi\rangle$, we will use ψ to denote the density matrix $|\psi\rangle\langle\psi|$. When we write $x \leftarrow \mathcal{X}$, we mean a variable x is drawn from a distribution \mathcal{X} . For example, by $U \leftarrow \mu$, we mean to draw a unitary U (from the unitary group $\mathbf{U}(\mathcal{H})$) using the Haar measure μ ; we will also use μ to denote the distribution of Haar random states, e.g., when we write $|\psi\rangle \leftarrow \mu$. We use $\Pr_{x \leftarrow \mathcal{X}}$ to denote taking the probability of some outcome occurring for x drawn from \mathcal{X} . We will often work with a parameter k , which in the cryptographic literature is called a “key”. This should be thought of as some small amount of random classical information that can be used to encode or decode data. Concretely, it is a bit string that is sampled uniformly from a *key space* \mathcal{K} , a set of bit strings. The dimension of \mathcal{K} scales with a parameter κ , which is frequently referred to in the literature as the “security parameter”. We will always have in mind that we are scaling $\log \dim \mathcal{H} = \text{poly}(\kappa)$ and will eventually take the $\kappa \rightarrow \infty$ limit. A function $f(x)$ is said to be *negligible* in its parameter x if for any constant $c > 0$, $|f(x)| \leq x^{-c}$ for any sufficiently large x . We will say that a quantity $f(x)$ is $O(g(x))$, denoted by $f(x) \sim O(g(x))$, if for some positive constants c_1, c_2 and for any sufficiently large x , $c_1 g(x) \leq |f(x)| \leq c_2 g(x)$.¹⁴ Finally, we will call an operator V *approximately norm-preserving* if it approximately preserves state norms, i.e., for any state $|\psi\rangle$, $\langle\psi|V^\dagger V|\psi\rangle$ is exponentially close to 1 for sufficiently large κ . This of course includes any unitary operators.

Our results will make heavy use of quantum algorithms, which we now define. A quantum algorithm \mathcal{A} is a map from an input Hilbert space \mathcal{H}_{in} to some output set \mathcal{C} . Note that the set \mathcal{C} can be a Hilbert space, a classical space, or some other set altogether. For example, a

¹⁴In computer science literature this is denoted by $f = \Theta(g(x))$.

quantum algorithm could take as input a quantum state $|\psi\rangle \in \mathcal{H}_{\text{in}}$ and output a classical bit representation of an operator \mathcal{O} (e.g., the coefficients of the representation of \mathcal{O} in some fixed choice of basis). A quantum algorithm \mathcal{A} acting on some \mathcal{H}_{in} is said to be *efficient*, *polynomially complex*, or equivalently *computationally bounded* if it can be implemented by a quantum circuit of size no greater than polynomial in $\log \dim \mathcal{H}_{\text{in}}$.¹⁵ Here we will typically take \mathcal{H}_{in} to be the code subspace under consideration. We will often drop the reference to $\log \dim \mathcal{H}_{\text{in}}$ when describing complexity or efficiency of such algorithms.

Consider now an efficient algorithm designed to take in a state $|\psi\rangle$ and output 0 if it determines that $|\psi\rangle$ was sampled from distribution μ_1 and 1 if it determines $|\psi\rangle$ came from μ_2 . In this case, we say the two ensembles are computationally indistinguishable — which we define precisely for general algorithms below — if any efficient algorithm succeeds at this task with probability upper bounded by 1/2 up to negligible corrections. It should be clear that this is equivalent to saying that the quantum algorithm outputs 1 (or equivalently 0) with almost the same probability when acting on both μ_1 and μ_2 .

Note that the authors of [33], whose definitions of pseudorandom unitaries and states we will use, demand that pseudorandom ensembles are indistinguishable from the Haar ensemble even if the algorithm is allowed a reasonable (polynomial) number of copies of the state. Having multiple copies of the state is essential to allow the algorithm to probe coherent information about the quantum state.

With this intuition, we present the exact definition of computational indistinguishability of ensembles of states for general quantum algorithms:

Definition 1 (Computational Indistinguishability). *Two ensembles of states, \mathcal{S}_1 and \mathcal{S}_2 , in some Hilbert space \mathcal{H} , with size parameterized by a parameter κ , are computationally indistinguishable if for any efficient quantum algorithm \mathcal{A} and for all $m = \text{poly}(\kappa)$,*

$$\left| \Pr_{|\psi_1\rangle \leftarrow \mathcal{S}_1} \left(\mathcal{A}(|\psi_1\rangle^{\otimes m}) = 1 \right) - \Pr_{|\psi_2\rangle \leftarrow \mathcal{S}_2} \left(\mathcal{A}(|\psi_2\rangle^{\otimes m}) = 1 \right) \right| = \text{negl}(\kappa) . \quad (1.5)$$

When we write $\mathcal{A} = 1$, we are considering that the algorithm is given some binary task (e.g., deciding whether a state is pseudorandom or not), and outputs either 0 or 1; we write $\mathcal{A} = 1$ to indicate that it outputs 1.

The specific quantum algorithms that we will consider are learning algorithms; these are efficient algorithms that take certain data as input, implement some processing with the goal of learning some outcome, and finally produce their best guess for the outcome. In our particular context, the learning algorithms will aim to learn a circuit (or algebraic) representation of a unitary operator U . During the initial learning phase, the algorithms are given “oracle access” to some unknown unitary U , denoted by \mathcal{A}^U : this means that the algorithms have the ability to query U a polynomial number of times on a polynomial number of (efficiently preparable) states $|\varphi_i\rangle \in \mathcal{H}$ to obtain $U|\varphi_i\rangle$ during its learning phase. During this phase, the algorithms can perform any polynomially complex computation between queries, on both the $U|\varphi_i\rangle$ from its queries thus far, and on its own registers. At the end of this phase, the algorithms’ oracle access to U is terminated, and they are tasked with

¹⁵This involves some canonical choice of gates; however, this choice does not change the scaling in $\log \dim \mathcal{H}_{\text{in}}$.

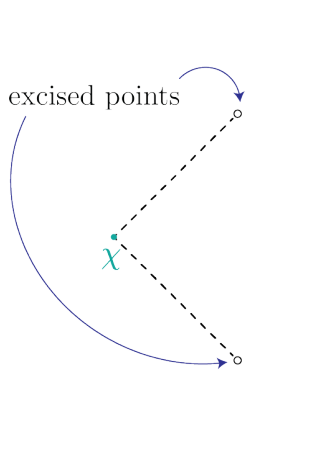


Figure 4. An example where null geodesics fired from χ are geodesically incomplete, but there is no horizon.

producing a maximally accurate representation of U or a particular aspect or feature of it: e.g., U acting on some initial state $|\psi_{\text{initial}}\rangle$ which would now be given to the algorithm.

Let us briefly make a note on the different parameters and dimensions used throughout this paper. In the context of AdS/CFT, we will be working with code subspaces of dimension $\dim \mathcal{H}_{\text{code}} = 2^\kappa$, where the security parameter $\kappa = \text{poly}(N)$. For concreteness, we will consider $G_N \sim 1/N^2$, but the relation between G_N and N can be changed to the reader’s liking. We will often use N instead of κ to parameterize functions, e.g., for negligible functions we will write $\eta(N)$.

Note that theorems and proofs in this paper are well above the physics level of rigor but fall short of a mathematician’s level of rigor in a few places. We identify these gaps where possible and leave a fully rigorous formulation thereof to the curious mathematician.

2 Warmup: QESs lie behind horizons

We will show in this section that if χ is a nonempty QES homologous to \mathcal{S} in a spacetime (M, g) , then it lies behind both past and future event horizons. The quantum focusing conjecture (QFC) [43] with reflecting boundary conditions at \mathcal{S} guarantees by the quantum singularity theorem [52] that at least one null geodesic fired from χ is geodesically incomplete.¹⁶ However, in the absence of a stringent causality assumption such as AdS hyperbolicity, null geodesic incompleteness does not guarantee the existence of horizons. See figure 4 for an illustration.

When χ is the minimal (and nonempty) QES for \mathcal{S} in M , it immediately follows from the QES formula relating $S_{\text{gen}}[\chi]$ to $S_{\text{vN}}[\rho_{\mathcal{S}}]$ that bulk perturbations locally propagating from \mathcal{S} (future- or past-directed) cannot modify $S_{\text{gen}}[\chi]$ and thus cannot reach χ : in this case, $\chi \cap I^\pm[\mathcal{S}] = \emptyset$.¹⁷ See figure 5 for a review of the argument, originally presented in [42].

¹⁶Strictly speaking, the restricted QFC [53] is sufficient to guarantee this.

¹⁷Note that we have excluded the future horismos $\partial I^\pm[\mathcal{S}]$ from this conclusion, as i^\pm are not precluded from null communication with χ .




Figure 5. A QES outside the horizon, leading to a contradiction with the unitary invariance of von Neumann entropy: $S_{\text{vN}}[\rho_{\mathcal{S}}] = S_{\text{vN}}[U^\dagger \rho_{\mathcal{S}} U]$.

However, from the perspective of cosmic censorship, global minimality of a QES is of little relevance. We expect that a similar result should apply to any QES independently of its global properties. In [40] some strides were made towards this end by CPT-conjugating the spacetime around the nonminimal QES χ to generate a new spacetime in which χ was minimal. This allowed [40] to prove that a large class of extremal surfaces and trapped surfaces must lie behind horizons. These results, however, relied on an assumption that we aim to relax in this section: that there can be no “evaporating singularities” (see [40] for the precise definition).

To relax this assumption, we will instead assume the Python’s Lunch proposal [28, 41, 54] in addition to the QES formula to argue that (quantum) extremal surfaces must generically be hidden behind both past and future horizons.¹⁸ Or more precisely, if they lie outside a horizon, this distance must shrink to zero in the $G_N \rightarrow 0$ limit. Let us briefly review some facts about the Python’s Lunch proposal. A Python’s Lunch L with respect to a conformal boundary \mathcal{S} is a region of spacetime defined by (at least) three mutually spacelike separated QESs homologous to \mathcal{S} : χ_{min} , χ_{bulge} , and χ_{aptz} . The first is the globally minimal QES, while χ_{aptz} is a candidate QES, but having larger generalized entropy than χ_{min} . The surfaces χ_{min} and χ_{aptz} are so-called throats: QESs whose generalized entropy is roughly locally minimal in space and locally maximal in time (see [54] for details). The bulge surface χ_{bulge} is also a QES, but it is locally maximal in time and space (see again [54] for details). The region L is defined as the domain of dependence of any spacelike slice bounded by χ_{min} and χ_{aptz} , and the Python’s Lunch proposal conjectures that the CFT computational complexity required to implement any operator in L scales as

$$\exp \left[\frac{S_{\text{gen}}[\chi_{\text{bulge}}] - S_{\text{gen}}[\chi_{\text{aptz}}]}{2} \right]. \tag{2.1}$$

We will only identify a Python’s Lunch in a portion of spacetime that is globally hyperbolic, and we will make the reasonable assumption that the presence of potential spacetime singularities away from L does not make operator reconstruction any easier.

Besides the Python’s Lunch proposal, our assumptions in this section (and this section only) are as follows.

¹⁸Note that the assumption of the Python’s Lunch proposal applies to this section only; it will not be used in the proof of Cryptographic Censorship in section 3.3.

Assumptions and Definitions. We assume the following in addition to our general assumptions stated at the beginning of section 1.1:

- (M, g) has reflecting boundary conditions at each connected boundary component \mathcal{I} .
- Smooth initial data: (M, g) has a smooth inextendible timeslice Σ_0 containing two QESs χ_{\min} and χ homologous to \mathcal{I} , and where χ_{\min} is the QES of minimal generalized entropy that is homologous to \mathcal{I} .
- A generic condition: if χ' is a bulge QES and χ'' a throat QES, then their associated generalized entropies satisfy $S_{\text{gen}}[\chi'] - S_{\text{gen}}[\chi''] \sim O(G_N^{-1})$.
- In any globally hyperbolic region of spacetime that we consider, the (restricted) QFC holds, and the quantum maximin [25, 55] and maximinimax [41] prescriptions converge.
- We denote by $W_O[\chi]$ the outer wedge of a QES χ [27, 56].
- We define $t = 0$ at \mathcal{I} as $\Sigma_0 \cap \mathcal{I}$.
- If a point p in the bulk is such that any future (past) directed causal curve from p to \mathcal{I} arrives at \mathcal{I} at a boundary time that diverges to $+\infty$ ($-\infty$) in the $G_N \rightarrow 0$ limit, or if no curve to \mathcal{I} exists, we say that p is behind an “effective future (past) horizon”.

Theorem 1. *Let (M, g) be a spacetime as above. Assuming the Python’s Lunch conjecture for complexity of operator reconstruction, the QES χ lies behind past and future effective horizons.*

Proof. Consider first working strictly within the globally hyperbolic spacetime given by the domain of dependence of Σ_0 within M , denoted as $D[\Sigma_0]$. By Proposition 1 of [29], there always exists an outermost minimal QES χ_{aptz} that lies in $W_O[\chi] \cap W_O[\chi_{\min}]$.¹⁹ This surface may or may not coincide with χ itself. Let now H be a spatial slice with $\partial H = \chi_{\text{aptz}} \cup \chi_{\min}$, which exists since $\chi_{\text{aptz}} \subset W_O[\chi_{\min}]$, and let $L = D[H]$. See figure 6 for an illustration. By the maximinimax construction in the appendix of [41], there is another QES χ_{bulge} in L that is of the bulge type, and so L must be a Python’s Lunch ($\chi_{\text{bulge}} = \chi$ is allowed). L must contain χ , since there exists a choice of H that contains χ_{\min}, χ , and χ_{aptz} . If χ is equal to χ_{aptz} or χ_{bulge} , this is trivial. If χ is neither of these surfaces, then we construct H by gluing together two pieces. The first piece is the subset of Σ_0 lying between χ_{\min} and χ . The second is any spacelike slice between χ and χ_{aptz} which exist by global hyperbolicity of $D[\Sigma_0]$ and the fact that χ_{aptz} lies strictly in the outer wedge of χ .

Let us now return to the full spacetime M . Define a code subspace $\mathcal{H}_{\text{code}}$ consisting of all low energy perturbative quantum states that backreact on Σ_0 at subleading order in G_N . By the genericity assumption, we have $S_{\text{gen}}[\chi_{\text{bulge}}] - S_{\text{gen}}[\chi_{\text{aptz}}] \sim O(G_N^{-1})$ so by the Python’s Lunch proposal, the complexity of reconstructing operators in $\mathcal{H}_{\text{code}}$ with support in L scales as $\exp(cG_N^{-1})$ for some constant $c > 0$. Assume now for contradiction that χ lies outside the

¹⁹Outermost minimal means that there is no surface in the outer wedge $W_O[\chi_{\min}]$ homologous to χ_{\min} with less S_{gen} .

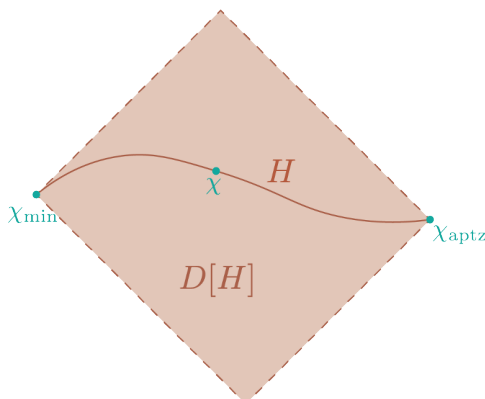


Figure 6. The setup in the proof of Theorem 1 where χ is nonminimal, and where H is chosen to contain χ .

effective past horizon. Then there exists a future-directed causal curve γ from \mathcal{S} to χ that is fired from \mathcal{S} at a time $t \sim O(G_N^0)$. Then we can find a simple operator Φ that modifies the quantum fields at χ . Just let Φ be an operator that injects a perturbative amount of matter with rocket boosters near \mathcal{S} , so that the matter travels along the curve γ . Because this matter reaches χ in a time that does not scale with G_N , this matter backreacts on Σ_0 only at subleading order, and so it acts within the code subspace. Φ can be constructed in two steps: an injection of matter near the boundary using HKLL [57–59] followed by an $O(G_N^0)$ amount of CFT time evolution. Neither is exponentially hard, and so we reach a contradiction with reconstruction complexity of operators with support in L . The argument excluding χ from intersecting the exterior of the effective future horizon proceeds *mutatis mutandis*. \square

Note that there are setups where islands [44, 45], and thus QESs χ , lie outside of event horizons. A universal aspect of these constructions is that the QES in question is a QES of a different subsystem than the one whose horizon is under consideration. We may also comment more specifically on each example: first, ref. [44] found an island outside of the black hole horizon when the CFT at \mathcal{S} was coupled to a reservoir. In this case boundary conditions were not reflecting, which was assumed in our result. Furthermore, by virtue of the coupling, time evolution of the black hole system alone is no longer unitary. Thus, while a signal sent from \mathcal{S} initially can be turned near the boundary with a unitary, propagating it to χ requires a non-unitary operation on the CFT, and so the entropy is expected to change. Additionally, decoupling the CFT from the auxiliary system produces a shockwave that alters the spacetime at leading order, which does push the island behind the horizon. Next, in [45] an island was found lying a distance of $O(\sqrt{\ell_{\text{Pl}} r_{\text{hor}}})$ outside the horizon. In this case the island was associated to angular momentum modes in the bulk, but qualitatively this construction is the same as ref. [44] as far as it relates to the above theorem.

3 Cryptographic Censorship

In this section we introduce and prove our main result: that on any sufficiently large code subspace, (pseudo)random dynamics guarantee event horizon formation in typical

(pseudo)random states with a semiclassical bulk dual. The precise statement (Theorem 5: Cryptographic Censorship) is given in section 3.3. In presenting and proving this theorem, we will be agnostic about the types of code subspaces that in a gravitational theory could correspond to this degree of pseudorandomness; in section 3.4 we will give two examples of code subspaces in gravity that are compatible with sufficiently pseudorandom dynamics.

Intuitively, pseudorandom unitary ensembles are ensembles of unitaries that are *computationally indistinguishable* from the Haar ensemble. As an example, upon evolution for some time longer than a scrambling time, a chaotic system is generally expected to become approximately pseudorandom. Our argument for the existence of a horizon hinges upon the assumption that the fundamental time evolution is well modeled by a unitary that is sufficiently pseudorandom on the code subspace. In a sense described below, we will require the unitary to be pseudorandom by an amount that scales inversely with G_N , which we will call *gravitationally* pseudorandom.

How is this assumption of gravitational pseudorandomness sufficient to guarantee the existence of event horizons? We will outline the logic here while omitting many of the technicalities (which we leave to appendices A and B); the rest of this section is devoted to the technical points necessary for the proof of Cryptographic Censorship and its application to gravitational code subspaces.

Recall that observables in the causal — and more generally, simple [27, 29, 56] — wedge of the boundary are computationally simple to reconstruct. This is a non-trivial aspect of bulk reconstruction in AdS/CFT [27, 60].²⁰ We argue that, given some reasonable technical assumptions outlined in section 3.3, there exists a learning algorithm \mathcal{A}_{CW} that can reconstruct the causal wedge of any geometric state in a code subspace of size $e^{O(G_N^a)}$ for some $-1 \leq a < 0$. However, a theorem of [32] (see Theorem 2 below for its precise statement) combined with results from measure concentration guarantees (roughly) that all efficient quantum algorithms must fail at predicting the expectation value of some operator Q , which we show exists in the large- N limit, after pseudorandom time evolution. It then immediately follows that for the algorithm \mathcal{A}_{CW} , there exists an operator Q that \mathcal{A}_{CW} cannot reconstruct and therefore must lie outside of the causal wedge; in other words, an exterior to the causal wedge exists — by extension, a boundary to the causal wedge exists. Since the boundary of the causal wedge of \mathcal{S} coincides with the event horizon, an event horizon must exist. In section 3.3, we will formalize this rough argument into a proof.

We start by briefly reviewing pseudorandom unitaries, as well as the appropriate learning result in [32] that we will use, in section 3.1. We then explain the required alterations and extensions necessary to apply the learning result to the holographic, $N \rightarrow \infty$ setting in section 3.2. In section 3.3 we state and prove our theorem of Cryptographic Censorship. Finally, in section 3.4 we comment on two code subspaces to which our theorem applies.

²⁰For the likely limited number of readers who believe in the Quantum Extended Church-Turing (QECT) thesis outside of horizons but are skeptical of arguments relying on AdS/CFT, we note that the QECT thesis by itself also implies that causal observables must be simple.

3.1 Review: pseudorandomness and complexity of learning

Pseudorandom states and unitaries will play an important role in what follows, so we will briefly review these concepts. Within the context of quantum cryptography, pseudorandom states were first defined in [33]. Informally, an ensemble of quantum states is pseudorandom if the states are efficiently constructible, but indistinguishable from Haar random by any polynomially bounded quantum algorithm — even if that algorithm is given polynomially many copies of the state. A pseudorandom state generator G is a map from some space \mathcal{K} into a Hilbert space \mathcal{H} . For each element $k \in \mathcal{K}$, it outputs $G(k) = |\psi_k\rangle \in \mathcal{H}$. The intuitive goal of the definition is to say that the action of feeding a random sample from \mathcal{K} to G to get a state $|\psi_k\rangle$ is indistinguishable from sampling the Haar random distribution of states (in the computational sense defined in section 1.1). The technical definition is as follows:

Definition 2 (Pseudorandom State Ensemble (PRS); Definition 2 in [33]). *Consider an ensemble of states $\Psi = \{|\psi_k\rangle \in \mathcal{H}\}_{k \in \mathcal{K}}$, labeled by a parameter k drawn from a key space \mathcal{K} , where both \mathcal{H} and \mathcal{K} are parameterized by a security parameter κ . Then the ensemble Ψ is a pseudorandom state ensemble if*

- *The states can be efficiently generated: there is an efficient quantum algorithm G that generates the state $|\psi_k\rangle$ given k as input. That is, for all $k \in \mathcal{K}$, $G(k) = |\psi_k\rangle$.*
- *The states are computationally indistinguishable from Haar random: for any polynomially many copies of $|\psi_k\rangle$ with the same $k \in \mathcal{K}$ and any efficient quantum algorithm \mathcal{A} , there exists a negligible function $\text{negl}(\kappa)$ such that*

$$\left| \Pr_{k \leftarrow \mathcal{K}} \left[\mathcal{A} \left(|\psi_k\rangle^{\otimes \text{poly}(\kappa)} \right) = 1 \right] - \Pr_{|\psi\rangle \leftarrow \mu} \left[\mathcal{A} \left(|\psi\rangle^{\otimes \text{poly}(\kappa)} \right) = 1 \right] \right| = \text{negl}(\kappa), \quad (3.1)$$

where μ is the Haar measure on \mathcal{H} .

We can define pseudorandom unitaries as in [33] in a similar fashion:

Definition 3 (Pseudorandom Unitary Ensemble (PRU); Definition 5 in [33]). *Consider an ensemble of unitaries, $\mathcal{U} = \{U_k \in \mathbf{U}(\mathcal{H})\}_{k \in \mathcal{K}}$, labeled by a parameter k drawn from a key space \mathcal{K} , where both \mathcal{K} and \mathcal{H} are parameterized by a security parameter κ . We say \mathcal{U} is a pseudorandom unitary ensemble if*

- *The unitaries are efficiently representable: there is an efficient quantum algorithm Q such that for every $k \in \mathcal{K}$ and any $|\psi\rangle \in \mathcal{H}$, $Q(k, |\psi\rangle) = U_k |\psi\rangle$.*
- *The unitaries are computationally indistinguishable from Haar random: for any efficient quantum algorithm \mathcal{A}^U given oracle access to some U , there exists a negligible function $\text{negl}(\kappa)$ such that*

$$\left| \Pr_{k \leftarrow \mathcal{K}} \left[\mathcal{A}^{U_k} = 1 \right] - \Pr_{U \leftarrow \mu} \left[\mathcal{A}^U = 1 \right] \right| = \text{negl}(\kappa). \quad (3.2)$$

We note that technically, PRSs and PRUs are actually sequences of ensembles, $\Psi = \{\Psi_\kappa\}_\kappa$ and $\mathcal{U} = \{\mathcal{U}_\kappa\}_\kappa$, respectively, each parametrized by $\kappa \in \mathbb{N}$, and the guarantees in the above

definitions are for large κ . Since we will take $\kappa = \text{poly}(N)$ in section 3.2 and later sections, this aligns well with the framework of asymptotically isometric codes outlined in section 1.1. For ease of notation, we will instead simply write Ψ and \mathcal{U} , as in the above definitions, where the key k should be understood to be drawn from a key space \mathcal{K}_κ where κ is sufficiently large. We will make an exception to this rule in some parts of the proof in appendix B, where it will be convenient to make the dependence on the security parameter explicit to take the $\kappa \rightarrow \infty$ limit.

Finally, we remark that while explicit constructions of PRSs exist (and were given first in [33]), there are currently no known constructions of PRUs, although candidate constructions were proposed in [33]. We will not require any explicit constructions of either PRSs or PRUs in this work.

Pseudorandom states vs. k -designs. For the reader familiar with the related notion of a k -design, we briefly comment on the two main differences with PRSs. First, the associated notion of indistinguishability is much stronger for k -designs: the states need to be indistinguishable from Haar random even to computationally *unbounded* observers, e.g., algorithms of exponential complexity. However, k -designs only require that the moments of these ensembles are indistinguishable up to a certain value (k), whereas PRS ensembles are computationally indistinguishable for any $\text{poly}(\kappa)$ number of moments. In this sense, PRSs are comparable to a high-order k -design but with a parametrically weaker notion of indistinguishability. Note that k -designs are less relevant for our purposes as they apply to unbounded observers, whereas our focus here is on bounded computations.

We now state the learning result of [32], which was briefly reviewed at a lower level of detail in section 1. This result (and our Cryptographic Censorship result) also holds for Haar random U and $|\psi\rangle$, but we will focus on pseudorandom states and unitaries as we believe they present a more realistic characterization of the time evolution of physical systems.

Theorem 2 (Complexity of Learning Pseudorandom Unitaries, Theorem 2 of [32]). *Let $\kappa \in \mathbb{N}$ be the security parameter, and let \mathcal{U} be any pseudorandom unitary ensemble and Ψ be any pseudorandom state ensemble in a Hilbert space \mathcal{H} where $|\mathcal{H}| = 2^\kappa$. For any quantum algorithm \mathcal{A} of $\text{poly}(\kappa)$ complexity that, given oracle access to the pseudorandom unitary $U \leftarrow \mathcal{U}$, produces (a circuit representation of) an approximately norm-preserving operator²¹ \hat{U} as its best guess for U , there exists a constant $\alpha > 0$ independent of κ such that for any $\kappa > \kappa_0$, the following quantity is bounded:*

$$\text{avg}_{\substack{U \leftarrow \mathcal{U} \\ |\psi\rangle \leftarrow \Psi \\ \hat{U} \leftarrow \mathcal{A}^U}} \left[F(\hat{U}|\psi\rangle, U|\psi\rangle) \right] \leq 1 - \alpha .$$

Here, $F(\hat{U}|\psi\rangle, U|\psi\rangle) = |\langle \psi | \hat{U}^\dagger U | \psi \rangle|^2$ which is the fidelity when \hat{U} is unitary. In this context, it is approximately the fidelity since we are working with approximately norm-preserving operators.

²¹We note that [32] showed this result for a wider class of algorithms — those producing any operator \hat{U} with column norms bounded above by ≈ 1 — which includes the algorithms here. We will not need this more general result in this paper.

This theorem shows that it is impossible for any efficient quantum learning algorithm to use oracle access to a pseudorandom U to produce a polynomially complex circuit or (succinct) algebraic expression that accurately implements U on a pseudorandom state $|\psi\rangle$. As we will explain in more detail in the next section, we will be interested in applying this result to a sequence of code subspaces $\{\mathcal{H}_{\text{code}}^{(\kappa)}\}_\kappa$ each of dimension $|\mathcal{H}_{\text{code}}^{(\kappa)}| = 2^\kappa$, where we will take $\kappa = \text{poly}(N)$.

We now comment on two aspects of this theorem that will appear extensively in the subsequent section. First, Theorem 2 implies that, *on average*, there exists a distinguishing operator that the learning algorithm fails to predict (see appendix B in [32] for an explanation).²² In particular, when the fidelity between $\widehat{U}|\psi\rangle$ and $U|\psi\rangle$ is $1 - O(1)$, then there exists an operator Q such that²³

$$\text{avg}_{\widehat{U} \leftarrow \mathcal{A}^U} \left| \text{tr}(Q\widehat{U}\psi(\widehat{U})^\dagger) - \text{tr}(QU\psi(U)^\dagger) \right| \geq O(1) . \tag{3.3}$$

In the next section, we will use measure concentration to show that for *nearly all* of the unitaries in the ensemble \mathcal{U} , such a fidelity bound holds and thus there is a distinguishing operator.²⁴ This operator will play a crucial role in the proof of Cryptographic Censorship.

Second, it immediately follows from Theorem 2 that

$$\sup_{\kappa \rightarrow \infty} \text{avg}_{\substack{U \leftarrow \mathcal{U} \\ |\psi\rangle \leftarrow \Psi \\ \widehat{U} \leftarrow \mathcal{A}^U}} \left[F(\widehat{U}|\psi\rangle, U|\psi\rangle) \right] \leq 1 - \alpha , \tag{3.4}$$

which crucially depends on the upper bound of $1 - \alpha$ being independent of κ . This will be important in extending and applying this result to holography, where we will take the large κ (or equivalently, large N) limit, and want the guarantee that the large N limits of these states also have fidelity at most $1 - \alpha$, and thus a distinguishing operator exists.

3.2 Complexity of learning in holography

We now turn to the application of this learning result to holography. Since our goal is to constrain semiclassical gravitational phenomena, our focus will be on the large- N limit in general and on quantities that scale with N in this limit in particular. For concreteness, and without loss of generality, we use $G_N \sim 1/N^2$, and consider code subspaces $\mathcal{H}_{\text{code}}$ of dimension $e^{O(G_N^a)}$, where $-1 \leq a < 0$. To this end, we will take the security parameter $\kappa = \text{poly}(N)$ such that our code subspace $\mathcal{H}_{\text{code}}$ is of dimension 2^κ . Unitary operators that are pseudorandom on such code subspaces at any large but finite N will be called gravitationally pseudorandom (defined precisely below); we will also resolve various technical challenges that follow from taking the large- N limit (e.g., with regards to the existence of certain quantities). We begin with a formal definition of gravitational pseudorandomness.

²²As a difference of language, [32] writes this as the distinguishing ‘‘POVM’’ and lower bounds the distance between measurement outcomes obtained via this POVM.

²³Note that the quantum algorithm \mathcal{A}^U , as a result of, e.g., intrinsic randomness in the outcome of measurements during the learning process, generally outputs a \widehat{U} that is not a deterministic function of U . Therefore, we average over the algorithm’s output $\widehat{U} \leftarrow \mathcal{A}^U$.

²⁴The result on average in [32] only implies that a constant fraction of the unitaries in \mathcal{U} have such an operator, whereas we will show this holds for a fraction of the unitaries that goes to 1 as $\kappa \rightarrow \infty$.

Definition 4 (Gravitationally Pseudorandom Unitary Ensemble (GPRU)). *Let $\mathcal{U} = \{U_k \in \mathbf{U}(\mathcal{H})\}_{k \in \mathcal{K}}$ be a pseudorandom unitary ensemble, with key space \mathcal{K} parameterized by a security parameter κ , where \mathcal{H} is the Hilbert space of a large N , holographic quantum system where $N \sim G_N^{-1/2}$, with G_N the bulk Newton’s constant. We say that \mathcal{U} is a gravitationally pseudorandom unitary ensemble (GPRU) if $\kappa = \text{poly}(N)$.*

As noted previously, at strictly infinite N there is no notion of pseudorandom unitaries. When discussing operators we shall be primarily interested in asymptotic scaling at large N rather than the actual operators at infinite N . When working in the actual infinite N regime, we will only discuss limits of pseudorandom *states*, using the asymptotically isometric codes framework outlined in section 1.1.

Recall that we aim to apply the results of [32] to the holographic setup. The careful reader may have by now noticed two potential points of concern: first, Theorem 2 was proved on average for a PRU ensemble rather than for some typical instance; second, it is not obvious that in the large- N limit, for a typical PRU, a distinguishing operator exists. Let us briefly elaborate on each concern before proceeding to state and prove our generalization of Theorem 2 that addresses both.

On the first concern, the learning result of [32] is a statement about the *average* fidelity between the evolution of a pseudorandom state (PRS) evolved with the actual PRU and the same state evolved with the algorithm’s best guess. The averaging is over the unitary drawn from any PRU ensemble \mathcal{U} , the state drawn from any PRS ensemble Ψ , and over the output of any quantum learning algorithm \mathcal{A} .

While much of the literature (see, e.g., for a few instances [34, 61, 62]) does model gravitational systems via randomly sampled states or dynamics, this approach is in some tension with the fact that there is just *one* (fixed) unitary operator that describes the time evolution of $\mathcal{N} = 4$ SYM. We therefore generalize the bound from an average to a *typical* pseudorandom unitary and state by bounding the fluctuations of typical members of the ensemble, first for Haar random unitaries and states and then for pseudorandom unitaries and states. We emphasize that this means a learning bound applies to any *single fixed* unitary and state that are both typical, as we state below in Theorem 3.

The second concern, involving the existence of a “non-simple observable” — i.e., the distinguishing observable that the algorithm fails to predict — for typical instances of the ensemble at large- N , is a bit more subtle. As a corollary to our theorem for learning typical pseudorandom unitaries at large but finite N , we prove, under a reasonable assumption on the limit of states as $N \rightarrow \infty$, that our theorem also bounds the fidelity at $N \rightarrow \infty$. Consequently, our theorem guarantees that for any simple boundary algorithm, there exists a *semiclassical* observable that cannot be reconstructed by the algorithm.

We now state the holographic version of Theorem 2, incorporating the generalizations advertised above. We will subsequently use this version in section 3.3 to prove Cryptographic Censorship.

Consider a CFT with a large but finite UV energy cutoff, so that when N is finite the dimension of the Hilbert space is finite.²⁵ We prove the theorem below for any large but

²⁵We are free to consider other ways of regulating the dimension of the Hilbert space if needed.

finite value of N , under the assumption that the algorithm's output as a function of U is (randomized) K -Lipschitz for some fixed constant K . We will describe this assumption in greater length in appendix A, but it boils down to a certain notion of continuity controlled by a parameter K of the algorithm as a function of U .²⁶ We subsequently state and prove a corollary of this theorem for infinite-dimensional Hilbert spaces, under a class of assumptions. To increase readability, we postpone both of the proofs to appendices A and B.

Theorem 3 (Complexity of Learning for Typical Pseudorandom Unitaries). *Let $\kappa = \text{poly}(N)$ and consider a code subspace of dimension 2^κ . Consider any (gravitationally) pseudorandom unitary ensemble \mathcal{U} , any pseudorandom state ensemble Ψ , and any constant $\epsilon > 0$. For any (randomized) K -Lipschitz quantum algorithm \mathcal{A} of $\text{poly}(N)$ complexity that has oracle access to $U \leftarrow \mathcal{U}$ and produces an approximately norm-preserving operator \hat{U} as its best guess for U , and for any sufficiently large N , the following holds: for a fraction $\geq 1 - e^{-\epsilon^2 2^\kappa / (48(1+K)^2)} - e^{-\epsilon^2 2^\kappa / 192} - \text{negl}'(N)$ of $U \in \mathcal{U}$ and $|\psi\rangle \in \Psi$, we have that:*

1. the following quantity is bounded

$$\text{avg}_{\hat{U} \leftarrow \mathcal{A}^U} \left[F(U|\psi\rangle, \hat{U}|\psi\rangle) \right] \leq 1 - \alpha + 2\epsilon + \text{negl}(N), \tag{3.5}$$

where α is a constant independent of N , and

2. there exists an operator Q (which may depend on U , $|\psi\rangle$, and \hat{U}) such that

$$\text{avg}_{\hat{U} \leftarrow \mathcal{A}^U} \left| \text{tr}(Q\hat{U}\psi(\hat{U})^\dagger) - \text{tr}(QU\psi(U)^\dagger) \right| \geq 2\alpha - 4\epsilon - \text{negl}(N). \tag{3.6}$$

Proof. See appendix A. □

We now extend this theorem to infinite-dimensional Hilbert spaces. For this purpose, in what follows, we reinstate the sequences of PRSs and PRUs to make the N -dependence explicit. We now assume that for every N we have a code subspace \mathcal{H}_N on which the time evolution operator $e^{-iH_N t}$ of the CFT, for some t (possibly scaling with N), has the same matrix elements on the code subspace as another unitary operator $U_N : \mathcal{H}_N \rightarrow \mathcal{H}_N$, up to exponential corrections. As in Assumption 2, we also assume that there exists an asymptotically isometric code with map V_N from the bulk Hilbert space, \mathcal{H} , to \mathcal{H}_N . Furthermore, we assume there exists a (gravitationally) PRU ensemble $\{\mathcal{U}_N\}_N$ and some N_0 such that for all $N \geq N_0$, the ensemble \mathcal{U}_N contains U_N as a typical member.

Next, we assume there exists a sequence of pseudorandom states $\{|\psi_N\rangle \in \mathcal{H}_N\}_N$ that are represented on the code. Namely, there exists a $|\xi\rangle \in \mathcal{H}$ such that for every $|\psi_N\rangle$ in the sequence with sufficiently large N , $|\psi_N\rangle = V_N |\xi\rangle$ up to exponentially small corrections in the one-norm. Finally, letting the learning algorithm's best guess for U_N be \hat{U}_N , we assume that $\hat{U}_N |\psi_N\rangle$ is also represented by the asymptotically isometric code, i.e., there exists a $\hat{\xi}$ such that $\hat{U}_N |\psi_N\rangle = V_N |\hat{\xi}\rangle$. We also assume that the subset of states at finite N that are the fraction of states that are typical converge to a subset of states in the infinite- N Hilbert space that is of measure 1 (as the fraction of typical states limits to 1 as $N \rightarrow \infty$ in Theorem 3).

²⁶It is possible that this assumption actually follows directly from our other assumptions; we leave determination of this question to future work.

Corollary 3.1 (Distinguishing Operator at $N \rightarrow \infty$). *Let $\{U_N |\psi_N\rangle\}_N$ and $\{\widehat{U}_N |\psi_N\rangle\}_N$ be two sequences of states that are represented on the code, namely there exist $|\xi\rangle$ and $|\widehat{\xi}\rangle$ such that $U_N |\psi_N\rangle = V_N |\xi\rangle$ and $\widehat{U}_N |\psi_N\rangle = V_N |\widehat{\xi}\rangle$. Furthermore, assume that $\{U_N |\psi_N\rangle\}_N$ and $\{\widehat{U}_N |\psi_N\rangle\}_N$ satisfy the assumptions stated above, then both 1. and 2. in Theorem 3 also hold for the states $|\xi\rangle$ and $|\widehat{\xi}\rangle$. In particular, by taking the constant ϵ to be $< \alpha/2$, we have for any $|\widehat{\xi}\rangle$ the existence of an operator $Q \in \mathcal{B}(\mathcal{H})$ such that*

$$\text{avg}_{\widehat{\xi}} \left| \text{tr}(Q\widehat{\xi}) - \text{tr}(Q\xi) \right| \geq O(1) .$$

Proof. See appendix B. □

3.3 Pseudorandomness implies horizons

We are now ready to argue that when time evolution in the boundary theory is sufficiently (i.e., gravitationally) pseudorandom, the bulk has an event horizon. Recall that the logic in this argument is that for a typical pseudorandom time evolution operator U , (a) efficient quantum learning algorithms cannot accurately compute the time evolution of $\langle Q \rangle$ for some operator Q that may depend on U , and (b) there exists an efficient quantum algorithm that can accurately compute the time evolution of all local operators — i.e., operators in the causal wedge — under U . Point (a) follows directly from Corollary 3.1.

Point (b) is in a sense obvious holographically given the developments on complexity of reconstruction in AdS/CFT. The general expectation following from [29] is that a simple quantum algorithm should be able to reconstruct the causal wedge. We might worry that the existence of a naked singularity complicates this story, but there is a simple physical argument suggesting otherwise. To measure an operator in the causal wedge, we can send rocket ships from the conformal boundary that travel into the bulk, carry out measurements anywhere in the causal wedge, and then travel back out to the boundary with the result. Since we are only interested in measuring operators that survive at strictly infinite N , the complexity to implement these rocket ships from the CFT should not scale exponentially with N .²⁷ Under the assumption that there exists an efficient process that can produce a guess V for the time evolution operator U using these measurements, we can then combine this with point (a) above to formulate our theorem of Cryptographic Censorship, Theorem 5 below.

In many cases, we can motivate this assumption — the existence of an efficient learning algorithm that can produce such a guess V — from a computer science point of view. For instance, consider first a domain of dependence corresponding to a boundary time band. If we have access to the time-evolved expectation values of single-trace operators in that time band, we can use the HKLL protocol [57–59] to reconstruct this domain of dependence. See figure 7. In the case where the causal wedge can be generated from a union of such wedges, and if there exists a simple algorithm that can aggregate all these reconstructed wedges into a guess for the time evolution operator, then we can efficiently reconstruct the full causal wedge. There is in fact a learning algorithm that comes close to providing such a guess V ,

²⁷One might worry that the bulk measurements take too long for the rocket ships to be able to get back out to the boundary. Again, this should not be the case for bulk operators that survive the infinite N limit.




Figure 7. In the leftmost and middle panels, the reconstruction of operators inside the green wedge (which is globally hyperbolic) can be done using the HKLL protocol, which should be simple, insofar as we can compute the expectation values of local operators in the boundary theory on the time band Δt . We can keep doing this for subsequent time bands until we have reconstructed the full causal wedge, as in the rightmost panel. We use this as an ingredient to motivate the existence of an algorithm that can output an operator V which reproduces the time evolution of operators in the causal wedge.

as we will explain in more detail below; this thus provides strong motivation that learning the causal wedge is simple, even in the presence of such naked singularities.

However, it is possible that there exist naked singularities with unusual topologies where this reasoning falls short and we cannot use this ‘progressive HKLL’ protocol combined with the aforementioned existing learning algorithm. An illustration of such a hypothetical singularity is given in figure 8, where the interior of a cylindrical singularity does not lie inside any HKLL wedge. In such cases, however, we revert to our initial holographic reasoning to motivate our assumption: we can still use rocket ships to measure operators in the entire causal wedge.

For the case where HKLL in progressive wedges does work, we will now illustrate more precisely that under a mild technical assumption, there exists a quantum learning algorithm satisfying the assumptions of Theorem 2 that, given oracle access to U , can output an operator V where $V|\psi\rangle$ can accurately reproduce the expectation values of all simple operators in the state $U|\psi\rangle$. Applicability to our strengthened version, Theorem 3, does not immediately follow and would require techniques from measure concentration. This falls outside the scope of the current work, but we expect that this may go through if the algorithm in [63] satisfies certain assumptions. Readers willing to accept the existence of the relevant algorithm — i.e., an efficient algorithm for reconstructing the causal wedge — can skip directly to the statement of Cryptographic Censorship, Theorem 5 below.

Efficiently learning the causal wedge. The starting point for our argument for (b) is the learning algorithm of [63], which we review below, together with the assumption referenced above that allows us to bridge the gap between [63]’s work and our Theorem 2.

Given an n -qubit system with n large, the authors of [63] showed that there exists an efficient learning algorithm that can learn a *bounded degree* representation of any unknown




Figure 8. The region of spacetime inside this cylindrical naked singularity does not lie inside of any HKLL wedge. It is possible, however, for a rocket ship to enter the region behind the singularity and come back out to the boundary.

operator $\mathcal{O}^{(\text{unk})}$. The bounded degree representation $\mathcal{O}^{(k)}$ takes the form of

$$\mathcal{O}^{(k)} = \sum_{P:|P|\leq k} \alpha_P P, \tag{3.7}$$

where P is a Pauli operator given by a string of Pauli operators X, Y, Z , or I at each site, and k is a constant (independent of n). By $|P|$ we mean the number of qubits that are acted on non-trivially by P . A bounded-degree operator is then just a sum of individual Pauli operators with $|P|$ not scaling with n in any way.

The algorithm of [63] works by forming a guess for the coefficients of the unknown operator $\mathcal{O}^{(\text{unk})}$ in the Pauli basis, using a polynomial or subexponential number of measurements involving the operator. The authors show that this guess $\mathcal{O}^{(k)}$ for the operator reproduces the correct expectation values in certain distributions of quantum states. Quantitatively:

Theorem 4 (Theorem 13 of [63]). *There exists an algorithm that, for any $n, \varepsilon, \delta > 0$, any unknown n -qubit observable $\mathcal{O}^{(\text{unk})}$ and any n -qubit state distribution \mathcal{D} that is invariant under single qubit gates, given training data $\{\rho_\ell, \text{tr}(\mathcal{O}^{(\text{unk})}\rho_\ell)\}_{\ell=1}^M$ of size*

$$M = \log\left(\frac{n}{\delta}\right) 2^{\mathcal{O}(\log(\frac{1}{\varepsilon})\log(n))}, \tag{3.8}$$

the algorithm is of $\text{poly}(M)$ complexity and can learn a function $h(\rho)$ that achieves a prediction error

$$\text{avg}_{\rho \leftarrow \mathcal{D}} \left| h(\rho) - \text{tr}(\mathcal{O}^{(\text{unk})}\rho) \right|^2 \leq \varepsilon \left\| \mathcal{O}^{(\text{unk})} \right\|^2, \tag{3.9}$$

with probability at least $1 - \delta$. Here, $\left\| \mathcal{O}^{(\text{unk})} \right\|$ is the operator norm of $\mathcal{O}^{(\text{unk})}$. In particular, the algorithm produces a bounded-degree operator $\mathcal{O}^{(k)} = \sum_{P:|P|\leq k} \alpha_P P$ and thereby learns $h(\rho) := \text{tr}(\mathcal{O}^{(k)}\rho)$.

One example of such a distribution \mathcal{D} is the distribution of Haar random states, which is of particular relevance to our setting.

In general, a bounded-degree operator is any operator of the form in (3.7). In the context of AdS/CFT, we expect that bounded-degree operators correspond to local operators in the CFT. In particular, expectation values of these operators fully determine the causal wedge of the entire boundary in the large- N limit. Recall that our aim is the illustration of an algorithm that falls under the purview of Theorem 2 and can learn a V that approximates the time-evolved expectation values of bounded-degree operators under U . Taking $\mathcal{O}^{(\text{unk})} = U$, the algorithm in the above theorem (given the specified training data) learns an operator $U^{(k)}$ such that

$$\text{tr}(U^{(k)}\rho) \approx \text{tr}(U\rho) \tag{3.10}$$

(on average over ρ). Note that this algorithm requires access to training data comprised of expectation values of the unknown operator $\mathcal{O}^{(\text{unk})} = U$ in a collection of states $\{\rho_\ell\}$. We expect that there exists an efficient algorithm to estimate the expectation values of U in this collection of efficiently preparable states given oracle access to U so that the learning algorithm in Theorem 4 then falls under the purview of Theorem 2.

In order to obtain a guarantee that this $U^{(k)}$ is such a V , we will make a mild assumption that (3.10) holds when multiple operators are concatenated. In general, we may expect that for any operators $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3$, using the learning algorithm in the theorem above to produce $\mathcal{O}_1^{(k)}, \mathcal{O}_2^{(k)}, \mathcal{O}_3^{(k)}$, the following holds:

$$\text{tr}(\mathcal{O}_1^{(k)}\mathcal{O}_2^{(k)}\mathcal{O}_3^{(k)}\rho) \approx \text{tr}(\mathcal{O}_1\mathcal{O}_2\mathcal{O}_3\rho) .$$

However, to bridge the gap to Theorem 2, it suffices that the $U^{(k)}$ produced by the algorithm satisfies the following for any bounded-degree operator \mathcal{O} :

$$\text{tr}((U^{(k)})^\dagger\mathcal{O}U^{(k)}\rho) \approx \text{tr}(U^\dagger\mathcal{O}U\rho) , \tag{3.11}$$

where the error scales with $1/N$ if the algorithm complexity is subexponential in N , which is a complexity to which Theorem 2 (and its strengthening Theorem 3) can be applied.²⁸ The upshot is that under the assumption (3.11) above, there is an efficient algorithm that obeys the assumptions of Theorem 2²⁹ and outputs a guess $V = U^{(k)}$ for the time evolution unitary U , which correctly predicts expectation values of local operators in the CFT, up to polynomial corrections in $1/N$.

To summarize: we use the above (and expectations from holography) to motivate that there exists an algorithm, which we will call \mathcal{A}_{CW} , that can always learn time evolution within the causal wedge but falls under the purview of Theorem 2, which guarantees that for pseudorandom time evolution, this algorithm cannot learn the time evolution of *all* operators. This forms the basis of our theorem of Cryptographic Censorship. Under the assumptions

²⁸Note that although in previous sections we have considered polynomial complexity algorithms, we expect our results there to also apply to algorithms of subexponential complexity.

²⁹Note that strictly speaking, Theorem 2 only applies to this algorithm if V is approximately norm-preserving. By taking e.g. $\mathcal{O} = I$ in (3.11) (which holds for most ρ in the Haar distribution), we expect that the guess V should satisfy this condition.




Figure 9. By assumption, the time evolution operator $U_{\text{fund},N} = e^{-iH_N \Delta t}$ on $\mathcal{H}_{\text{code}}$ over a time Δt is well approximated by some $U_N \leftarrow \mathcal{U}_N$, such that $|\psi_N(t')\rangle$ is typical in the PRS ensemble $\Psi = \mathcal{U}_N |\psi_N(t)\rangle$. In this case, Theorem 5 guarantees the existence of a distinguishing operator Q_* that has support outside the causal wedge: there must be a horizon (indicated by the green dotted line).

above, which simply boil down to the existence of \mathcal{A}_{CW} , together with the assumptions of section 1.1, we have the following theorem:

Theorem 5 (Cryptographic Censorship). *As in Assumption 3, let $U_{\text{fund},N} = V_N U_{\text{bulk}} V_N^\dagger$ where U_{bulk} is the time evolution operator of the bulk quantum fields acting on \mathcal{H} . Furthermore, assume that $U_{\text{fund},N}$ is well approximated on $V_N \mathcal{H} \subseteq \mathcal{H}_N$ by a unitary U_N which is typical in a gravitationally pseudorandom unitary ensemble, \mathcal{U}_N . Assume also (as motivated above) that there exists an efficient algorithm for causal wedge reconstruction that is (randomized) K -Lipschitz for some $K > 0$, and which outputs (approximately) norm-preserving operators. Let $|\psi_N(t)\rangle$ be a state at any time t whose complete time evolution is dual to a strongly causal geometric bulk (M, g) .³⁰ Under the assumption that $U_N |\psi_N(t)\rangle$ is itself the image of a bulk state $|\psi(t')\rangle$ under V_N , then if there exists a single time t such that $U_N |\psi_N(t)\rangle$ is typical in the pseudorandom state ensemble $\Psi = \mathcal{U}_N |\psi_N(t)\rangle$,³¹ then there is an event horizon in (M, g) .*

Proof. Let U_N be the pseudorandom unitary approximating $V_N e^{-iH_{\text{bulk}} \Delta t} V_N^\dagger$ over a time Δt . Let the fundamental quantum state of the system be $|\psi_N(t)\rangle$ at some initial time t , and $|\psi_N(t')\rangle = U_N |\psi_N(t)\rangle$ at some time $t' = t + \Delta t$. Let the bulk duals be $|\psi(t)\rangle$ and $|\psi(t')\rangle$ such that $|\psi_N(t)\rangle = V_N |\psi(t)\rangle$ and $|\psi_N(t')\rangle = V_N |\psi(t')\rangle$. See figure 9. By assumption, we can find a t such that $|\psi_N(t')\rangle$ is typical in the PRS ensemble $\Psi = \mathcal{U}_N |\psi_N(t)\rangle$.

Consider a particular (possibly randomized) algorithm, $\mathcal{A}_{CW}^{U_N}$, which by the reasoning and assumptions above based on Theorem 4 we can choose to correctly predict the values of

³⁰When we say that $|\psi_N(t)\rangle$ is dual to a geometric state, we have in mind the definition in Ass. 2.

³¹Recall that the typical states constitute a $\lim_{N \rightarrow \infty} 1 - e^{-\epsilon^2 2^\kappa / (48(1+\kappa)^2)} - e^{-\epsilon^2 2^\kappa / 192} - \text{negl}'(N) = 1$ (recall $\kappa = \text{poly}(N)$) fraction of states in Ψ , and similarly for typical $U \in \mathcal{U}$.

all operators \mathcal{O} that are in the causal wedge algebra \mathcal{N}_{cw} for any output $\widehat{U}_N \leftarrow \mathcal{A}_{CW}^U$. By Corollary 3.1, on average over the \widehat{U}_N there exists a distinguishing bulk operator, $Q_* \in \mathcal{B}(\mathcal{H})$, which distinguishes between the bulk states $|\psi(t')\rangle$ and $|\widehat{\psi}(t')\rangle$ such that $U_N |\psi_N(t)\rangle = V_N |\psi(t')\rangle$ and $\widehat{U}_N |\psi_N(t)\rangle = V_N |\widehat{\psi}(t')\rangle$. Thus there must exist at least one \widehat{U}_N for which the corresponding Q_* distinguishes between $|\psi(t')\rangle$ and $|\widehat{\psi}(t')\rangle$. Note that since $\mathcal{B}(\mathcal{H})$ is the set of all bulk quantum field theory operators, we are guaranteed that Q_* is itself a good semiclassical operator. Since, by construction, \mathcal{A}_{CW}^U can correctly predict all causal wedge operators, the operator Q_* must have support in a bulk spacetime region that is outside the causal wedge. Thus a horizon exists. \square

We end this section by pointing out that we in fact get a more general result: the proof of Theorem 5 goes through for any time evolution operator U that cannot be predicted (in the sense of Theorem 3) by efficient quantum algorithms. In particular, the diagnostic does not have to be pseudorandomness: our theorem holds for any U that is exponentially complex to learn.

3.4 Possible subspaces with pseudorandom time evolution

We now discuss possible code subspaces which accommodate gravitationally pseudorandom dynamics. We will outline two cases of interest. First, we consider a microcanonical window about some energy above the black hole threshold. Second, we will argue that there can exist a code subspace where all the states have the same classical geometry, with low energy modes on top of it, and for which it is possible to obtain gravitational pseudorandomness within an $O(1)$ time on the boundary.

The microcanonical window.

An obvious candidate for a subspace in a holographic quantum system which admits pseudorandom time evolution is a microcanonical window about some energy $E_0 \sim N^2$ with energy width $\Delta E \sim O(N)$. The microcanonical ensemble defined by this window is known to have a stationary two-sided black hole dual [64]. Here we are interested in applying our result to any state in this energy window with a semiclassical geometric bulk description.

For holographic theories that exhibit maximal chaos, it is an interesting question how large T needs to be for $U(T) = e^{-iHT}$ to be well approximated by a pseudorandom unitary ensemble. A related (but different) question is how large T needs to be in a maximally chaotic system for the unitary to be well approximated by a unitary k -design. For $k = 2$, this question was answered in [65] by directly relating out-of-time-order (OTO) correlators to the question of how close a given ensemble is to a unitary 2-design. The authors found that if the (square of the) 2-point OTO correlator is small when averaged over all possible operators, then the time evolution is well modeled by a 2-design.³²

As discussed above, k -designs are a different but related notion from pseudorandom unitary ensembles. That said, this intuition suggests that the time T at which $U(T)$ becomes well modeled by a pseudorandom unitary is at least the scrambling time for chaotic quantum

³²The authors proved this statement for an n -qubit system and the average is over all Pauli operators on this n -qubit Hilbert space, which provide a basis for all operators.

systems, since this is the time when the OTO is expected to match those of a 2-design. Note that if $U(T)$ is ever going to be modeled by a pseudorandom unitary, it better be at times T sub-exponential in the entropy, otherwise U will not be efficiently generated. Thus, the expectation is that $U(T)$ should become pseudorandom on the microcanonical subspace for sufficiently large T larger than the scrambling time. It would of course be nice to know what exactly this timescale is in certain systems.

Applying Theorem 5 to this subspace, we then see that for any state in the microcanonical window, $|\psi_0\rangle$, that is both dual to a geometric state and is also typical in the pseudorandom state ensemble, then *with a probability that is negligibly (in N) close to one there must exist a horizon in the bulk dual to $|\psi_0\rangle$* . It is natural to ask what *fraction* of the states in the microcanonical window this corresponds to. We cannot give a definitive answer to this question in the case of pseudorandom ensembles; let us just note that for Haar random unitaries and states, this would include all but a doubly-exponentially small fraction of the states. However, since it is not clear whether typical Haar random states can have semiclassical geometric duals,³³ we prefer to use pseudorandom ensembles. It is not implausible that a similar fraction of the states in some pseudorandom ensemble would obey this property, but we shall remain agnostic on this point in the absence of a proof of this statement.

A code subspace for $O(1)$ time bands on the boundary.

We will now consider a different subspace, where we restrict to $O(1)$ time bands on the boundary in order to preserve the geometry of states in our code subspace within exponentially suppressed errors. We will start by making some assumptions that will apply to what follows, as well as in section 4. These assumptions will describe how we construct a code subspace to describe time evolution of probe fields for $O(1)$ times around geometries with naked singularities.

Consider a CFT labeled by a parameter N , such that the theory in the $N \rightarrow \infty$ limit is dual to Einstein gravity coupled to matter.³⁴ By this we mean that the theory participates in an asymptotically isometric code given by $V_N : \mathcal{H} \rightarrow \hat{\mathcal{H}}_N$, where \mathcal{H} is the bulk Hilbert space of Einstein gravity perturbatively coupled to matter and $\hat{\mathcal{H}}_N$ is the finite- N Hilbert space of the CFT. We are now going to consider a finite-dimensional subspace \mathcal{H}_N of $\hat{\mathcal{H}}_N$ that we will assume to have certain properties. This will formalize assumptions often implicitly made in the study of AdS/CFT, but since we will (1) study somewhat exotic setups, and (2) apply certain rigorous theorems, we are careful to state our assumptions.

Take $|\psi_N^i\rangle$ to be a basis of \mathcal{H}_N labeled by i . We will assume that there exists a one-parameter family of unitary operators

$$U_N(t) : \mathcal{H}_N \rightarrow \mathcal{H}_N \tag{3.12}$$

such that for all $t \sim O(N^0)$,

$$V_N |\psi^i(t)\rangle = U_N(t) |\psi_N^i\rangle \tag{3.13}$$

³³Note that if the fraction of geometric states which violate our theorem is much larger than doubly exponentially small, then there would need to be a conspiracy between having a geometry and the ability to reconstruct all operators. Since we expect spacetimes with horizons to be much more “numerous” than those without, this would be rather surprising.

³⁴We always assume that any effective 't Hooft coupling in the theory has been taken to infinity.

with $|\psi^i(t)\rangle \in \mathcal{H}$ a bulk state, time-evolved using the bulk Hamiltonian. In other words, $|\psi^i(t)\rangle$ is a good bulk semiclassical state. Next, we will assume that if $e^{-iH_N t}$ is the true time evolution operator in the full CFT at finite N , then for any fixed $t \sim O(N^0)$ and any state $|\varphi\rangle \in \widehat{\mathcal{H}}_N$,

$$\langle \varphi | e^{-iH_N t} | \psi_N^i \rangle = \langle \varphi | U_N(t) | \psi_N^i \rangle + O(e^{-cG^a}) \tag{3.14}$$

for some $c > 0, a < 0$. For $|\varphi\rangle$ in the orthogonal complement of \mathcal{H}_N , the first term is identically zero, so our statement implies that at large N , CFT time evolution for an $O(1)$ amount of time preserves our code subspace \mathcal{H}_N , up to exponentially small corrections. This is just the usual intuition that different geometries have exponentially small overlaps, so unless we work to exponential precision, we can consider time evolution to preserve our code subspace. Instead of working with e^{-iHt} , we replace this operator by a unitary that genuinely acts on \mathcal{H}_N .

Then, under the above assumptions, Theorem 5 implies the following. Consider an asymptotically AdS spacetime (M, g) where the boundary dynamics in an $O(1)$ time band evolve according to a unitary $U_{\text{fund}} = e^{-iH_N t}$ and as in the above, can be approximated by $U_N(t)$ which is a typical gravitationally pseudorandom unitary (as defined in Def. 4). Let the state dual to (M, g) be typical in a pseudorandom state ensemble. Then there must be an event horizon in (M, g) .

This is the version of Cryptographic Censorship that we will use in the next section.

4 From Cryptographic to cosmic censorship

In section 3.4, we adapted our main result — Cryptographic Censorship — to code subspaces of states that have the same classical geometry and undergo gravitationally pseudorandom boundary dynamics in an $O(1)$ time band. In this section we will argue that a potential example of such a code subspace is a spacetime with a ‘large’ naked singularity. Our goal is to connect Cryptographic Censorship to the more familiar notion of cosmic censorship, which is roughly the expectation that naked singularities formed from collapse must be hidden behind horizons.

Cryptographic Censorship guarantees that singularities that are well-described by gravitationally pseudorandom dynamics must incur event horizons. The central task in connecting Cryptographic Censorship to cosmic censorship is thus understanding whether — or when — singularities are associated with (sufficiently) pseudorandom dynamics. If generic singularities formed from matter collapse do indeed require pseudorandom time evolution, then they cannot exist in horizonless spacetimes. This version of censorship of naked singularities is distinct from the original version of cosmic censorship: even when a singularity is associated with pseudorandom dynamics, our result does not imply that the singularity is completely hidden. This is not a bug but a feature, since certain naked singularities do exist: e.g., the Gregory-Laflamme (GL) instability, which occurs when a black string in $D > 4$ pinches off and fragments into multiple black holes [6, 7, 66].

In the following, we take some first steps towards connecting a version of cosmic censorship to Cryptographic Censorship, culminating in section 4.3, where we state a cryptographically-inspired formulation of quantum cosmic censorship. We begin by defining a mathematically

precise notion of the size of a singularity, and specifically identify how large the singularity must be in the $G_N \rightarrow 0$ limit to accommodate gravitationally pseudorandom fundamental dynamics. This definition usefully distinguishes between a Planckian and a classical naked singularity (and similarly for non-naked singularities of large and small sizes). This is in itself useful, as it separates naked singularities into classes that are expected to be excluded and those that we expect may be permitted in a quantum theory of gravity (e.g., evaporating black holes or Gregory-Laflamme instabilities).³⁵ Our second step uses this quantification to provide some evidence that large classical singularities can indeed be associated with gravitationally pseudorandom dynamics and thus incur horizons by Cryptographic Censorship. Finally, we propose a Quantum Weak Cosmic Censorship Conjecture, which formally requires large naked singularities to be atypical.

4.1 The size of singularities

We are interested in an unambiguous and diffeomorphism-invariant way of measuring the extent of a naked singularity. Since a naked singularity can only be distinguished from a cloaked one via access to \mathcal{I} , it is clear that this must be defined asymptotically. A definition on \mathcal{I} would be automatically diffeomorphism-invariant and would more readily admit a straightforward dual CFT interpretation. To find this criterion, we will leverage the definition of a naked singularity — i.e., causal visibility to \mathcal{I} — to define what will roughly correspond to the time evolution in the CFT that is necessary, in the Schrödinger picture, to reconstruct near-singularity physics.

To make this precise, we first use an extant mathematical construction to single out the singularity as a set, a nontrivial task given that the singularity is not part of the spacetime manifold. Second, we identify the time intervals on \mathcal{I} that are causally-separated from this set, and finally find the minimal boundary time bands that are necessary for a full reconstruction of the singularity. The spacetime volume of these time bands in a fixed conformal frame is essentially our definition of the size of the singularity in that frame. There is thus a conformal frame ambiguity in our definition of singularity size. This, however, is immaterial for our desired application: we shall only care about scaling of the singularity size with G_N , which is independent of the choice of conformal frame. Using this scaling, we separate the singularities into three categories: classical, semi-Planckian, and Planckian. The first type survives the $G_N \rightarrow 0$ limit, while the other two do not. These latter two sizes are distinguished via the power of G_N by which they vanish. Let us now turn to the details, starting with an identification of the set comprising the singularity.

The singularity set. The primary difficulty in this task is that singularities are not locations in some geodesically-incomplete manifold M but are rather associated to the endpoints of incomplete inextendible geodesics in M . However, as shown by [67], it is possible to treat singularities as if they are made up of a set of points Γ , although this set does not lie in M .³⁶ We will briefly review this construction, as it is instrumental to our classification of

³⁵There is some extant literature dedicated to investigating the size of a singularity; see, e.g., [24]. These approaches differ from ours, as we are motivated by quantum computational considerations in the fundamental quantum description of the system.

³⁶ Γ also comes with a natural choice of topology [67].

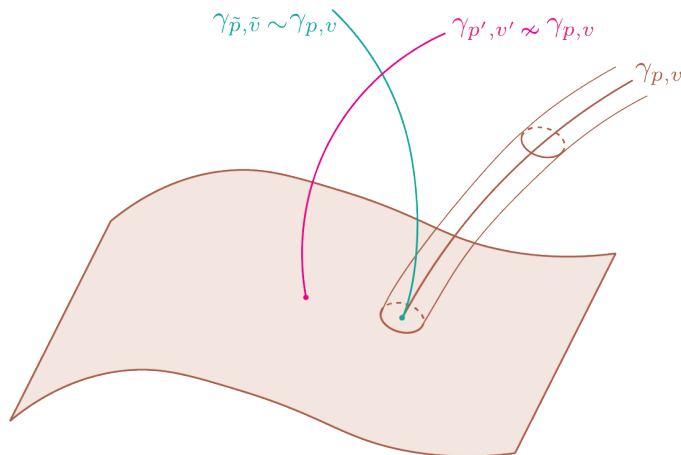


Figure 10. Example spacetime where the colored sheet represents a singularity. The two geodesics $\gamma_{p,v}$ and $\gamma_{\bar{p},\bar{v}}$ by definition end on the same singular point, since $\gamma_{\bar{p},\bar{v}}$ must enter any tubular neighborhood of $\gamma_{p,v}$, and vice versa.

singularities. The idea is to build an equivalence class of geodesics associated to an endpoint. Consider a geodesic $\gamma_{p,v}$ in M fired from a starting point $p \in M$ along a non-vanishing tangent vector v^a at p . We will continue $\gamma_{p,v}$ so it is inextendible in the direction of v^a (with starting point p). We then identify two distinct geodesics $\gamma_{p,v} \sim \gamma_{\bar{p},\bar{v}}$ if $\gamma_{p,v}$ enters and remains inside every tubular neighborhood of $\gamma_{\bar{p},\bar{v}}$. Explicitly, if $\lambda \in [0, \lambda_0)$ is an affine parameter for $\gamma_{p,v}$, then for every given tubular neighborhood U of $\gamma_{\bar{p},\bar{v}}$, there exist a $0 \leq \hat{\lambda} < \lambda_0$ such that $\gamma_{p,v}(\lambda) \in U$ for all $\lambda \in (\hat{\lambda}, \lambda_0)$. See figure 10. The equivalence classes under \sim then form the points in Γ , and we say that $\gamma_{p,v}, \gamma_{\bar{p},\bar{v}}$ are geodesics that end on the singular point $[\gamma_{p,v}] \in \Gamma$. Importantly, if $p_s \in \Gamma$, the future $I^+(p_s)$ and past $I^-(p_s)$ of p_s in M is well-defined even though p_s is not in M . Namely, we say that $p \in I^+(p_s)$ if there is a future-directed timelike curve γ through p and with past “endpoint” on p_s , where the latter means that γ , sufficiently far in its past direction, enters and remains within any tubular neighborhood of any representative geodesic in p_s . Armed with a set identified as the singularity, we now extract the part of the singularity that is naked with respect to a given conformal boundary \mathcal{S} as follows:

$$\Gamma_{n(\mathcal{S})} = \{p_s \in \Gamma \mid I^+(p_s) \cap \mathcal{S} \neq \emptyset \text{ and } I^-(p_s) \cap \mathcal{S} \neq \emptyset\}. \quad (4.1)$$

If this is nonempty, we say that Γ is naked with respect to \mathcal{S} .

The singularity size. We now proceed to classify the size of any (possibly improper) subset of a singularity. Let $\hat{\Gamma}$ be a subset of Γ ; below we will mostly be interested in taking $\hat{\Gamma} = \Gamma_{n(\mathcal{S})}$, but it will be useful to take our definition to be more general. Consider the sets

$$\begin{aligned} T_+ &= \{p \in \mathcal{S} \mid \exists \text{ an achronal past-directed null geodesic from } p \text{ to } \hat{\Gamma}\}, \\ T_- &= \{p \in \mathcal{S} \mid \exists \text{ an achronal future-directed null geodesic from } p \text{ to } \hat{\Gamma}\}. \end{aligned} \quad (4.2)$$

Figure 11 illustrates some examples of these sets for both naked singularities and ordinary singularities behind horizons.




Figure 11. Some examples of spacetimes with (naked) singularities. In (a) the singularity is timelike and classically naked. Panel (b) depicts the evaporating black hole. The non-naked part of the singularity is classical and behind a horizon. For the naked part, see figure 13(a). Panel (c) is Reissner-Nordström-AdS; here $|\Gamma|_{\mathcal{I}} = |T_-|$; note, however, that $\Gamma_{n(\mathcal{I})} = \emptyset$, as $I^\pm \cap \mathcal{I} = \emptyset$.

We will now define the size of $\hat{\Gamma}$ according to properties of T_+ , T_- . We assume that (M, g) is conformal to the Einstein static universe on each conformal boundary, and we choose the unique conformal frame where the boundary metric is identical to the Einstein static universe in which the spatial spheres have radius equal to the AdS radius.³⁷ This fixes a unique Lorentzian metric on \mathcal{I} , and when T_+ is nonempty, we define $|T_+|$ as the absolute value of the spacetime volume of T_+ in this metric; similarly for T_- . As examples, if T_+ is a timeslice or a union of a finite number of points, then $|T_+| = 0$, since these sets are measure zero. We can now define the size $|\hat{\Gamma}|_{\mathcal{I}}$ of $\hat{\Gamma}$ with respect to \mathcal{I} as follows: first, if T_+, T_- are both empty, then $|\hat{\Gamma}|_{\mathcal{I}} = 0$. Second, if only one of T_+, T_- is nonempty, say T_+ , then we define $|\hat{\Gamma}|_{\mathcal{I}} = |T_+|$. Third, if both T_\pm are nonempty, we define $|\hat{\Gamma}|_{\mathcal{I}} = \min(|T_+|, |T_-|)$. This finally leads us to our definition of classical and Planckian singularities, as well as an in-between regime, which we term semi-Planckian:

Definition 5. Let (M, g) be a geodesically incomplete spacetime of dimension $D = d + 1$ with a singularity $\hat{\Gamma}$, and let

$$|\hat{\Gamma}|_{\mathcal{I}} \sim O(\ell_{\text{Pl}}^a). \tag{4.3}$$

If $a \leq 0$, we say that $\hat{\Gamma}$ is a classical singularity, while if $0 < a < d - 1$, we say that $\hat{\Gamma}$ is a semi-Planckian singularity. Otherwise, it is a Planckian singularity.

Singularity size from CFT time evolution. If $|\psi(t)\rangle$ is the state whose complete past and future history is dual to (M, g) , then without access to the time evolution operator, we can only reconstruct operators at points in spacetime that are not causally related to the

³⁷For asymptotically locally AdS spacetimes where \mathcal{I} has another conformal class, we can also define $|\hat{\Gamma}|_{\mathcal{I}}$, once we pick some preferred conformal frame.

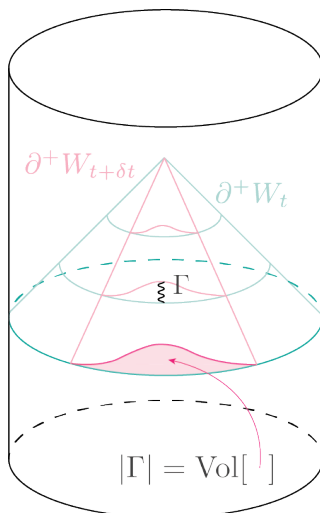


Figure 12. To get all information about the small singularity Γ , we only need to evolve the state forwards in a small spatial subregion of the boundary Cauchy slice, such that the Wheeler-de Witt patch now includes the singularity. Note that, as our artistic capabilities are limited, we here only indicate the future boundaries of the Wheeler-de Witt patches for times t and $t + \delta t$.

boundary time slice on which $|\psi(t)\rangle$ lives — i.e., the Wheeler-de Witt patch W_t .³⁸ Let us now specialize to a spacetime with a naked singularity Γ . For some fixed t , generally only a (possibly empty) subset of Γ will be contained in W_t : to learn everything about Γ (for example by reconstruction of all near-singularity correlators), a certain amount of boundary time evolution is necessary. More precisely, if T_- is nonempty, and we are evolving forwards from the past, then the past-most boundary of T_- corresponds to the earliest time that some part of Γ is encoded in the Wheeler-de Witt patch W_t . As we evolve forward through T_- , new pieces of Γ enter W_t and are encoded in the time-evolved state. Finally, once we evolve past the futuremost boundary of T_- , no new parts of Γ enter W_t that we have not already seen, and so further forward time evolution reveals nothing new about the singularity. Thus, everything about the naked singularity can be learned by carrying out CFT time evolution through the band T_- (plus bulk reconstruction not using the time-evolution operator). For a singularity where T_+ is also nonempty, everything about the singularity could alternatively have been learned by evolving through T_+ . So as advertised, $|\Gamma|_{\mathcal{S}} = \min\{|T_-|, |T_+|\}$ measures the minimal amount of CFT time evolution required to learn everything about the naked singularity. Note that $|\Gamma|_{\mathcal{S}}$ has units of spacetime volume rather than time. This is a feature: if all information about the singularity is available upon evolution of the state forwards in a small spatial subregion of a boundary Cauchy slice, then this should be accounted for in the smallness of a singularity. See figure 12. Note also that if $|\Gamma|_{\mathcal{S}}$ is large, this always is caused by a large amount of time evolution, since timeslices have bounded volume.

Let us now make a few remarks on our definition, starting with a question. Under what

³⁸This assumes for simplicity that there is only one asymptotic boundary, and that the state of the CFT on that boundary is pure — so the entanglement wedge of $|\psi\rangle$ is the entire bulk. If there are multiple boundaries or the state on one boundary is mixed for a different reason, we simply take the intersection of W_t with the entanglement wedge of $|\psi\rangle$.




Figure 13. Two examples of $\lim_{\ell_{\text{Pl}} \rightarrow 0} |\hat{\Gamma}|_{\mathcal{I}} = 0$: (a) For the naked part of the singularity at the endpoint of black hole evaporation, we expect that T_{\pm} are just timeslices in the $G_N \rightarrow 0$ limit. (b) A ‘hidden’ naked timelike singularity; here, $T_{\pm} = \emptyset$.

circumstances do Planckian or semi-Planckian singularities exist? The most natural case is that we have $\lim_{\ell_{\text{Pl}} \rightarrow 0} |\hat{\Gamma}|_{\mathcal{I}} = 0$ because the CFT state acquires all information about the singularity in a CFT timestep that shrinks to zero as $\ell_{\text{Pl}} \rightarrow 0$. An example would be the singularity in the evaporating black hole, illustrated in figure 13(a), or a naked timelike singularity $\hat{\Gamma}$ lying behind another naked timelike singularity $\tilde{\Gamma}$, like in figure 13(b). In this latter case, $\hat{\Gamma}$ is such that $T_+ = T_- = \emptyset$, since all causal geodesics from the boundary to $\hat{\Gamma}$ are necessarily chronal. This latter example shows that we probably should not assign much meaning to the size of distinct connected components, and that we instead just should look at the size of the full naked singularity $\Gamma_{n(\mathcal{I})}$, which by definition contains all connected components.

4.2 Accommodating gravitational pseudorandomness

Our motivation for classifying a classical singularity via the scaling $a \leq 0$ is simple: such singularities are visible for a boundary time that diverges or remains $O(1)$ in the $G_N \rightarrow 0$ limit. For the remaining range of a , we have chosen to separate semi-Planckian singularities (with $0 < a < d - 1$) from Planckian singularities (with $a \geq d - 1$). In both cases, $|\hat{\Gamma}|_{\mathcal{I}} \rightarrow 0$ as $G_N \rightarrow 0$, but as we will show below, singularities with $a < d - 1$ are compatible with an EFT code subspace which is sufficiently large that it can accommodate gravitational pseudorandomness, while singularities with $a \geq d - 1$ are not.

We will now show that singularities with $a < d - 1$ are compatible with gravitational pseudorandomness. Consider an asymptotically AdS spacetime with a timelike naked singularity $\Gamma = \Gamma_{n(\mathcal{I})}$. We will work with large (but possibly finite) N and demand only that effective field theory as defined by a bulk UV cutoff be valid in the asymptotic region. At finite N we will consider the failure of EFT due to large UV corrections as the diagnostic of a singularity.




Figure 14. From the low-energy perspective, there is an unknown, UV-sensitive boundary condition assigned at the singularity.

We will decompose our bulk state into modes with energy below the cutoff, and we will focus on modes localized near the boundary, some of which fall into the naked singularity. By unitarity, there must also be outgoing occupied modes as well, which can be traced back to the singularity. However, the relation between the modes falling into the singularity and the modes coming out is dependent on UV physics: it is inaccessible within low energy EFT. So while the low energy modes in the asymptotic region could be semiclassical as $G_N \rightarrow 0$, the time evolution operator that relates them is not known in the bulk EFT — it cannot be expressed in terms of the low-energy operators. See figure 14. From the low-energy perspective, there is a boundary condition assigned at the singularity whose precise nature is unknown.³⁹

In what follows, we will be agnostic about the UV-sensitive boundary condition that implements pseudorandomness in the fundamental theory. Our approach instead will be to show that once the singularity is of classical or semi-Planckian size, then there are enough EFT modes in a code subspace in which all of the states have the same leading order geometry to accommodate gravitational pseudorandomness. On the other hand, if the singularity is Planckian, we show that we do not expect this. As our analysis is asymptotic, we do not require a sharply defined geometry near the singularity.

We now give a parametric estimate of the number of modes within semiclassical EFT that can potentially get scrambled by the naked timelike singularity. That is, we want to know the number of EFT modes that are affected by the highly UV-sensitive boundary condition. Assume that T_+, T_- are nonempty, and assume without loss of generality $|\Gamma_{n(\mathcal{S})}|_{\mathcal{S}} = |T_-|$. Near \mathcal{S} , the leading part of the metric can be written

$$ds^2 = \frac{-dt^2 + d\rho^2 + \sin^2 \rho d\Omega^2}{\cos^2 \rho} + \dots, \tag{4.4}$$

where we work in units where the AdS radius is unity, and where \mathcal{S} lies at $\rho = \frac{\pi}{2}$. For pedagogical simplicity we take the matter in our theory to be a conformal massless scalar

³⁹Since quantum gravity is expected to lack global symmetries [68–70], any approximate symmetries emergent in the semiclassical EFT should be badly violated by this boundary condition.




Figure 15. To count the ingoing modes sent from T_- , which here is the time band $[-\Delta t, 0] \times S^{d-1}$, we can simply count the number of modes in the strip indicated by the red dotted line.

and assume that the system is spherically symmetric, so that T_- can be taken to be a time band $[-\Delta t, 0] \times S^{d-1}$ without loss of generality. We can now count roughly how many ingoing single-particle modes with energy below the cutoff can be sent from T_- into the singularity. As is clear from figure 15, this simply corresponds to the number of orthogonal modes in the strip given by $\rho \in [\frac{\pi}{2} - \Delta t, \frac{\pi}{2}]$ at $t = 0$.⁴⁰ As long as Δt is less than roughly $G_N M \sim O(1)$ where M is the spacetime mass, we are always restricted to an asymptotic range of ρ where the metric is approximately pure AdS, so the number of modes can be counted using the pure AdS metric. We always take G_N small with $G_N M = \text{fixed}$, and we are interested in a parametric counting that can only change if $\Delta t \rightarrow 0$ sufficiently fast in G_N so we can without loss of generality restrict to this regime. Assume now that \tilde{u}_i is a collection of orthonormal flat-space positive frequency ingoing wavepackets with mean frequency ω_i in the metric $-dt^2 + d\rho^2 + \sin^2 \rho d\Omega^2$; furthermore, we will assume that they are well localized within $\rho \in (\frac{\pi}{2} - \Delta t, \frac{\pi}{2})$ at some central $\rho = \rho_0$. These have total energy approximately equal to ω_i , as long as

$$\omega_i \gg \Delta\omega \geq \frac{1}{\Delta t}, \tag{4.5}$$

where $\Delta\omega$ is the frequency spread of each wavepacket. Now, since we have a conformally coupled scalar, $u_i = (\cos \rho)^{d-1} \tilde{u}_i$ is also an orthonormal set of positive frequency mode functions in AdS.⁴¹ Each wavepacket has a total energy approximately equal to ω_i , also in AdS.⁴² Imposing the cutoff $\omega_i \leq \epsilon^{-1}$,⁴³ we want to count the orthogonal wavepackets. This is

⁴⁰Greybody factors cause some modes to scatter off background curvature before they get close to the singularity, but this should not change the parametric counting.

⁴¹The scaling dimension of the scalar is $\frac{d-1}{2}$, and the conformal factor is $\cos^2 \rho$.

⁴²There might be an overall conformal anomaly, so that $H_{\text{AdS}} = H_{\text{Mink}} + H_{\text{anomaly}}$, but this is irrelevant for our purposes.

⁴³This corresponds to a proper length cutoff in AdS of $\epsilon_{\text{proper}} \sim \epsilon / \cos \rho_0$. We do not scale ρ_0 with ℓ_{Pl} , so $\epsilon, \epsilon_{\text{proper}}$ scale the same way.

equivalent to counting orthogonal flat-space modes; this is of the order

$$\frac{\Delta t \cdot \text{Vol}(S^{d-1})}{\epsilon^d} . \tag{4.6}$$

The numerator is $|\Gamma_{n(\mathcal{S})}|_{\mathcal{S}}$, and in the absence of spherical symmetry it is clear that $|\Gamma_{n(\mathcal{S})}|_{\mathcal{S}}$ is the correct replacement for $\Delta t \cdot \text{Vol}(S^{d-1})$. So for a general naked timelike singularity, we find that an order

$$\frac{|\Gamma_{n(\mathcal{S})}|_{\mathcal{S}}}{\epsilon^d} \tag{4.7}$$

number of distinct EFT modes can propagate into the singularity from T_- . Next, we want to take the UV cutoff to scale as $\epsilon = \ell_{\text{Pl}}^{1-\delta}$ for some $\delta \in (0, 1)$ so that our modes are parametrically below the Planck scale as $\ell_{\text{Pl}} \rightarrow 0$. To be cautious, we demand that $\delta > 1/d$, which ensures that the naïve QFT entropy does not exceed the Bekenstein-Hawking entropy. To see this, note that if V is the spatial volume occupied by our wavepackets in AdS, then we need

$$\frac{V}{\epsilon_{\text{proper}}^d} = \cos^d \rho_0 \frac{V}{\epsilon^d} \lesssim \frac{A}{\ell_{\text{Pl}}^{d-1}} . \tag{4.8}$$

Remembering that volume and surface area of a large shell near the boundary of AdS scale the same way as we approach the conformal boundary, we thus need

$$\cos^d \rho_0 < O(1) \frac{\epsilon^d}{\ell_{\text{Pl}}^{d-1}} = O(1) \ell_{\text{Pl}}^{1-d\delta} . \tag{4.9}$$

If we take $\delta < \frac{1}{d}$, then wavepackets with mean frequency near the cutoff would give a naïve QFT entropy exceeding the Bekenstein-Hawking entropy for any $\rho_0 < \pi/2$. For $\delta = \frac{1}{d}$ we get an $O(1)$ bound on ρ_0 that we want to avoid, and so we require $\delta > \frac{1}{d}$. Thus, assuming finally that $|\Gamma_{n(\mathcal{S})}|_{\mathcal{S}} \sim O(\ell_{\text{Pl}}^a)$, we get that the number of distinct semiclassical modes that can probe the naked singularity during time evolution through T_- is

$$\frac{|\Gamma_{n(\mathcal{S})}|_{\mathcal{S}}}{\epsilon^d} \sim \ell_{\text{Pl}}^{a-d-d\delta} . \tag{4.10}$$

If $a < d - 1$, there is always some $\delta > 1/d$ so that this diverges in the $\ell_{\text{Pl}} \rightarrow 0$ limit. Hence, in the regime $0 < a < d - 1$, even when the singularity shrinks to zero size as $\ell_{\text{Pl}} \rightarrow 0$, the singularity is in principle compatible with influencing an infinite number of EFT modes in the $G_N \rightarrow 0$ limit: it is compatible with a gravitational amount of pseudorandomness in the time evolution.

Let us now comment on a few well known cases of naked singularities, some of which were briefly mentioned in section 1: the singularity at the endpoint of black hole evaporation, the Gregory-Laflamme (GL) instability [6, 66], which occurs as the pinch-off endpoint of the evolution of higher dimensional extended black strings and other black objects, and Choptuik’s critical collapse [71, 72] resulting from a collapsing massless scalar field which is tuned to lie exactly at the threshold of black hole formation. In the case of the endpoint of black hole evaporation, the naked part of the singularity is not extended in the $G_N \rightarrow 0$ limit. See figure 13(a). Thus the naked part of the singularity is measure zero in this

limit, so this naked singularity is either Planckian or semi-Planckian. In the cases of the Gregory-Laflamme instability and critical collapse, it is impossible to evolve past the pinch-off using classical gravity, so in the absence of a quantum gravitational description of the system, we cannot ascertain the nature of the naked singularity. However, it is quite plausible that these singularities end up being Planckian or semi-Planckian. The critical collapse singularity is in a sense a zero mass black hole, which within the resolution of semiclassical gravity is indistinguishable from a Planck mass black hole. So if we could carry out critical collapse at finite G_N in full nonperturbative quantum gravity, and then afterwards take the $G_N \rightarrow 0$ limit, it is reasonable to expect that the singularity becomes a Planckian or semi-Planckian singularity. The Gregory-Laflamme singularity could plausibly behave similarly: it is highly analogous to the fluid pinch-off phenomenon, where only a minimal amount of nonperturbative evolution is necessary to re-enter the EFT regime. It is tempting to speculate that naked singularities in a holographic theory of quantum gravity can never be classical, and possibly not even semi-Planckian. We will formalize a conjecture to this effect in the next section.

Finally, let us briefly tie this discussion back to our primary theorem of Cryptographic Censorship. The fact that classical singularities are always compatible (by our definition) with gravitational pseudorandomness provides additional motivation for the intuition that the dynamics of strongly gravitating systems in the $G_N \rightarrow 0$ limit may admit a description in terms of pseudorandom unitaries. Under this supposition, Cryptographic Censorship guarantees that any such systems — including those with naked singularities — must also have event horizons. We leave the quantification of the subset of the singularity that is necessarily behind a horizon to future work.

We wish to make clear that we do not claim that all classical singularities must act as sources of pseudorandomness; a clear counterexample is given by orbifold singularities in compact dimensions. It is the ones that weak cosmic censorship is supposed to rule out — naked classical singularities resulting from the dynamical evolution of generic, smooth initial data and matter collapse — that we expect to have, to a large extent, gravitationally pseudorandom time evolution. If the singularity is static, the mode counting argument in the previous section becomes misleading. In this case, measuring later modes scattered off the singularity does not provide information we could not have already gained from scattering earlier modes, unlike in the case where the singularity is evolving. Thus, there might be a significantly smaller amount of data that is required to describe scattering off the singularity than what the naïve mode counting argument above would suggest, making static singularities less likely candidate sources of pseudorandom dynamics.

4.3 Quantum cosmic censorship

The combination of Cryptographic Censorship and our results above regarding compatibility of naked singularities with gravitational pseudorandomness suggest a particular reformulation of the Weak Cosmic Censorship Conjecture as a *quantum gravity* conjecture rather than a classical gravity conjecture. This point bears emphasizing: there is no *ab initio* reason that classical gravity must obey cosmic censorship. Even if there had been no known counterexamples, one of the best motivations for the Weak Cosmic Censorship Conjecture is empirical: the apparent absence of naked singularities in the night sky. Cosmic censorship is

not required by consistency of known theoretical structures within General Relativity. By contrast, if naked singularities were typical in quantum gravity, then foundational results in AdS/CFT — such as the behavior of the CFT thermal state OTOC being precisely attributed to bulk event horizon dynamics — would be some remarkable conspiracy: either naked singularities engineer themselves to mimic black holes precisely in the OTOC signature, or black holes engineer themselves to mimic naked singularities. A similar coincidence would be required for various results on CFT thermalization and its bulk interpretation.⁴⁴ We therefore find that a quantum gravity formulation of cosmic censorship is *both* theoretically and empirically well-motivated; our results relating horizons to pseudorandomness provide a clear first step towards formulating weak cosmic censorship as an emergent phenomenon in terms of fundamental quantities in quantum gravity. (For example, this would require replacing the classical gravity notion of genericity — an open set in the space of solutions to the Einstein equation — with typicality.)

For completeness, we first state the classical gravity proposal for weak cosmic censorship in AdS, which is now known to be violated in classical gravity. (It is also trivially violated in quantum gravity due to, e.g., evaporating black holes, although this requires violations of classical energy conditions.)

Conjecture 1 (Classical Weak Cosmic Censorship — false). *[73, 74] Consider a geodesically complete, asymptotically AdS initial data set satisfying the dominant⁴⁵ energy condition. The maximal causal development of this initial data under the Einstein equation, together with boundary conditions at the AdS boundary, is generically a strongly asymptotically predictable asymptotically AdS spacetime.*

Reformulating this proposal in terms of a fundamental quantum description of the theory requires a complete revamping of the statement and the scope of its applicability. Motivated by our results in this paper, the existing set of counterexamples to cosmic censorship, and the top down motivation from thermalization and chaos in AdS/CFT, we propose the following formulation of the Weak Cosmic Censorship Conjecture (below we work in the Heisenberg picture, so a given state is its entire history):

Conjecture 2 (Quantum Cosmic Censorship). *Typical states with an emergent geometry in quantum gravity do not have classical or semi-Planckian naked singularities.*

5 Discussion

Cryptographic Censorship is the first precise and general result identifying concrete CFT behavior with bulk horizon formation. (While the *vox populi* associates CFT thermalization to black hole equilibration,⁴⁶ this hypothetical link is at best established in certain highly symmetric and nearly static black holes.) Our first application of the theorem guarantees

⁴⁴We thank D. Harlow for extensive discussions on this point.

⁴⁵This is the weakest reasonable statement: we could assume only the null energy condition, but counterexamples exist even when we demand the dominant energy condition.

⁴⁶Note that black hole equilibration was only recently established in four-dimensional vacuum asymptotically flat space [75] for a certain range of spins and remains an open problem in AdS.

that in the microcanonical window, under assumption of pseudorandomness, states without horizons are exponentially suppressed in probability (within their pseudorandom ensemble).

As an immediate consequence of Cryptographic Censorship, we obtain a partial resolution to the black hole teleology problem for AdS spacetimes with a CFT dual. Event horizons can be diagnosed, using Cryptographic Censorship, in finite time. The implications are significant: first, black hole mechanics in AdS/CFT may admit a legitimate thermodynamic interpretation in highly dynamical states. Second, as our primary focus in this paper, we establish a typicality result for event horizons, which is highly suggestive of a form of Quantum Cosmic Censorship. The fact that naïve models of code subspaces containing classical (or semi-Planckian) timelike singularities are compatible with a sufficient amount of gravitational pseudorandomness serves as strong motivation for our conjecture that typical states do not have naked singularities that survive the classical limit. We emphasize that our proof of Cryptographic Censorship in fact provides a guarantee for event horizon formation whenever the CFT time evolution is exponentially hard to learn; our focus in this article was on complexity of learning as a result of pseudorandomness, but any such sufficient source of complexity would result in an equivalent theorem under the same proof. Below we will elaborate on potential applications of our results and discuss some aspirational expansions thereof.

Towards a converse.

Since information escapes from the black hole interior under inclusion of quantum corrections, there is no obvious definition of a quantum event horizon (since all information eventually reaches \mathcal{S} and causal structures are only emergent in the classical limit). It is tempting to use our results to *define* quantum event horizons — and by extension quantum black holes (i.e., at finite N): states undergoing gravitationally pseudorandom time evolution with a semiclassical bulk dual in the large- N limit. A definition, however, would require a proof of a converse to Cryptographic Censorship: a statement that CFT states with a semiclassical emergent bulk featuring a horizon typically exhibit gravitational pseudorandomness. A proof of such a converse would additionally facilitate a complete resolution of the teleology problem in AdS. The tentative link between chaos, thermalization, and pseudorandomness [34, 65, 76–79] could then potentially be used to understand the horizon thermodynamics⁴⁷ of highly dynamical black holes and the statistical interpretation of the Hawking area theorem, which remains elusive [26, 81].

Gravitational pseudorandomness from a bulk quantum computer?

The reader might wonder: what if we put a quantum computer in the bulk, generating a large amount of pseudorandomness directly from dynamics of the bulk fields? Must this incur a horizon? If the quantum computer does not incur violations of EFT in the causal wedge, then our theorem applies and indeed, a gravitationally pseudorandom computer must incur a horizon. If the quantum computer generates gravitational pseudorandomness in such a way that the causal wedge is no longer described by EFT — e.g., if geometric invariants such as the volume scale inversely with G_N as in the “quantum volatile” spacetimes of [60] —

⁴⁷As opposed to the thermodynamics of the apparent horizon, which are by now well-understood [27, 56, 80].

then the assumptions of our theorem are violated, namely that the $N \rightarrow \infty$ limit is described by QFT in curved space, so a horizon is not required to form.

It is interesting that bulk processes that source a large amount of pseudorandomness (e.g., a bulk quantum computer with the power to do gravitationally pseudorandom computations) and do not violate EFT in the causal wedge must typically collapse into black holes. Prior descriptions of black hole formation typically involve energetics; to our knowledge, this is the first proof establishing a causal relation between bulk quantum computations and black hole formation (though see, e.g., [82] for a discussion of constraints on bulk quantum computations behind horizons). Interestingly, this appears to bear some relation to the proposal of [30] that effective field theory requires both an energy and a complexity cutoff: a bulk process sourcing large amounts of pseudorandomness must result in high reconstruction complexity. While pseudorandomness is itself not particularly complex to implement, reconstruction of pseudorandom quantities is by definition of high complexity class. In this case, Cryptographic Censorship may act to protect a computationally bounded asymptotic observer from detecting large violations of effective field theory due to computational complexity by forcing formation of an event horizon, depending on where the event horizon forms.

Stringy and cosmological applications.

It would be interesting to consider stringy effects and ask if (and how) our theorem might apply to, e.g., the D1-D5 system, or orbifold singularities, which are horizonless (see, e.g., [83]). Exact BPS states have trivial dynamics: we expect that to produce chaos and pseudorandom dynamics we would need to sufficiently excite these states. If it is indeed the case that chaotic, pseudorandom time evolution is produced as a result of heavily exciting such states, then our theorem suggests that a horizon will appear. More generally, it is not clear that horizonless geometries have any relation to chaos, and furthermore such states are probably not typical in the measure concentration sense that we are using in this work.

Let us also briefly comment on non-AdS cosmological spacetimes. In particular, it is interesting to consider why our results fail for flat cosmologies, which do not have horizons. An obvious complaint is that our AdS/CFT arguments should not necessarily apply. They do not apply, however, in an interesting manner: the relevant feature of AdS/CFT that fails here is the lack of a notion of unitary time evolution on the conformal boundary (which in this case is null). Our results crucially rely on a fundamental description of the system which is a *unitary* quantum field theory. This in all likelihood does not apply to flat cosmologies. It does, however, apply to AdS big-crunch cosmologies, and in those, we do indeed have a horizon. See figure 16 for a comparison.

What about de Sitter? de Sitter cosmologies do have cosmological horizons, although those are not unambiguous in the same sense as AdS ones. There have been some suggestions (most recently [84, 85]) that de Sitter time evolution — insofar as such a notion can be defined — is well-modeled by an isometry rather than a unitary. A generalization from pseudorandom unitaries to pseudorandom isometries was recently proposed by [86]. It would be interesting to see if our results can be generalized to accommodate this notion and potentially have some bearing on de Sitter black holes and cosmologies; in particular, if there is any sense in which the relaxation to isometries could be mapped to the observer-dependence of de Sitter horizons.




Figure 16. The AdS-big-crunch cosmology (a) versus the flat big-bang cosmology (b). In the asymptotically flat case, there is no notion of unitary time evolution and no horizon. In the AdS case, we have both.

Some open questions.

Cryptographic Censorship paves the way for several interesting new investigations. We have expanded on a few of those above. We briefly enumerate a few others here:

1. The fact that gravitational pseudorandomness associated to a quantum computer in the bulk spacetime must result in a horizon suggests a relationship between backreaction and computational power. Does such a relation between energy and computation exist, perhaps analogous to a gravitational version of Lloyd’s bound [87]?
2. Are singularities in general in fact associated with large amounts of pseudorandomness?⁴⁸ This idea features prominently in our models of timelike naked singularities, but has traditionally appeared in the context of postselection models of black hole singularities and dynamics.
3. Our definition of the size of a singularity admits a middle class of singularities that disappear in the strict large- N limit but which could accommodate gravitational pseudorandomness in code subspaces constructed from $O(1)$ time bands. It would be interesting to understand the significance, if any, and frequency of semi-Planckian naked singularity portions.
4. Explicit constructions of pseudorandom unitary ensembles to date have not been found. Still, as assumed in this work, we expect explicit examples of AdS/CFT to be well modeled by pseudorandom dynamics. Can we use complexity theoretic conjectures and aspects of AdS/CFT to demonstrate that holographic quantum systems like the Sachdev-Ye-Kitaev (SYK) model actually generate a pseudorandom unitary ensemble at long enough times?
5. Is there anything we can infer about code subspace-dependent cryptographic conditions/constructions necessary for pseudorandomness from the holographic manifestation of our results? That is, does the existence of subspaces without horizons (e.g., the microcanonical window around the vacuum) have any meaningful implication for cryptography?

⁴⁸See [88] for an unrelated attempt to connect singularities to holographic CA/CV-complexity conjectures.

Acknowledgments

It is a pleasure to thank Chris Akers, Jordan Cotler, Jordan Docter, Roberto Emparan, Daniel Harlow, Gary Horowitz, Hsin-Yuan (Robert) Huang, Andrea Puhm, Juan Maldacena, Arvin Shahbazi-Moghaddam, Natalie Paquette, Jon Sorce, Edward Witten, and Michelle Xu for valuable discussions. This work is supported in part by the Department of Energy under Early Career Award DE-SC0021886 (NE, ÅF, and LY) and QuantISED DE-SC0020360 contract no. 578218 (NE), by the Heising-Simons Foundation under grant no. 2023-4430 (NE and AL), the Sloan Foundation (NE), the Templeton Foundation via the Black Hole Initiative (NE and EV), the Gordon and Betty Moore Foundation (EV), the Packard Foundation (AL), and the MIT department of physics (NE). The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton or Moore Foundations.

A Proof of Theorem 3

In this appendix, we prove Theorem 3, introduced in section 3.2. We begin with a definition of randomized K -Lipschitz, which we will require our algorithms to satisfy.

Definition 6 (Randomized K -Lipschitz). *An algorithm \mathcal{A} is said to be randomized K -Lipschitz if it satisfies the following condition: there exists a constant $K > 0$ such that*

$$\text{avg}_{\substack{\widehat{U}_1 \leftarrow \mathcal{A}^{U_1} \\ \widehat{U}_2 \leftarrow \mathcal{A}^{U_2}}} \|\widehat{U}_1 - \widehat{U}_2\|_\infty \leq K \|U_1 - U_2\|_\infty . \tag{A.1}$$

When the algorithm is deterministic, i.e., it maps a given unitary U to a single output \widehat{U} , we will just refer to this property as K -Lipschitz.

In addition, here we will work with operators that are exactly norm-preserving, i.e., $\|\widehat{U}\|_\infty = 1$. We leave it to the mathematician to confirm that this result applies equally well to approximately norm-preserving operators.

We now repeat Theorem 3 for the reader's convenience.

Theorem 3 (Complexity of Learning for Typical Pseudorandom Unitaries). *Let $\kappa = \text{poly}(N)$ and consider a code subspace of dimension 2^κ . Consider any (gravitationally) pseudorandom unitary ensemble \mathcal{U} , any pseudorandom state ensemble Ψ , and any constant $\epsilon > 0$. For any (randomized) K -Lipschitz quantum algorithm \mathcal{A} of $\text{poly}(N)$ complexity that has oracle access to $U \leftarrow \mathcal{U}$ and produces an approximately norm-preserving operator \widehat{U} as its best guess for U , and for any sufficiently large N , the following holds: for a fraction $\geq 1 - e^{-\epsilon^2 2^\kappa / (48(1+K)^2)} - e^{-\epsilon^2 2^\kappa / 192} - \text{negl}'(N)$ of $U \in \mathcal{U}$ and $|\psi\rangle \in \Psi$, we have that:*

1. the following quantity is bounded

$$\text{avg}_{\widehat{U} \leftarrow \mathcal{A}^U} \left[F(U|\psi), \widehat{U}|\psi \right] \leq 1 - \alpha + 2\epsilon + \text{negl}(N), \tag{3.5}$$

where α is a constant independent of N , and

2. there exists an operator Q (which may depend on U , $|\psi\rangle$, and \widehat{U}) such that

$$\text{avg}_{\widehat{U} \leftarrow \mathcal{A}^U} \left| \text{tr}(Q\widehat{U}\psi(\widehat{U})^\dagger) - \text{tr}(QU\psi(U)^\dagger) \right| \geq 2\alpha - 4\epsilon - \text{negl}(N) . \quad (3.6)$$

Proof. We prove 1. and 2. consecutively for Hilbert spaces of any finite (but sufficiently large) dimension d , where we always have in mind that d scales with κ . Note that in the theorem above, and in the main body of the paper, we specialized to a code subspace of dimension $d = 2^\kappa$ with $\kappa = \text{poly}(N)$; here we will leave d general.

1. Fidelity for typical pseudorandom unitaries. We start by proving that the bound on the averaged fidelity in Theorem 2 also holds for *typical* Haar random unitaries and states, and subsequently prove this for pseudorandom unitaries and states. To do so, we employ techniques from measure concentration. An introduction to the theory of measure concentration, including numerous references, is given in appendix B of [30]. Here we only need the following lemma for measure concentration on the unitary group.

Lemma 1 (Meckes (Corollary 17 in [89])). *Let $f : U(d) \rightarrow \mathbb{R}$ be K -Lipschitz in the sense that $|f(U_1) - f(U_2)| \leq K\|U_1 - U_2\|_2$, with $\|X\|_2 \equiv \sqrt{\text{tr } X^\dagger X}$. Then in the Haar measure on $U(d)$ we have the following concentration inequalities for any $\epsilon > 0$:*

$$\begin{aligned} \Pr[f \geq \langle f \rangle + \epsilon] &\leq \exp\left(-\frac{\epsilon^2 d}{12K^2}\right), \\ \Pr[f \leq \langle f \rangle - \epsilon] &\leq \exp\left(-\frac{\epsilon^2 d}{12K^2}\right), \\ \Pr[|f - \langle f \rangle| \geq \epsilon] &\leq 2 \exp\left(-\frac{\epsilon^2 d}{12K^2}\right). \end{aligned} \quad (\text{A.2})$$

For the purposes of Theorem 3, we are interested in applying this to the function $F = F(U|\psi\rangle, \widehat{U}|\psi\rangle) = |\langle \psi|U^\dagger\widehat{U}|\psi\rangle|^2$, where we first consider Haar random U and $|\psi\rangle$, i.e., $U \leftarrow \mu$ and $|\psi\rangle \leftarrow \mu$, and furthermore \widehat{U} is the guess that the algorithm \mathcal{A}^U outputs. Thus, our task is to show that the function F obeys the requirements for Lemma 1 to hold, such that fluctuations around the average are exponentially suppressed. That is, we need to check that F is K -Lipschitz for some fixed K .

Note that the algorithm's output \widehat{U} depends on U . This dependence is complicated: the quantum algorithm \mathcal{A}^U , as a result of, e.g., intrinsic randomness in the outcome of measurements during the learning process, generally outputs a \widehat{U} that is not a deterministic function of U . In the following, we will start by considering instead a deterministic algorithm, i.e., the algorithm is some function $\mathcal{A}(U)$ that maps every unitary U to some norm-preserving operator \widehat{U} .

We then prove the following lemma:

Lemma 2. *Let $\mathcal{A}(U)$ denote any deterministic K -Lipschitz function that maps every unitary U to some norm-preserving operator \widehat{U} . Let $|\psi\rangle$ and $|\varphi\rangle$ be any two fixed states in the Hilbert space \mathcal{H} . Then the function $F(U) \equiv F(U|\psi\rangle, \mathcal{A}(U)|\varphi\rangle)$ is $(2 + 2K)$ -Lipschitz in U .*

Proof. Our aim is to show that

$$|F(U_1) - F(U_2)| \leq (2 + 2K) \|U_1 - U_2\|_2 . \quad (\text{A.3})$$

To show (A.3), first, notice that we can rewrite

$$\begin{aligned} F(U) &= F(U|\psi\rangle, \mathcal{A}(U)|\varphi\rangle) = \text{tr} \left[\mathcal{A}(U)|\varphi\rangle\langle\varphi| \mathcal{A}(U)^\dagger U|\psi\rangle\langle\psi| U^\dagger \right] \\ &= \text{tr} \left[(\mathcal{A}(U)^\dagger U)^\dagger \varphi \mathcal{A}(U)^\dagger U \psi \right] \\ &= \text{tr} \left[\left((\mathcal{A}(U)^\dagger U)^\dagger \otimes \mathcal{A}(U)^\dagger U \right) (\varphi \otimes \psi) \mathbf{S} \right] . \end{aligned} \quad (\text{A.4})$$

Here we used $\varphi = |\varphi\rangle\langle\varphi|$ and $\psi = |\psi\rangle\langle\psi|$. In the third line we introduced a second copy of the Hilbert space; \mathbf{S} is the swap operator. For readability, in what follows, we will abbreviate $\mathcal{A}(U_i) = \mathcal{A}_i$. Then, using Hölder's inequality,

$$\begin{aligned} |F(U_1) - F(U_2)| &= \left| \text{tr} \left[\left((\mathcal{A}_1^\dagger U_1)^\dagger \otimes \mathcal{A}_1^\dagger U_1 - (\mathcal{A}_2^\dagger U_2)^\dagger \otimes \mathcal{A}_2^\dagger U_2 \right) (\varphi \otimes \psi) \mathbf{S} \right] \right| \\ &\leq \|\varphi \otimes \psi\|_1 \|\mathbf{S} \left((\mathcal{A}_1^\dagger U_1)^\dagger \otimes \mathcal{A}_1^\dagger U_1 - (\mathcal{A}_2^\dagger U_2)^\dagger \otimes \mathcal{A}_2^\dagger U_2 \right)\|_\infty . \end{aligned} \quad (\text{A.5})$$

Since $\varphi \otimes \psi$ is a state, its trace norm is 1. Now we can combine submultiplicativity of the operator norm with the fact that \mathbf{S} is unitary, such that

$$\begin{aligned} &\|\mathbf{S} \left((\mathcal{A}_1^\dagger U_1)^\dagger \otimes \mathcal{A}_1^\dagger U_1 - (\mathcal{A}_2^\dagger U_2)^\dagger \otimes \mathcal{A}_2^\dagger U_2 \right)\|_\infty \\ &\leq \|\mathbf{S}\|_\infty \| \left((\mathcal{A}_1^\dagger U_1)^\dagger \otimes \mathcal{A}_1^\dagger U_1 - (\mathcal{A}_2^\dagger U_2)^\dagger \otimes \mathcal{A}_2^\dagger U_2 \right) \|_\infty \\ &= \| \left((\mathcal{A}_1^\dagger U_1)^\dagger \otimes \mathcal{A}_1^\dagger U_1 - (\mathcal{A}_2^\dagger U_2)^\dagger \otimes \mathcal{A}_2^\dagger U_2 \right) \|_\infty . \end{aligned} \quad (\text{A.6})$$

Recall that we want to show that this is upper bounded by some $K' \|U_1 - U_2\|_2$. Rewrite the last line as

$$\begin{aligned} &\| \left((\mathcal{A}_1^\dagger U_1)^\dagger \otimes \mathcal{A}_1^\dagger U_1 - (\mathcal{A}_1^\dagger U_1)^\dagger \otimes \mathcal{A}_2^\dagger U_2 + (\mathcal{A}_1^\dagger U_1)^\dagger \otimes \mathcal{A}_2^\dagger U_2 - (\mathcal{A}_2^\dagger U_2)^\dagger \otimes \mathcal{A}_2^\dagger U_2 \right) \|_\infty \\ &\leq \| \left((\mathcal{A}_1^\dagger U_1)^\dagger \otimes \mathcal{A}_1^\dagger U_1 - (\mathcal{A}_1^\dagger U_1)^\dagger \otimes \mathcal{A}_2^\dagger U_2 \right) \|_\infty + \| \left((\mathcal{A}_1^\dagger U_1)^\dagger \otimes \mathcal{A}_2^\dagger U_2 - (\mathcal{A}_2^\dagger U_2)^\dagger \otimes \mathcal{A}_2^\dagger U_2 \right) \|_\infty \\ &= 2 \| \mathcal{A}_1^\dagger U_1 - \mathcal{A}_2^\dagger U_2 \|_\infty . \end{aligned} \quad (\text{A.7})$$

In the second line we used the triangle inequality; in the third we used $\|A \otimes B\| = \|A\| \|B\|$, and for unitaries $\|U\|_\infty = 1$, and we are working with exactly norm-preserving operators \widehat{U} such that also $\|\widehat{U}\|_\infty = 1$. We now again use the triangle inequality and submultiplicativity of the operator norm to rewrite the final line as

$$\begin{aligned} 2 \| \mathcal{A}_1^\dagger U_1 - \mathcal{A}_2^\dagger U_2 \|_\infty &= 2 \| \mathcal{A}_1^\dagger U_1 - \mathcal{A}_2^\dagger U_1 + \mathcal{A}_2^\dagger U_1 - \mathcal{A}_2^\dagger U_2 \|_\infty \\ &\leq 2 \| \mathcal{A}_1^\dagger - \mathcal{A}_2^\dagger \|_\infty \|U_1\|_\infty + 2 \| \mathcal{A}_2^\dagger \|_\infty \|U_1 - U_2\|_\infty \\ &\leq (2 + 2K) \|U_1 - U_2\|_\infty \\ &\leq (2 + 2K) \|U_1 - U_2\|_2 . \end{aligned} \quad (\text{A.8})$$

In the third line we used that $\mathcal{A}(U)$ is K -Lipschitz for some K , and finally we used $\|\cdot\|_\infty \leq \|\cdot\|_2$. This completes the proof that the function $F(U) \equiv F(U|\psi\rangle, \mathcal{A}(U)|\varphi\rangle)$ is $(2 + 2K)$ -Lipschitz:

$$|F(U_1) - F(U_2)| \leq (2 + 2K) \|U_1 - U_2\|_2 . \quad (\text{A.9})$$

□

Then, by Lemma 1 and Lemma 2, we have the following measure concentration result: that for fraction $\geq 1 - e^{-\epsilon^2 d / (48(1+K)^2)}$ of Haar random unitaries,

$$F(U|\psi\rangle, \mathcal{A}(U)|\varphi\rangle) - \text{avg}_{V \leftarrow \mu} [F(V|\psi\rangle, \mathcal{A}(V)|\varphi\rangle)] < \epsilon. \quad (\text{A.10})$$

We now wish to extend this to include measure concentration over Haar random states. To this end, we prove the following lemma.

Lemma 3. *Let $\mathcal{A}(U)$ denote any deterministic K -Lipschitz function that maps every unitary U to some norm-preserving operator \hat{U} . For fraction $\geq 1 - e^{-\epsilon^2 d / (48(1+K)^2)}$ of Haar random unitaries $U \leftarrow \mu$ and fraction $\geq 1 - e^{-\epsilon^2 d / 192}$ of Haar random states $|\psi\rangle \leftarrow \mu$, we have*

$$F(U|\psi\rangle, \mathcal{A}(U)|\psi\rangle) - \text{avg}_{\substack{V \leftarrow \mu \\ |\varphi\rangle \leftarrow \mu}} [F(V|\varphi\rangle, \mathcal{A}(V)|\varphi\rangle)] < 2\epsilon. \quad (\text{A.11})$$

Proof. Consider first fixing some U and $\mathcal{A}(U)$. Let V be a Haar random unitary and $|\psi_0\rangle$ some basis state, such that $V|\psi_0\rangle$ is a Haar random state. We show that the function

$$F(V) \equiv F(UV|\psi_0\rangle, \mathcal{A}(U)V|\psi_0\rangle) = F(V^\dagger \mathcal{A}(U)^\dagger UV|\psi_0\rangle, |\psi_0\rangle) \quad (\text{A.12})$$

is 4-Lipschitz. Following the same methods as in the proof of Lemma 2, we find the analogue of (A.7):

$$|F(V_1) - F(V_2)| \leq 2\|V_1^\dagger \mathcal{A}(U)^\dagger UV_1 - V_2^\dagger \mathcal{A}(U)^\dagger UV_2\|_\infty. \quad (\text{A.13})$$

Then, by the triangle inequality,

$$\begin{aligned} |F(V_1) - F(V_2)| &\leq 2\|V_1^\dagger \mathcal{A}(U)^\dagger UV_1 - V_2^\dagger \mathcal{A}(U)^\dagger UV_1 + V_2^\dagger \mathcal{A}(U)^\dagger UV_1 - V_2^\dagger \mathcal{A}(U)^\dagger UV_2\|_\infty \\ &\leq 2\|V_1^\dagger - V_2^\dagger\|_\infty \|\mathcal{A}(U)^\dagger UV_1\|_\infty + 2\|V_2^\dagger \mathcal{A}(U)^\dagger U\|_\infty \|V_1 - V_2\|_\infty \\ &\leq 4\|V_1 - V_2\|_2. \end{aligned} \quad (\text{A.14})$$

Note that in the last line we used again that $\mathcal{A}(U)$ is exactly norm-preserving, such that $\|\mathcal{A}(U)^\dagger UV_i\|_\infty = 1$, as well as $\|\cdot\|_\infty \leq \|\cdot\|_2$.

By Lemma 1, for fraction $\geq 1 - e^{-\epsilon^2 d / 192}$ of $|\psi\rangle \equiv V|\psi_0\rangle \leftarrow \mu$,

$$F(U|\psi\rangle, \mathcal{A}(U)|\psi\rangle) - \text{avg}_{|\varphi\rangle \leftarrow \mu} [F(U|\varphi\rangle, \mathcal{A}(U)|\varphi\rangle)] < \epsilon. \quad (\text{A.15})$$

We now combine this with measure concentration for Haar random unitaries. Starting from (A.10) for two (fixed) states $|\varphi\rangle$, and then averaging over $\varphi \leftarrow \mu$, we have for fraction $\geq 1 - e^{-\epsilon^2 d / 48(1+K)^2}$ of $U \leftarrow \mu$,

$$\text{avg}_{|\varphi\rangle \leftarrow \mu} [F(U|\varphi\rangle, \mathcal{A}(U)|\varphi\rangle)] - \text{avg}_{\substack{V \leftarrow \mu \\ |\varphi\rangle \leftarrow \mu}} [F(V|\varphi\rangle, \mathcal{A}(V)|\varphi\rangle)] < \epsilon. \quad (\text{A.16})$$

Now add and subtract $F(U|\psi\rangle, \mathcal{A}(U)|\psi\rangle)$. This gives

$$\begin{aligned} &F(U|\psi\rangle, \mathcal{A}(U)|\psi\rangle) - \text{avg}_{\substack{V \leftarrow \mu \\ |\varphi\rangle \leftarrow \mu}} [F(V|\varphi\rangle, \mathcal{A}(V)|\varphi\rangle)] \\ &< \epsilon + F(U|\psi\rangle, \mathcal{A}(U)|\psi\rangle) - \text{avg}_{|\varphi\rangle \leftarrow \mu} [F(U|\varphi\rangle, \mathcal{A}(U)|\varphi\rangle)]. \end{aligned} \quad (\text{A.17})$$

Now we can use (A.15) to bound the right hand side to 2ϵ for fraction $\geq 1 - e^{-\epsilon^2 d/192}$ of $|\psi\rangle \leftarrow \mu$. This concludes the proof of Lemma 3. \square

To prove Theorem 3, it remains to extend the above results to (1) non-deterministic functions $\mathcal{A}(U)$, e.g., quantum algorithms \mathcal{A}^U , and (2) to pseudorandom unitaries and states. We first briefly address point (1). We can run the proofs of Lemma 2 and Lemma 3 analogously for $\widehat{U}_i \leftarrow \mathcal{A}^{U_i}$ until we need a smoothness property of \widehat{U}_i . There, instead of K -Lipschitz, we use the randomized K -Lipschitz property in Definition 6. Furthermore, using the inequality

$$|\text{avg } F(U_1) - \text{avg } F(U_2)| \leq \text{avg } |F(U_1) - F(U_2)|, \tag{A.18}$$

we have that the function $F(V) = \text{avg}_{\widehat{V} \leftarrow \mathcal{A}^V} [F(V|\psi), \widehat{V}|\psi]$ is itself $(2 + 2K)$ -Lipschitz. This allows us to extend our measure concentration results to this function $F(V)$ where previously we proved them for $F(U) = F(U|\psi), \widehat{U}|\psi$.

Then we arrive at the following result: that for fraction $\geq 1 - e^{-\epsilon^2 d/(48(1+K)^2)}$ of $U \leftarrow \mu$, and for fraction $\geq 1 - e^{-\epsilon^2 d/192}$ of $|\psi\rangle \leftarrow \mu$ (alternatively, we can think of this as $\geq 1 - e^{-\epsilon^2 d/192} - e^{-\epsilon^2 d/(48(1+K)^2)}$ fraction of $U \leftarrow \mu$ and $|\psi\rangle \leftarrow \mu$), we have the following concentration bound:

$$\text{avg}_{\widehat{U} \leftarrow \mathcal{A}^U} [F(U|\psi), \widehat{U}|\psi] - \text{avg}_{\substack{V \leftarrow \mu \\ |\varphi\rangle \leftarrow \mu}} \left[\text{avg}_{\widehat{V} \leftarrow \mathcal{A}^V} [F(V|\varphi), \widehat{V}|\varphi] \right] < 2\epsilon. \tag{A.19}$$

Finally, we address point (2) above, and show that to preserve indistinguishability, the above result must also hold for *pseudorandom* unitaries and states. Intuitively, if the fluctuations around the Haar random averaged fidelity are exponentially suppressed, but the fluctuations around the pseudorandom averaged fidelity are not, then one could measure these fluctuations and thereby distinguish pseudorandom unitaries and states from Haar random ones, breaking the pseudorandomness.

We proceed to formally show this below.

Lemma 4. *Let $\kappa = \text{poly}(N)$ be the security parameter. Let \mathcal{U} be any pseudorandom unitary (PRU) ensemble and let Ψ be any pseudorandom state (PRS) ensemble. Let \widehat{U} be the output of the learning algorithm \mathcal{A}^U , which we assume to be (randomized) K -Lipschitz, and where $U \leftarrow \mathcal{U}$. Then the following concentration inequality holds, in the measure induced on \mathcal{U} and Ψ , for any $\epsilon > 0$, with probability $\geq 1 - e^{-\epsilon^2 d/48(1+K)^2} - e^{-\epsilon^2 d/192} - \text{negl}'(N)$ over $U \leftarrow \mathcal{U}$ and $|\psi\rangle \leftarrow \Psi$,*

$$\text{avg}_{\widehat{U} \leftarrow \mathcal{A}^U} [F(U|\psi), \widehat{U}|\psi] - \text{avg}_{\substack{V \leftarrow \mathcal{U} \\ |\varphi\rangle \leftarrow \Psi \\ \widehat{V} \leftarrow \mathcal{A}^V}} [F(V|\varphi), \widehat{V}|\varphi] < 2\epsilon + \text{negl}(N). \tag{A.20}$$

Proof. To prove Lemma 4, we first extend (A.19) to pseudorandom unitaries, and then to pseudorandom states.

First, let Ψ be any ensemble of states. For ease of notation, we write

$$\beta := \text{avg}_{\substack{U \leftarrow \mu \\ |\psi\rangle \leftarrow \Psi \\ \widehat{U} \leftarrow \mathcal{A}^U}} \left[F(U|\psi), \widehat{U}|\psi\rangle \right] \tag{A.21}$$

and

$$\beta' := \text{avg}_{\substack{U \leftarrow \mathcal{U} \\ |\psi\rangle \leftarrow \Psi \\ \widehat{U} \leftarrow \mathcal{A}^U}} \left[F(U|\psi), \widehat{U}|\psi\rangle \right] . \tag{A.22}$$

By Lemma 7 of [32], $|\beta - \beta'| \leq \eta(N)$ where η is a negligible function.

Assume towards contradiction that there exists a pseudorandom unitary ensemble \mathcal{U} , any distribution of states Ψ , and a positive polynomial p such that

$$\Pr_{\substack{U \leftarrow \mathcal{U} \\ |\psi\rangle \leftarrow \Psi}} \left[\text{avg}_{\widehat{U} \leftarrow \mathcal{A}^U} F(U|\psi), \widehat{U}|\psi\rangle - \beta' \geq \epsilon + \eta(N) \right] \geq e^{-\epsilon^2 d / 48(1+K)^2} + \frac{1}{p(N)} . \tag{A.23}$$

If the above equation holds, then there is a state $|\psi_0\rangle$ such that this equation holds with $|\psi\rangle := |\psi_0\rangle$:

$$\Pr_{U \leftarrow \mathcal{U}} \left[\text{avg}_{\widehat{U} \leftarrow \mathcal{A}^U} F(U|\psi_0), \widehat{U}|\psi_0\rangle - \beta' \geq \epsilon + \eta(N) \right] \geq e^{-\epsilon^2 d / 48(1+K)^2} + \frac{1}{p(N)} . \tag{A.24}$$

We now construct an efficient algorithm Dist that uses oracle access to the pseudorandom unitary U and also can run the learning algorithm \mathcal{A}^U .⁴⁹ We show that if our assumption (A.23) is true, then Dist can distinguish between $U \leftarrow \mu$ and $U \leftarrow \mathcal{U}$, thus breaking the pseudorandomness of \mathcal{U} .

The distinguisher Dist gets $\text{poly}(N)$ copies of $|\psi_0\rangle$, and is also handed β .⁵⁰ Then it proceeds as follows:

1. Dist runs \mathcal{A} and simulates the oracle for \mathcal{A} : it receives each query from \mathcal{A} , sends the query to its oracle U , and returns the response to \mathcal{A} .
2. Dist gives $|\psi_0\rangle$ to \mathcal{A} and receives back $\widehat{U}|\psi_0\rangle$. It also gives $|\psi_0\rangle$ to its oracle U and gets back $U|\psi_0\rangle$.
3. Dist repeats this $\text{poly}(N)$ times and then uses a quantum algorithm (e.g., [90]) to efficiently estimate $\beta'' := F(U|\psi_0), \widehat{U}|\psi_0\rangle$ up to some additive error $\hat{\epsilon} > 0$, i.e., $|\beta''_{\text{est}} - \beta''| \leq \hat{\epsilon}$.
4. If $\beta''_{\text{est}} - \beta \geq \tilde{\epsilon}$, where $\tilde{\epsilon} > \hat{\epsilon}$, then Dist outputs 1.

⁴⁹Note that the distinguisher can depend on N , as ours does here, since e.g. $|\psi_0\rangle$ depends on N .

⁵⁰The careful reader might worry that handing the distinguisher (multiple copies of) *any* state in \mathcal{H} , or β , breaks the computational boundedness of the algorithm. Note, however, that although the algorithms we consider are limited to a polynomial number of queries in $\log \dim \mathcal{H}$, each query can be for any state in \mathcal{H} that the algorithm a) can prepare or b) has been given as “advice”. This advice can also include quantities like β . Note that the learning algorithm we consider in, e.g. Theorem 3 gets only *one* copy of the physical state $|\psi\rangle$; the distinguisher that we describe here is a different algorithm and can receive multiple copies as advice.

Note that Dist runs in $\text{poly}(N)$ -time since it can efficiently query the oracle, and furthermore the quantum algorithm for fidelity estimation is efficient.⁵¹

We now show that Dist distinguishes between $U \leftarrow \mu$ and $U \leftarrow \mathcal{U}$. For Haar random $U \leftarrow \mu$ as the oracle,

$$\begin{aligned} \Pr_{U \leftarrow \mu} [\text{Dist} = 1] &= \Pr [\beta''_{\text{est}} - \beta \geq \tilde{\epsilon}] = \Pr [\beta'' - \beta \geq \tilde{\epsilon} - (\beta''_{\text{est}} - \beta'')] \\ &\leq \Pr [\beta'' - \beta \geq \tilde{\epsilon} - \hat{\epsilon}] \\ &\leq e^{-(\tilde{\epsilon} - \hat{\epsilon})^2 d / 48(1+K)^2}, \end{aligned} \quad (\text{A.25})$$

where we used $-\hat{\epsilon} \leq \beta''_{\text{est}} - \beta'' \leq \hat{\epsilon}$ in the second line, and (A.10) (concentration in the Haar measure) for $\epsilon = \tilde{\epsilon} - \hat{\epsilon} > 0$ in the last line.

For pseudorandom $U \leftarrow \mathcal{U}$ as the oracle,

$$\begin{aligned} \Pr_{U \leftarrow \mathcal{U}} [\text{Dist} = 1] &= \Pr [\beta''_{\text{est}} - \beta \geq \tilde{\epsilon}] = \Pr [\beta'' - \beta' \geq \tilde{\epsilon} - (\beta''_{\text{est}} - \beta'') - (\beta' - \beta)] \\ &\geq \Pr [\beta'' - \beta' \geq \tilde{\epsilon} + \hat{\epsilon} + \eta(N)] \\ &\geq e^{-(\tilde{\epsilon} + \hat{\epsilon})^2 d / 48(1+K)^2} + \frac{1}{p(N)}, \end{aligned} \quad (\text{A.26})$$

where we used $-\eta(N) \leq \beta - \beta' \leq \eta(N)$ in the third line, and our assumption (A.23) in the last line for $\epsilon = \tilde{\epsilon} + \hat{\epsilon}$.

Therefore,

$$\Pr_{U \leftarrow \mathcal{U}} [\text{Dist} = 1] - \Pr_{U \leftarrow \mu} [\text{Dist} = 1] \geq \frac{1}{p(N)} + e^{-\frac{(\tilde{\epsilon} + \hat{\epsilon})^2 d}{48(1+K)^2}} - e^{-\frac{(\tilde{\epsilon} - \hat{\epsilon})^2 d}{48(1+K)^2}}. \quad (\text{A.27})$$

This shows that Dist distinguishes between $U \leftarrow \mu$ and $U \leftarrow \mathcal{U}$. Since Dist runs in $\text{poly}(N)$ time, this contradicts the pseudorandomness of \mathcal{U} (Definition 3). Thus based on the pseudorandomness of \mathcal{U} , we must have that for fraction $\leq e^{-\epsilon^2 d / 48(1+K)^2} + \eta'(N)$ of $U \leftarrow \mathcal{U}$

$$\text{avg}_{\hat{U} \leftarrow \mathcal{A}^U} F(U|\psi\rangle, \hat{U}|\psi\rangle) - \text{avg}_{\substack{V \leftarrow \mathcal{U} \\ |\varphi\rangle \leftarrow \Psi \\ \hat{V} \leftarrow \mathcal{A}^V}} F(V|\varphi\rangle, \hat{V}|\varphi\rangle) \geq \epsilon + \eta(N), \quad (\text{A.28})$$

where $\eta'(N)$ is negligible. In particular, this holds when $|\psi\rangle \leftarrow \Psi$ is Haar random.

Thus, we have proven a measure concentration result for fluctuations around the average fidelity where we average over pseudorandom unitaries (and Haar random states). An analogous proof holds for fluctuations around the average fidelity where we average over pseudorandom states (and any distribution of unitaries \mathcal{D} ; in particular, it holds for Haar random unitaries $U \leftarrow \mu$). It requires Lemma 8 in [32]. This would result in the following measure concentration result: that for fraction $\leq e^{-\epsilon^2 d / 192} + \tilde{\eta}'(N)$ of $|\psi\rangle \leftarrow \Psi$,

$$\text{avg}_{\hat{U} \leftarrow \mathcal{A}^U} F(U|\psi\rangle, \hat{U}|\psi\rangle) - \text{avg}_{\substack{V \leftarrow \mathcal{U} \\ |\varphi\rangle \leftarrow \Psi \\ \hat{V} \leftarrow \mathcal{A}^V}} [F(V|\varphi\rangle, \hat{V}|\varphi\rangle)] \geq \epsilon + \tilde{\eta}(N), \quad (\text{A.29})$$

where we explicated that the negligible functions need not be the same as in (A.28). Combining (A.28) and (A.29) concludes the proof of Lemma 4. \square

⁵¹For example, the quantum algorithm of [90] runs in $\text{poly}(1/\hat{\epsilon})$ -time.

We will now use Lemma 4 to prove the first part of Theorem 3. Consider any constant $\epsilon > 0$. Consider any sufficiently large value of $N \in \mathbb{N}$ so that Theorem 2 from [32] holds. Let $U \leftarrow \mathcal{U}$ and $|\psi\rangle \leftarrow \Psi$ denote the pseudorandom unitary and state. Let \widehat{U} denote the operator that \mathcal{A}^U produces. To connect to the learning result (Theorem 2), we average over the output of the algorithm, $\widehat{U}|\psi\rangle \leftarrow \mathcal{A}^U(|\psi\rangle)$. We abbreviate this by $\widehat{U} \leftarrow \mathcal{A}^U$. By Lemma 4, we have that with probability $\geq 1 - e^{-\epsilon^2 d/48(1+K)^2} - e^{-\epsilon^2 d/192} - \text{negl}'(N)$,

$$\text{avg}_{\widehat{U} \leftarrow \mathcal{A}^U} \left[F(U|\psi), \widehat{U}|\psi\rangle \right] - \text{avg}_{\substack{V \leftarrow \mathcal{U} \\ |\varphi\rangle \leftarrow \Psi \\ \widehat{V} \leftarrow \mathcal{A}^V}} \left[F(V|\varphi), \widehat{V}|\varphi\rangle \right] < 2\epsilon + \text{negl}(N). \quad (\text{A.30})$$

We call $(U, |\psi\rangle)$ *typical* if the above equation holds. The complexity of learning result (Theorem 2) gives us an upper bound on the average fidelity. Thus we have that for any typical $(U, |\psi\rangle)$,

$$\text{avg}_{\widehat{U} \leftarrow \mathcal{A}^U} \left[F(U|\psi), \widehat{U}|\psi\rangle \right] \leq 1 - \alpha + 2\epsilon + \text{negl}(N). \quad (\text{A.31})$$

This shows the first part of the theorem (Equation (3.5)).

Haar random unitaries and states. Note that we obtain a stronger result for Haar random unitaries and states: namely, that for fraction $\geq 1 - e^{-\epsilon^2 d/48(1+K)^2} - e^{-\epsilon^2 d/192}$ of $U \leftarrow \mu$ and $|\psi\rangle \leftarrow \mu$,

$$\text{avg}_{\widehat{U} \leftarrow \mathcal{A}^U} \left[F(U|\psi), \widehat{U}|\psi\rangle \right] \leq 1 - \alpha + 2\epsilon. \quad (\text{A.32})$$

Note that for $d = 2^\kappa$ with $\kappa = \text{poly}(N)$, this means that the above bound is violated for only a fraction that is doubly-exponentially suppressed in N .

2. Distinguishing operator. We now show the second part of the theorem. The relationship between fidelity and distinguishing measurements is a known result; for any finite dimensional Hilbert space, this is given in appendix B in [32]. Thus, for any typical $(U, |\psi\rangle)$, there exists an operator Q (for a given \widehat{U}) for which

$$\text{avg}_{\widehat{U} \leftarrow \mathcal{A}^U} \left| \text{tr}(Q\widehat{U}\psi\widehat{U}^\dagger) - \text{tr}(QU\psi U^\dagger) \right| \geq 2\alpha - 4\epsilon - \text{negl}(N). \quad (\text{A.33})$$

This concludes the proof of Theorem 3. □

B Proof of Corollary 3.1

In this appendix, we prove Corollary 3.1, the existence of a semiclassical distinguishing operator. We repeat Corollary 3.1 here for convenience:

Corollary 3.1 (Distinguishing Operator at $N \rightarrow \infty$). *Let $\{U_N|\psi_N\rangle\}_N$ and $\{\widehat{U}_N|\psi_N\rangle\}_N$ be two sequences of states that are represented on the code, namely there exist $|\xi\rangle$ and $|\hat{\xi}\rangle$ such that $U_N|\psi_N\rangle = V_N|\xi\rangle$ and $\widehat{U}_N|\psi_N\rangle = V_N|\hat{\xi}\rangle$. Furthermore, assume that $\{U_N|\psi_N\rangle\}_N$ and $\{\widehat{U}_N|\psi_N\rangle\}_N$ satisfy the assumptions stated above, then both 1. and 2. in Theorem 3 also*

hold for the states $|\xi\rangle$ and $|\hat{\xi}\rangle$. In particular, by taking the constant ϵ to be $< \alpha/2$, we have for any $|\hat{\xi}\rangle$ the existence of an operator $Q \in \mathcal{B}(\mathcal{H})$ such that

$$\text{avg}_{\hat{\xi}} \left| \text{tr}(Q\hat{\xi}) - \text{tr}(Q\xi) \right| \geq O(1) .$$

Proof. By the assumptions on the sequences $\{U_N |\psi_N\rangle\}_N$ and $\{\hat{U}_N |\psi_N\rangle\}_N$, for every sufficiently large N , $U_N |\psi_N\rangle$ and $\hat{U}_N |\psi_N\rangle$ satisfy 1. and 2. in Theorem 3.

The result (A.31) by definition guarantees that

$$\text{avg}_{\hat{\xi}} \left| \langle \xi | V_N^\dagger V_N | \hat{\xi} \rangle \right| \leq 1 - \alpha - \text{negl}(N). \tag{B.1}$$

Taking the limit of both sides and using the definition of an asymptotically isometric code in Assumption 2, we see that the two bulk states $|\xi\rangle$ and $|\hat{\xi}\rangle$ obey the inequality as well. Note that to commute the average over $\hat{\xi}$ with the $N \rightarrow \infty$ limit we had to use the dominated convergence theorem.

By the Fuchs-van de Graaf inequality, the overlap of the two states $\langle \xi | \hat{\xi} \rangle$ upper and lower bounds their one-norm distance:

$$1 - \sqrt{F(\xi, \hat{\xi})} \leq \frac{1}{2} \|\xi - \hat{\xi}\|_1 \leq \sqrt{1 - F(\xi, \hat{\xi})} . \tag{B.2}$$

When ξ is a pure state, the lower bound can be strengthened:

$$\frac{1}{2} \|\xi - \hat{\xi}\|_1 \geq 1 - F(\xi, \hat{\xi}) . \tag{B.3}$$

Averaging this inequality over $\hat{\xi}$ and using (B.1), we see that

$$\text{avg}_{\hat{\xi}} \|\xi - \hat{\xi}\|_1 \geq 2\alpha . \tag{B.4}$$

The one-norm distance between two states on a general algebra $\mathcal{B}(\mathcal{H})$ may be written as

$$\|\xi - \hat{\xi}\|_1 = \sup_{a \in \mathcal{B}(\mathcal{H}) : \|a\|_\infty \leq 1} |\text{tr}(a\xi) - \text{tr}(a\hat{\xi})| . \tag{B.5}$$

where the operators in the supremum are restricted to be self-adjoint. Combining (B.5) and (B.4), we see that there exists some operator $Q \in \mathcal{B}(\mathcal{H})$, depending on \hat{U} and therefore $\hat{\xi}$, such that

$$\text{avg}_{\hat{\xi}} |\text{tr}(Q\hat{\xi}) - \text{tr}(Q\xi)| \geq 2\alpha . \tag{B.6}$$

In particular, we get that there exists an operator whose expectation values for $\hat{\xi}$ and ξ must differ by a constant. This concludes the proof of Corollary 3.1. \square

Data Availability Statement. This article has no associated data or the data will not be deposited.

Code Availability Statement. This article has no associated code or the code will not be deposited.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] R. Penrose, *Gravitational collapse and space-time singularities*, *Phys. Rev. Lett.* **14** (1965) 57 [[INSPIRE](#)].
- [2] S.W. Hawking, *Singularities in the universe*, *Phys. Rev. Lett.* **17** (1966) 444 [[INSPIRE](#)].
- [3] S.W. Hawking and R. Penrose, *The singularities of gravitational collapse and cosmology*, *Proc. Roy. Soc. Lond. A* **314** (1970) 529 [[INSPIRE](#)].
- [4] S.W. Hawking, *Gravitational radiation from colliding black holes*, *Phys. Rev. Lett.* **26** (1971) 1344 [[INSPIRE](#)].
- [5] D. Marolf and J. Polchinski, *Gauge/Gravity Duality and the Black Hole Interior*, *Phys. Rev. Lett.* **111** (2013) 171301 [[arXiv:1307.4706](#)] [[INSPIRE](#)].
- [6] R. Gregory and R. Laflamme, *The instability of charged black strings and p-branes*, *Nucl. Phys. B* **428** (1994) 399 [[hep-th/9404071](#)] [[INSPIRE](#)].
- [7] L. Lehner and F. Pretorius, *Black Strings, Low Viscosity Fluids, and Violation of Cosmic Censorship*, *Phys. Rev. Lett.* **105** (2010) 101102 [[arXiv:1006.5960](#)] [[INSPIRE](#)].
- [8] R. Gregory, *The Gregory–Laflamme instability*, in *Black holes in higher dimensions*, G.T. Horowitz ed., Cambridge University Press (2012), p. 29–43 [[arXiv:1107.5821](#)] [[INSPIRE](#)].
- [9] G.T. Horowitz, J.E. Santos and B. Way, *Evidence for an Electrifying Violation of Cosmic Censorship*, *Class. Quant. Grav.* **33** (2016) 195007 [[arXiv:1604.06465](#)] [[INSPIRE](#)].
- [10] T. Crisford and J.E. Santos, *Violating the Weak Cosmic Censorship Conjecture in Four-Dimensional Anti-de Sitter Space*, *Phys. Rev. Lett.* **118** (2017) 181101 [[arXiv:1702.05490](#)] [[INSPIRE](#)].
- [11] T. Crisford, G.T. Horowitz and J.E. Santos, *Testing the Weak Gravity - Cosmic Censorship Connection*, *Phys. Rev. D* **97** (2018) 066005 [[arXiv:1709.07880](#)] [[INSPIRE](#)].
- [12] G.T. Horowitz and J.E. Santos, *Further evidence for the weak gravity — cosmic censorship connection*, *JHEP* **06** (2019) 122 [[arXiv:1901.11096](#)] [[INSPIRE](#)].
- [13] Å. Folkestad, *Penrose Inequality as a Constraint on the Low Energy Limit of Quantum Gravity*, *Phys. Rev. Lett.* **130** (2023) 121501 [[arXiv:2209.00013](#)] [[INSPIRE](#)].
- [14] F.C. Eperon, B. Ganchev and J.E. Santos, *Plausible scenario for a generic violation of the weak cosmic censorship conjecture in asymptotically flat four dimensions*, *Phys. Rev. D* **101** (2020) 041502 [[arXiv:1906.11257](#)] [[INSPIRE](#)].
- [15] J. Bardeen, *Non-singular general relativistic gravitational collapse*, in the proceedings of the 5th International Conference on Gravitation and the Theory of Relativity, (1968), p. 87.
- [16] N. Arkani-Hamed, L. Motl, A. Nicolis and C. Vafa, *The String landscape, black holes and gravity as the weakest force*, *JHEP* **06** (2007) 060 [[hep-th/0601001](#)] [[INSPIRE](#)].
- [17] N. Engelhardt and G.T. Horowitz, *Holographic argument for the Penrose inequality in AdS spacetimes*, *Phys. Rev. D* **99** (2019) 126009 [[arXiv:1903.00555](#)] [[INSPIRE](#)].
- [18] K.S. Thorne, *Nonspherical gravitational collapse: a short review*, W.H. Freeman & Co. (1972), pp. 231–258. [[INSPIRE](#)].

- [19] S.L. Shapiro and S.A. Teukolsky, *Formation of naked singularities: the violation of cosmic censorship*, *Phys. Rev. Lett.* **66** (1991) 994 [INSPIRE].
- [20] M.W. Choptuik and F. Pretorius, *Ultra Relativistic Particle Collisions*, *Phys. Rev. Lett.* **104** (2010) 111101 [arXiv:0908.1780] [INSPIRE].
- [21] P.M. Chesler and W. van der Schee, *Early thermalization, hydrodynamics and energy loss in AdS/CFT*, *Int. J. Mod. Phys. E* **24** (2015) 1530011 [arXiv:1501.04952] [INSPIRE].
- [22] A. Almheiri, D. Marolf, J. Polchinski and J. Sully, *Black Holes: Complementarity or Firewalls?*, *JHEP* **02** (2013) 062 [arXiv:1207.3123] [INSPIRE].
- [23] A. Almheiri et al., *An Apologia for Firewalls*, *JHEP* **09** (2013) 018 [arXiv:1304.6483] [INSPIRE].
- [24] R. Emparan, *Predictivity lost, predictivity regained: a Miltonian cosmic censorship conjecture*, *Int. J. Mod. Phys. D* **29** (2020) 2043021 [arXiv:2005.07389] [INSPIRE].
- [25] A.C. Wall, *Maximin Surfaces, and the Strong Subadditivity of the Covariant Holographic Entanglement Entropy*, *Class. Quant. Grav.* **31** (2014) 225007 [arXiv:1211.3494] [INSPIRE].
- [26] N. Engelhardt and A.C. Wall, *No Simple Dual to the Causal Holographic Information?*, *JHEP* **04** (2017) 134 [arXiv:1702.01748] [INSPIRE].
- [27] N. Engelhardt and A.C. Wall, *Coarse Graining Holographic Black Holes*, *JHEP* **05** (2019) 160 [arXiv:1806.01281] [INSPIRE].
- [28] N. Engelhardt, G. Penington and A. Shahbazi-Moghaddam, *Finding pythons in unexpected places*, *Class. Quant. Grav.* **39** (2022) 094002 [arXiv:2105.09316] [INSPIRE].
- [29] N. Engelhardt, G. Penington and A. Shahbazi-Moghaddam, *A world without pythons would be so simple*, *Class. Quant. Grav.* **38** (2021) 234001 [arXiv:2102.07774] [INSPIRE].
- [30] C. Akers et al., *The black hole interior from non-isometric codes and complexity*, *JHEP* **06** (2024) 155 [arXiv:2207.06536] [INSPIRE].
- [31] A. Bouland, B. Fefferman and U. Vazirani, *Computational pseudorandomness, the wormhole growth paradox, and constraints on the AdS/CFT duality*, arXiv:1910.14646 [INSPIRE].
- [32] L. Yang and N. Engelhardt, *The Complexity of Learning (Pseudo)random Dynamics of Black Holes and Other Chaotic Systems*, arXiv:2302.11013 [INSPIRE].
- [33] Z. Ji, Y.-K. Liu, and F. Song, *Pseudorandom quantum states*, in H. Shacham and A. Boldyreva eds., *Advances in Cryptology — CRYPTO 2018*, Springer International Publishing (2018), pp. 126–152.
- [34] P. Hayden and J. Preskill, *Black holes as mirrors: quantum information in random subsystems*, *JHEP* **09** (2007) 120 [arXiv:0708.4025] [INSPIRE].
- [35] I.H. Kim and J. Preskill, *Complementarity and the unitarity of the black hole S-matrix*, *JHEP* **02** (2023) 233 [arXiv:2212.00194] [INSPIRE].
- [36] S.H. Shenker and D. Stanford, *Black holes and the butterfly effect*, *JHEP* **03** (2014) 067 [arXiv:1306.0622] [INSPIRE].
- [37] D.A. Roberts, D. Stanford and L. Susskind, *Localized shocks*, *JHEP* **03** (2015) 051 [arXiv:1409.8180] [INSPIRE].
- [38] J. Maldacena, S.H. Shenker and D. Stanford, *A bound on chaos*, *JHEP* **08** (2016) 106 [arXiv:1503.01409] [INSPIRE].
- [39] J.S. Cotler et al., *Black Holes and Random Matrices*, *JHEP* **05** (2017) 118 [Erratum *ibid.* **09** (2018) 002] [arXiv:1611.04650] [INSPIRE].

- [40] N. Engelhardt and Å. Folkestad, *Holography abhors visible trapped surfaces*, *JHEP* **07** (2021) 066 [[arXiv:2012.11445](#)] [[INSPIRE](#)].
- [41] A.R. Brown, H. Gharibyan, G. Penington and L. Susskind, *The Python's Lunch: geometric obstructions to decoding Hawking radiation*, *JHEP* **08** (2020) 121 [[arXiv:1912.00228](#)] [[INSPIRE](#)].
- [42] N. Engelhardt and A.C. Wall, *Quantum Extremal Surfaces: Holographic Entanglement Entropy beyond the Classical Regime*, *JHEP* **01** (2015) 073 [[arXiv:1408.3203](#)] [[INSPIRE](#)].
- [43] R. Bousso, Z. Fisher, S. Leichenauer and A.C. Wall, *Quantum focusing conjecture*, *Phys. Rev. D* **93** (2016) 064044 [[arXiv:1506.02669](#)] [[INSPIRE](#)].
- [44] A. Almheiri, R. Mahajan and J. Maldacena, *Islands outside the horizon*, [arXiv:1910.11077](#) [[INSPIRE](#)].
- [45] R. Bousso and G. Penington, *Islands far outside the horizon*, *JHEP* **11** (2024) 164 [[arXiv:2312.03078](#)] [[INSPIRE](#)].
- [46] V.A. Belinsky, I.M. Khalatnikov and E.M. Lifshitz, *Oscillatory approach to a singular point in the relativistic cosmology*, *Adv. Phys.* **19** (1970) 525 [[INSPIRE](#)].
- [47] V. A. Belinsky and I. M. Khalatnikov, *General Solution of the Gravitational Equations with a Physical Singularity*, *Sov.Phys.JETP* **30** (1970) 1174.
- [48] E. Shaghoulian and H. Wang, *Timelike BKL singularities and chaos in AdS/CFT*, *Class. Quant. Grav.* **33** (2016) 125020 [[arXiv:1601.02599](#)] [[INSPIRE](#)].
- [49] T. Faulkner and M. Li, *Asymptotically isometric codes for holography*, [arXiv:2211.12439](#) [[INSPIRE](#)].
- [50] S. Leutheusser and H. Liu, *Subregion-subalgebra duality: emergence of space and time in holography*, [arXiv:2212.13266](#) [[INSPIRE](#)].
- [51] N. Engelhardt et al., *Spoofing Entanglement in Holography*, [arXiv:2407.14589](#) [[INSPIRE](#)].
- [52] A.C. Wall, *A proof of the generalized second law for rapidly-evolving Rindler horizons*, *Phys. Rev. D* **82** (2010) 124019 [[arXiv:1007.1493](#)] [[INSPIRE](#)].
- [53] A. Shahbazi-Moghaddam, *Restricted quantum focusing*, *Phys. Rev. D* **109** (2024) 066023 [[arXiv:2212.03881](#)] [[INSPIRE](#)].
- [54] N. Engelhardt, G. Penington and A. Shahbazi-Moghaddam, *Twice upon a time: timelike-separated quantum extremal surfaces*, *JHEP* **01** (2024) 033 [[arXiv:2308.16226](#)] [[INSPIRE](#)].
- [55] C. Akers, N. Engelhardt, G. Penington and M. Usatyuk, *Quantum Maximin Surfaces*, *JHEP* **08** (2020) 140 [[arXiv:1912.02799](#)] [[INSPIRE](#)].
- [56] N. Engelhardt and A.C. Wall, *Decoding the Apparent Horizon: Coarse-Grained Holographic Entropy*, *Phys. Rev. Lett.* **121** (2018) 211301 [[arXiv:1706.02038](#)] [[INSPIRE](#)].
- [57] A. Hamilton, D.N. Kabat, G. Lifschytz and D.A. Lowe, *Local bulk operators in AdS/CFT: a boundary view of horizons and locality*, *Phys. Rev. D* **73** (2006) 086003 [[hep-th/0506118](#)] [[INSPIRE](#)].
- [58] A. Hamilton, D.N. Kabat, G. Lifschytz and D.A. Lowe, *Holographic representation of local bulk operators*, *Phys. Rev. D* **74** (2006) 066009 [[hep-th/0606141](#)] [[INSPIRE](#)].
- [59] A. Hamilton, D.N. Kabat, G. Lifschytz and D.A. Lowe, *Local bulk operators in AdS/CFT: a holographic description of the black hole interior*, *Phys. Rev. D* **75** (2007) 106001 [Erratum *ibid.* **75** (2007) 129902] [[hep-th/0612053](#)] [[INSPIRE](#)].

- [60] N. Engelhardt and H. Liu, *Algebraic ER=EPR and complexity transfer*, *JHEP* **07** (2024) 013 [[arXiv:2311.04281](#)] [[INSPIRE](#)].
- [61] D. Harlow and P. Hayden, *Quantum Computation vs. Firewalls*, *JHEP* **06** (2013) 085 [[arXiv:1301.4504](#)] [[INSPIRE](#)].
- [62] I.H. Kim, E. Tang and J. Preskill, *The ghost in the radiation: robust encodings of the black hole interior*, *JHEP* **06** (2020) 031 [[arXiv:2003.05451](#)] [[INSPIRE](#)].
- [63] H.-Y. Huang, S. Chen and J. Preskill, *Learning to Predict Arbitrary Quantum Processes*, *PRX Quantum* **4** (2023) 040337 [[arXiv:2210.14894](#)] [[INSPIRE](#)].
- [64] D. Marolf, *Microcanonical Path Integrals and the Holography of small Black Hole Interiors*, *JHEP* **09** (2018) 114 [[arXiv:1808.00394](#)] [[INSPIRE](#)].
- [65] D.A. Roberts and B. Yoshida, *Chaos and complexity by design*, *JHEP* **04** (2017) 121 [[arXiv:1610.04903](#)] [[INSPIRE](#)].
- [66] R. Gregory and R. Laflamme, *Black strings and p-branes are unstable*, *Phys. Rev. Lett.* **70** (1993) 2837 [[hep-th/9301052](#)] [[INSPIRE](#)].
- [67] R.P. Geroch, *Local characterization of singularities in general relativity*, *J. Math. Phys.* **9** (1968) 450 [[INSPIRE](#)].
- [68] Y.B. Zeldovich, *A New Type of Radioactive Decay: Gravitational Annihilation of Baryons*, *Phys. Lett. A* **59** (1976) 254 [[INSPIRE](#)].
- [69] L.F. Abbott and M.B. Wise, *Wormholes and Global Symmetries*, *Nucl. Phys. B* **325** (1989) 687 [[INSPIRE](#)].
- [70] R. Kallosh, A.D. Linde, D.A. Linde and L. Susskind, *Gravity and global symmetries*, *Phys. Rev. D* **52** (1995) 912 [[hep-th/9502069](#)] [[INSPIRE](#)].
- [71] M.W. Choptuik, *Universality and scaling in gravitational collapse of a massless scalar field*, *Phys. Rev. Lett.* **70** (1993) 9 [[INSPIRE](#)].
- [72] D. Christodoulou, *Examples of naked singularity formation in the gravitational collapse of a scalar field*, *Annals Math.* **140** (1994) 607 [[INSPIRE](#)].
- [73] J. Santos, *Connecting the weak gravity conjecture to the weak cosmic censorship*, in *Strings 2018*, Okinawa, Japan, 25–29 June (2018).
- [74] R.P. Geroch and G.T. Horowitz, *Global structure of spacetimes*, in *General Relativity: an Einstein Centenary Survey*, Cambridge University Press (1979), pp. 212–293 [[INSPIRE](#)].
- [75] E. Giorgi, S. Klainerman and J. Szeftel, *Wave equations estimates and the nonlinear stability of slowly rotating Kerr black holes*, [arXiv:2205.14808](#) [[INSPIRE](#)].
- [76] F. Haake, *Quantum signatures of chaos*, Springer (1991).
- [77] D.N. Page, *Average entropy of a subsystem*, *Phys. Rev. Lett.* **71** (1993) 1291 [[gr-qc/9305007](#)] [[INSPIRE](#)].
- [78] J. Choi et al., *Preparing random states and benchmarking with many-body quantum chaos*, *Nature* **613** (2023) 468 [[arXiv:2103.03535](#)] [[INSPIRE](#)].
- [79] W.W. Ho and S. Choi, *Exact emergent quantum state designs from quantum chaotic dynamics*, *Phys. Rev. Lett.* **128** (2022) 060601 [[arXiv:2109.07491](#)] [[INSPIRE](#)].
- [80] R. Bousso, V. Chandrasekaran and A. Shahbazi-Moghaddam, *From black hole entropy to energy-minimizing states in QFT*, *Phys. Rev. D* **101** (2020) 046001 [[arXiv:1906.05299](#)] [[INSPIRE](#)].

- [81] N. Engelhardt and S. Fischetti, *Losing the IR: a Holographic Framework for Area Theorems*, *Class. Quant. Grav.* **36** (2019) 035008 [[arXiv:1805.08891](#)] [[INSPIRE](#)].
- [82] A.M. Kubicki, A. May and D. Pérez-Garcia, *Constraints on physical computers in holographic spacetimes*, *SciPost Phys.* **16** (2024) 024 [[arXiv:2304.09900](#)] [[INSPIRE](#)].
- [83] I. Bena et al., *Smooth horizonless geometries deep inside the black-hole regime*, *Phys. Rev. Lett.* **117** (2016) 201601 [[arXiv:1607.03908](#)] [[INSPIRE](#)].
- [84] J. Cotler and A. Strominger, *The Universe as a Quantum Encoder*, [arXiv:2201.11658](#) [[INSPIRE](#)].
- [85] J. Cotler and K. Jensen, *Isometric Evolution in de Sitter Quantum Gravity*, *Phys. Rev. Lett.* **131** (2023) 211601 [[arXiv:2302.06603](#)] [[INSPIRE](#)].
- [86] P. Ananth, A. Gulati, F. Kaleoglu and Y.-T. Lin, *Pseudorandom Isometries*, *Lect. Notes Comput. Sci.* **14654** (2024) 226 [[arXiv:2311.02901](#)] [[INSPIRE](#)].
- [87] S. Lloyd, *Ultimate physical limits to computation*, *Nature* **406** (2000) 1047 [[INSPIRE](#)].
- [88] G. Katoch, J. Ren and S.R. Roy, *Quantum complexity and bulk timelike singularities*, *JHEP* **12** (2023) 085 [[arXiv:2303.02752](#)] [[INSPIRE](#)].
- [89] E. Meckes and M. Meckes, *Spectral measures of powers of random matrices*, [arXiv:1210.2681](#).
- [90] Q. Wang et al., *Quantum Algorithm for Fidelity Estimation*, *IEEE Trans. Info. Theor.* **69** (2023) 273 [[arXiv:2103.09076](#)] [[INSPIRE](#)].