

A World without Pythons would be so Simple

Netta Engelhardt,¹ Geoff Penington,^{2,3} and Arvin Shahbazi-Moghaddam⁴

¹*Center for Theoretical Physics, Massachusetts Institute of Technology,
Cambridge, MA 02139, USA*

²*Center for Theoretical Physics and Department of Physics,
University of California, Berkeley, CA 94720, U.S.A. and*

³*Institute for Advanced Study, 1 Einstein Dr, Princeton, NJ 08540, U.S.A.*

⁴*Stanford Institute for Theoretical Physics,
Stanford University, Stanford, CA 94305 USA*

engelnn@mit.edu, geoffp@berkeley.edu, arvinshm@gmail.com

ABSTRACT: We show that bulk operators lying between the outermost extremal surface and the asymptotic boundary admit a simple boundary reconstruction in the classical limit. This is the converse of the Python’s lunch conjecture, which proposes that operators with support between the minimal and outermost (quantum) extremal surfaces – e.g. the interior Hawking partners – are highly complex. Our procedure for reconstructing this “simple wedge” is based on the HKLL construction, but uses causal bulk propagation of perturbed boundary conditions on Lorentzian timefolds to expand the causal wedge as far as the outermost extremal surface. As a corollary, we establish the Simple Entropy proposal for the holographic dual of the area of a marginally trapped surface as well as a similar holographic dual for the outermost extremal surface. We find that the simple wedge is dual to a particular coarse-grained CFT state, obtained via averaging over all possible Python’s lunches. An efficient quantum circuit converts this coarse-grained state into a “simple state” that is indistinguishable in finite time from a state with a local modular Hamiltonian. Under certain circumstances, the simple state modular Hamiltonian generates an exactly local flow; we interpret this result as a holographic dual of black hole uniqueness.

Contents

1	Introduction	1
2	Which Wedge?	9
3	Two Dimensions	14
4	The Perturbation	17
5	Zigzag	24
6	Simplicity Killed the Python	26
6.1	Time-Folded Simple Entropy	26
6.2	The Simple State	28
7	Discussion	31
7.1	Boundary-Anchored Surfaces	31
7.2	Robustness under Quantum Corrections	32
7.3	Asymptotically Flat Spacetimes	33
7.4	Complexity Censorship	33
A	uv and vi constraints	34
B	Alternative derivation of the vv constraint	35

1 Introduction

Recent developments in the black hole information frontier have pointed to a holographic geometrization of the degrees of freedom of the Hawking radiation [1–6]. For an AdS black hole evaporating into a bath, the “entanglement wedge of the radiation” after the Page time includes a large part of the black hole interior, bounded by the minimal quantum extremal surface (QES) [7].

This geometric description of the information naturally accounts for both the Page curve [8] and the Hayden-Preskill decoding criterion [9]. It also leads to a geometrical explanation for the expectation of Harlow-Hayden [10] that decoding Hawking radiation




Figure 1. An illustration of the Python’s lunch. On the left, the geometry of a Cauchy slice featuring the titular python’s lunch between the two constrictions, the dominant QES $X_{\min}$ and the non-minimal QES X , both of which lie behind the causal surface C . On the right, a spacetime diagram of the same.

should be exponentially complex. Even though the interior degrees of freedom lie on the radiation side of the minimal QES and so lie within the radiation entanglement wedge, they are still hidden behind a nonminimal QES; in the case of the single-sided black hole, the nonminimal QES is simply the empty set. The region between the nonminimal and minimal extremal surfaces was dubbed “the Python’s lunch” in [6], because appropriate Cauchy slices in the bulk (quantum) geometry have a constriction at each extremal surface, together with a bulge in the middle (the eponymous “lunch”). See Fig. 1 for an illustration.

The claim of [6] was that any bulk operator with support in the interior of a Python’s lunch should be exponentially difficult to decode, with an exponent that is controlled by the size of the bulge and grows as $O(1/G_N)$ in the semiclassical limit. The justification for this conjecture was based primarily on tensor network toy models, where the fastest known protocols for decoding operators inside a lunch use a Grover-search-based algorithm that takes exponential time.

An additional important consistency check comes from the quantum focusing conjecture [11], which is the quantum avatar of classical gravitational lensing when the null energy condition is violated due to quantum corrections. Quantum focusing and global hyperbolicity ensure that no causal semiclassical Lorentzian evolution of the bulk geometry can result in causal communication from behind a Python’s lunch to the asymptotic boundary, no matter what the asymptotic boundary conditions are. If such communication were possible, reconstruction of information from within the lunch could be implemented using only boundary time evolution with simple sources and the extrapolate dictionary relating bulk observables at the asymptotic boundary to local boundary operators. From a boundary perspective, this is a very simple procedure compared to the exponential complexity claimed to be necessary for operator

reconstruction inside the lunch.

Without input from nonperturbative quantum gravity (such as entanglement wedge reconstruction), simple reconstruction using only low-complexity, causally-propagating operators and sources is all that semiclassical Lorentzian gravity is capable of: i.e. all that semiclassical gravity “sees”. Calculations and arguments that rely exclusively on semiclassical gravity with no further input (e.g. Hawking’s original calculation) are thus certainly restricted by the Python’s lunch proposal to recover no more than the domain of dependence between the outermost quantum extremal surface – the “appetizer” of the lunch – and the asymptotic boundary.

So how much does purely semiclassical gravity *actually* recover? On the one hand, if the Python’s lunch conjecture is true, it is natural to expect that simple reconstruction can in fact obtain the entire bulk up to the outermost extremal surface. This “converse” to the Python’s lunch conjecture is certainly true in tensor network toy models where anything not in a lunch can be reconstructed using a simple unitary circuit. Any gap in gravitational theories between the simply reconstructible region and the start of the lunch would therefore be somewhat puzzling and demand explanation. On the other hand, simple reconstruction appears to be little more than a glorified version of the HKLL procedure [12–14], which is supposed to recover just the so-called causal wedge: the region that can simultaneously send and receive signals from the asymptotic boundary. And generically the causal wedge and the outermost extremal wedge do not coincide.

To clarify this conundrum, let us first briefly review HKLL here, as it will be instrumental for our work in this paper. The HKLL procedure is a reconstruction protocol for bulk matter on a fixed background, in which the bulk fields (which can include gravitons [15, 16], $\mathcal{O}(1/N)$ corrections [16, 17], interactions [18], and higher spins [16, 19]) are obtained via a non-standard Cauchy evolution from their boundary counterparts (related to them via the extrapolate dictionary). Quantitatively,

$$\phi(x) = \int dX K(x; X) \mathcal{O}(X), \tag{1.1}$$

where $K(x; X)$ is a smearing function that depends on the spacetime geometry supported on the set of boundary points spacelike-separated to x .¹ The sense in which HKLL is “simple” is evident: from a boundary perspective it consists of boundary

¹Note that the validity of this non-standard “rotated” Cauchy problem is far from well-established (though see [20, 21] for proofs in certain cases). Our purpose here is not to put HKLL on a firm footing, but rather to show that reconstruction of the simple wedge is as simple as HKLL. We shall therefore assume HKLL, but any other simple reconstruction procedure for operators in the causal wedge would do just as well for our purposes.

time evolution with local sources turned on. And local Hamiltonian evolution can be simulated efficiently using a quantum circuit.

The immediate prediction therefore, as expressed in [22], is that HKLL can reconstruct operators within the causal wedge. We might then expect that the simply reconstructible region – which we shall henceforth refer to as the *simple wedge* – is to be identified with the causal wedge. Since the causal wedge is always a subset of the outermost quantum extremal wedge [7] and is generically a proper subset ², this leads to the undesirable no-man’s land between the simple wedge and the outermost quantum extremal surface.

To see deeper into the bulk, we need to expand the causal wedge via the addition of simple boundary sources as proposed in [24, 25]. In the very special case when the gap between the outermost extremal surface and the causal wedge is Planckian, [26] showed that certain causal unitaries produce just enough backreaction to maximally expand the causal past or future. However, in generic spacetimes the gap region is non-empty even in the classical limit and can in fact be arbitrarily large. It was conjectured in [24, 25] that it should be possible to fine-tune simple sources in order to “turn off” any extant focusing and so expand the causal wedge up to an apparent horizon, all without violating the null energy condition.

A central result of this paper is an explicit, constructive derivation of this fact – in the limit where the bulk dynamics are classical and with a variety of matter fields. Furthermore, by evolving backwards and forwards in time using timefolds with different boundary conditions, one can continue to iteratively expand the causal wedge, from apparent horizon to apparent horizon, all the way to the outermost extremal surface. Combining this result with ordinary HKLL leads to simple reconstructions of arbitrary operators in the outermost extremal wedge.

It is easy to see how this works in the case of Jackiw-Teitelboim gravity [27, 28] minimally coupled to a (classical) massless free scalar. In this setup, the matter factorizes into left and right movers; by changing the boundary conditions, we may “absorb” the right movers and turn off focusing on the future event horizon; this pushes the future event horizon backwards. We can then repeat the same procedure for the past horizon by evolving backwards in time; this will now push the past causal horizon backwards. The shift will have likely revealed additional left-movers, so the procedure needs to be iteratively repeated until it converges on a stationary bifurcate horizon.

²As shown in [23], quantum effects can allow the causal wedge to be outside the outermost extremal wedge when defined using time evolution couples the asymptotic boundary to an auxiliary system. However, this is only true if we do not include the auxiliary coupled system when defining the outermost extremal wedge. When comparing apples to apples by doing so, one indeed finds that the causal wedge is still contained in the outermost extremal wedge.




Figure 2. A caricature of the procedure used to push the causal wedge towards the appertizer in JT gravity coupled to a classical massless scalar. The leftmost panel is the original spacetime including left- and right-movers with reflecting boundary conditions. The causal surface of the right boundary is C_1 . In the middle panel, the left movers have been turned off, which causes the future event horizon to shift inwards. The resulting causal surface is C_2 , which is null-separated from C_1 . The final panel shows that the right movers have been turned off, which causes the past event horizon to move inwards, shifting the causal surface to C_3 . This shift reveals new left-movers in the causal wedge, which will have to be removed in subsequent zigzags along the past and future event horizons.

This is illustrated in Fig. 2. The generalization to higher dimensions is significantly more technically challenging – rather than removing sources of focusing entirely, it is more practical to “stretch out” the focusing over the causal horizon and so dilute its effect – but the essential intuition is the same.

The original motivation in [24] for attempting to expand the causal wedge using simple sources was to understand the holographic dual of the *simple entropy*, defined as a maximization of the von Neumann entropy over all CFT density matrices with the same one-point functions – with arbitrary time-ordered simple sources turned on after some initial time t – as the original CFT state. In other words, the simple entropy coarse-grains over all of the details of the state, except for simple observables that can be measured in the future of the initial time t . It was conjectured in [24, 25] that the simple entropy is the boundary dual of the *outer entropy*, a bulk quantity that coarse-grains over the geometry behind the outermost apparent horizon null-separated from the boundary at time t ; it is equal to (one quarter of) the area of the apparent horizon. As a corollary of the results discussed above, we prove that this conjecture is indeed true whenever the bulk physics can be treated classically.

What if we generalize the definition of the simple entropy to allow not just time-ordered insertions of simple operators, but insertions on arbitrary timefolds (and at arbitrary time)? In this case, there is no obstacle to seeing behind apparent horizons.




Figure 3. The coarse-graining procedure of [24, 25] as applied to the outermost extremal surface. The spacetime behind the outermost extremal surface X is discarded and replaced with a CPT conjugate of the outermost extremal wedge. The rest of the spacetime is generated by standard Cauchy evolution.

By evolving the state backwards (and forwards) in time and then turning simple sources, it is possible to causally alter the spacetime near the apparent horizon, changing its location and “seeing” degrees of freedom that were originally hidden behind it. As per the discussion above, the first obstruction that cannot be bypassed in this way is the outermost extremal surface. Indeed, our results demonstrate that the simple entropy with arbitrary timefolds allowed is holographically dual to the area of the outermost extremal surface. Similarly, the density matrix ρ_{coarse} whose von Neumann entropy is the simple entropy with timefolds allowed reconstructs exactly the entire outermost extremal wedge and no more. In fact, we can actually construct a complete spacetime in which the outermost extremal wedge is the entire entanglement wedge of one connected asymptotic boundary obtained using the spacetime doubling procedure of [24, 25]; thus ρ_{coarse} is the actual CFT state dual to the canonical purification as proposed in [24, 25] and proven in [29]. See Fig. 3.

An immediate application of our result is then the construction of the CFT dual to the simple wedge in the final spacetime where the causal and outermost extremal wedges coincide. That is, this dual can be produced from ρ_{coarse} via a set of simple operations, with the dual bulk result being a two-sided black hole in which the bifurcation surface is extremal. The significance of this statement is manifold: we prove that the causal and entanglement wedges coincide if and only if the CFT state has a local modular Hamiltonian, which shows that finite time measurements cannot tell that the modular flow generated by the simple state is not local. In spacetimes with sufficient isotropy, the simple wedge CFT modular flow is in fact exactly local. This is analogous to a type of no-hair theorem: the set of holographic black holes with a stationary bifurcation surface is identical to the highly limited set of states with local modular Hamiltonians in the dual CFT.

From the perspective of holographic complexity, we may therefore interpret the

absence of a Python’s lunch in the dual theory as the CFT state being related by a simple circuit to a rather special state with local modular flow (or at least indistinguishable from local in finite time). The world, it would seem, is rarely simple; pythons are ubiquitous. An explicit example of how a python might spring on an unsuspecting holographer in what would *prima facie* appear to be a python-less spacetime will be provided in our upcoming work [30].

The paper is structured as follows. In Section 2, we define the outermost extremal wedge and the simple wedge, and we prove that the former is well-defined. In Section 3 we showcase our procedure for the simple case of JT gravity coupled to a massless (classical) scalar. In Section 4, we prepare the perturbation that moves the causal horizon backwards along a the future event horizon in higher dimensional gravity (with arbitrary, null energy condition-satisfying matter), and we prove that the required perturbation satisfies the constraint equations. Section 5 describes the zigzag portion of the procedure and completes the proof. We discuss the implications of our results, from the dual to the area of the outermost extremal surface to the nature of the simple state, in Section 6. We finish with a discussion of generalizations and other implications in Section 7.

Assumptions and Conventions: The bulk spacetime (M, g) is assumed to be classical with the dynamics governed by the Einstein field equation, i.e. we work in the large- N and large λ limit of AdS/CFT except where otherwise stated. We will assume the AdS analogue of global hyperbolicity [31]. We also assume that the initial spacetime under consideration is one that satisfies the Null Energy Condition (NEC):

$$T_{ab}k^ak^b \geq 0 \tag{1.2}$$

where T_{ab} is the stress energy tensor and k^a is any null vector. We will demonstrate that our perturbations of the spacetime maintain the NEC. All other conventions are as in [32] unless otherwise stated.

- We shall use $J^\pm$ to refer to the bulk causal future and past and $I^\pm$ to refer to the bulk chronological future and past. Given a closed achronal set S , we use $D[S]$ to denote its domain of dependence, which we shall take to contain its boundary, as in [32]. $D^+[S]$ and $D^-[S]$ refer to the future and past components of the domain of dependence.
- Hypersurfaces will refer to codimension-one embedded submanifolds of arbitrary signature.

- By a “surface” we will always mean an achronal, codimension-two embedded submanifold which is Cauchy-splitting [33]. Two surfaces σ_1 and σ_2 are homologous whenever there exists a hypersurface H such that $\partial H = \sigma_1 \cup \sigma_2$. We will be primarily interested in surfaces homologous to (partial) Cauchy slices of the asymptotic boundary (CFT (sub)regions).
- Let Σ be a Cauchy slice containing a surface σ homologous to a boundary (sub)region R . By definition, σ splits Σ into two disjoint components that we will denote $\text{Int}_\Sigma[\sigma]$ and $\text{Out}_\Sigma[\sigma]$, where the conformal completion of the latter contains the boundary subregion R . We define $W_\sigma \equiv D[\text{Out}_\Sigma[\sigma]]$, the outer wedge of σ . Similarly, we define $I_\sigma \equiv D[\text{Int}_\Sigma[\sigma]]$, the inner wedge of σ . See also [25].
- For a smooth surface σ homologous to a boundary (sub)region, we denote by k^a and ℓ^a the unique future-directed orthogonal null vector fields on the C^1 subsets of σ pointing towards $\text{Out}_\Sigma[\sigma]$ and towards $\text{Int}_\Sigma[\sigma]$ respectively.
- We define $\partial^+ W_\sigma = \partial D^+[\text{Out}_\Sigma[\sigma]]$ and $\partial^- W_\sigma = \partial D^-[\text{Out}_\Sigma[\sigma]]$. When σ is smooth, $\partial^+ W_\sigma$ and $\partial^- W_\sigma$ can be constructed by firing null congruences starting from k^a and ℓ^a , terminating the congruence at caustics and non-local self-intersections [32, 34].
- Given any orthogonal null vector field k^a on surface σ , $\theta_{(k)}$ denotes the expansion of σ along k^a .³ We will refer to the following types of σ based on its expansions:
 - A compact σ is trapped if $\theta_{(k)} < 0$ and $\theta_{(\ell)} < 0$, and marginally trapped if $\theta_{(\ell)} < 0$ and $\theta_{(k)} = 0$.
 - σ is *extremal* if $\theta_{(k)} = 0$ and $\theta_{(\ell)} = 0$. By linearity, σ is then stationary under deformations along any direction.
- The future and past causal horizons associated to any boundary spacetime region $R \subset \mathcal{I}$ are defined as $\partial J^- [R]$ and $\partial J^+ [R]$ respectively. By convention, we use k^a and ℓ^a to refer to the generators of the future and past horizons respectively. More specifically, the future and past *event* horizons are defined as $\mathcal{H}^+ \equiv \partial J^- [\mathcal{I}]$ and $\mathcal{H}^- \equiv \partial J^+ [\mathcal{I}]$. The causal surface can be defined as $C \equiv \mathcal{H}^+ \cap \mathcal{H}^-$ and the causal wedge is $\mathcal{W}_C \equiv J^+ [\mathcal{I}] \cap J^- [\mathcal{I}]$ [35].⁴ An important result in general relativity – which follows from NEC and cosmic censorship – is that future causal

³ k^a is not uniquely defined in places where σ is not C^1 . However, the expansion still has a definite sign, which can be computed via a limiting procedure.

⁴Note that although the term causal wedge is common in the literature, unlike the entanglement wedge it does not refer to a domain of dependence.

horizons satisfy an area law: the areas of their cross sections do not decrease as we move the cross section to the future [36, 37]. In particular, this means that any congruence of null generators on a future horizon has nonnegative expansion. By time-reversal symmetry, a “reverse” area law holds for past horizons.

- We define the terminated horizons $\mathcal{H}_C^+ \equiv \mathcal{H}^+ \cap J^-[\mathcal{I}]$ and $\mathcal{H}_C^- \equiv \mathcal{H}^- \cap J^-[\mathcal{I}]$. These are natural definition for us since we are interested in perturbations of $\mathcal{H}^+$ and $\mathcal{H}^-$ caused by causal boundary sources.

2 Which Wedge?

Three bulk regions are under consideration here: the *outermost extremal wedge*, the *causal wedge*, and the *simple wedge*. We will ultimately argue that the outermost extremal wedge is in fact the simple wedge, but in order to avoid subscribing to our own conclusions before we have demonstrated them, we introduce terminology that distinguishes between the two.

We will argue for the equivalence between the outermost extremal and simple wedges by showing that simple operations and sources (together with a finite number of time-folds) are sufficient to shift the causal wedge so that it comes arbitrarily close to coinciding with the outermost extremal wedge. While our primary results are for compact extremal surfaces, many of our intermediate results remain valid for boundary-anchored surfaces. In Sec. 7.1 we will discuss in more detail the extent to which our results apply to the latter case.

We have already defined the more familiar causal wedge in the introduction. Let us now give a precise definition of the outermost extremal wedge and the simple wedge.

Intuitively, the extremal wedge is defined as the analogue of the entanglement wedge for the outermost extremal surface – be it minimal or not. So before defining the outermost extremal wedge, we must prove that a unique outermost extremal surface exists in the first place:

Proposition 1. *If there exists more than one extremal surface homologous to a connected component of $\mathcal{I}$, then exactly one is outermost; i.e. there exists an extremal surface X contained in the outer wedge $W_{X'}$ of all other extremal surfaces X' homologous to $\mathcal{I}$.*

We will prove this proposition using a series of three lemmas:

Lemma 1. *Given two surfaces σ_1 and σ_2 homologous to $\mathcal{I}$, $W_{\sigma_1} \cap W_{\sigma_2}$ is a domain of dependence.*

Proof. Let Σ be a Cauchy slice containing σ_1 . We define a new surface $\Sigma' = \partial((J^+[\Sigma] \cup J^+[\sigma_2]) \cap \overline{I^-[\sigma_2]})$, where $\overline{I^-[\sigma_2]} \equiv M - I^-[\sigma_2]$. Since every inextensible timelike curve at some point in its past is outside $J^+[\Sigma] \cup J^+[\sigma_2]$, but eventually ends up in $J^+[\Sigma] \cap \overline{I^-[\sigma_2]}$, and since no timelike curve can exit $(J^+[\Sigma] \cup J^+[\sigma_2]) \cap \overline{I^-[\sigma_2]}$ after entering it, Σ' is Cauchy. Note that, despite appearances, the definition of Σ' is invariant under time-reversal symmetry.

We will now show that $W_{\sigma_1} \cap W_{\sigma_2}$ is the domain of dependence of $H = \Sigma' \cap W_{\sigma_1} \cap W_{\sigma_2}$. Any causal curve intersecting $W_{\sigma_1} \cap W_{\sigma_2}$ needs to intersect $\Sigma \cap W_{\sigma_1}$ either (i) in $W_{\sigma_2} - \partial W_{\sigma_2}$ or (ii) outside of it. In case (ii), the causal curve needs to leave $J^-[\sigma_2]$ in W_{σ_1} after intersecting Σ or enter $J^+[\sigma_1]$ in W_{σ_1} before intersecting Σ . Therefore, in both cases i and ii we conclude that the causal curve intersects H . $\square$

Let V now be a domain of dependence that intersects $\mathcal{J}$. Then there must exist an “edge” surface σ homologous to $\mathcal{J}$ such that $V = W_\sigma$. More precisely, σ can be defined as the set of points $p \in \partial V$ such that in any small neighborhood of p any inextensible timelike curve crossing p only intersects V at p .

Lemma 2. *Let $V = W_{\sigma_1} \cap W_{\sigma_2}$ and let σ be the edge of V as defined above. Then, $\sigma \subset \sigma_1 \cup \sigma_2 \cup (\partial^+ W_{\sigma_1} \cap \partial^- W_{\sigma_2}) \cup (\partial^- W_{\sigma_1} \cap \partial^+ W_{\sigma_2})$.*

Proof. Clearly, the edge is contained in $\partial V \subset \partial W_{\sigma_1} \cup \partial W_{\sigma_2}$. Say some point p in the edge of V were in $\partial^+ W_{\sigma_1} - \sigma_1$. Then, every timelike curve crossing p leaves W_{σ_1} (and therefore V) in $I^+(p)$ and enters W_{σ_1} in $I^-(p)$. Furthermore, for p to be in the edge of V it needs to be in W_{σ_2} , but $I^-(p)$ must not intersect W_{σ_2} . Therefore, p needs to be in $\partial^- W_{\sigma_2}$. Together with the time reverse of this argument and also switching σ_1 and σ_2 , we conclude that the edge of V must be contained in $\sigma_1 \cup \sigma_2 \cup (\partial^+ W_{\sigma_1} \cap \partial^- W_{\sigma_2}) \cup (\partial^- W_{\sigma_1} \cap \partial^+ W_{\sigma_2})$. $\square$

Lastly, we will state a lemma from Sec. 2.2 of [38], without providing the proof (see also Appendix B of [6] for a similar discussion). The lemma assumes the existence of a stable maximin surface [31] in any domain of dependence.

Lemma 3. *If σ is a surface homologous to $\mathcal{J}$ satisfying $\theta_{(k)} \leq 0$ and $\theta_{(\ell)} \geq 0$, then there exists an extremal surface $Y \subset W_\sigma$ homologous to $\mathcal{J}$.*

Let us provide some intuition for Lemma 3. The restricted-maximin prescription returns a surface in W_σ which is homologous to $\mathcal{J}$ and is minimal on some Cauchy slice of W_σ . If σ satisfies $\theta_{(k)} < 0$ and $\theta_{(\ell)} > 0$, then the maximin surface cannot intersect σ , since its area would get smaller under deformations away from such intersections.

In [38], it was further shown that the max property of the maximin surface prohibits intersections with $\partial W_\sigma - \sigma$. This shows that the maximin surface is in the interior of W_σ and thus extremal. Furthermore, it was argued that even when the inequalities are not strict, there still exists an extremal surface homologous to $\mathcal{J}$ in W_σ even though in the surface might lie on ∂W_σ .

We are now ready to prove Proposition 1:

Proof. In this proof any “surface” will mean a surface that is homologous to $\mathcal{J}$. Let an extremal surface X be called *exposed* if there does not exist any extremal surface $Y \neq X$ such that $Y \subseteq W_X$. Let us first argue that there must always exist at least one exposed surface. Define a partial ordering on extremal surfaces by declaring $X \geq Y$ if and only if $W_X \subseteq W_Y$. Note that exposed surfaces would correspond to maximal elements with respect to this partial order, while an outermost extremal surface would be a greatest element. Upper bounds exist for any chain because monotonicity and boundedness (there are no extremal surfaces near asymptotic infinity) ensure that any sequence $\{X_n\}$ of extremal surfaces with $W_{X_{n+1}} \subseteq W_{X_n}$ converges to an extremal surface X_∞ with $W_{X_\infty} \subseteq W_{X_n}$ for any finite n . Hence, by Zorn’s lemma, at least one maximal element, i.e. exposed surface, exists.

Now suppose, by way of contradiction, that there exists an exposed surface X_1 that is not outermost, i.e. there exists some other extremal surface X_2 such that $W_{X_1} \not\subseteq W_{X_2}$. Let $V = W_{X_1} \cap W_{X_2}$ (by definition then $V \subset W_{X_1}$). By Lemma 2 and 3, $V = W_\sigma$ for some surface $\sigma \subseteq X_1 \cup X_2 \cup (\partial^+ W_{X_1} \cap \partial^- W_{X_2}) \cup (\partial^- W_{X_1} \cap \partial^+ W_{X_2})$. Let us consider each component of the set to which σ belongs separately. The first two are subsets of extremal surfaces and hence have zero expansion. By focusing, the null hypersurface $\partial^+ W_{X_1}$ and $\partial^+ W_{X_2}$ are non-expanding towards the future, while the null hypersurface $\partial^- W_{X_1}$ and $\partial^- W_{X_2}$ are non-expanding towards the past. Therefore, σ satisfies $\theta_{(k)} \leq 0$ and $\theta_{(\ell)} \geq 0$ ⁵, and hence by Lemma 3, there exists an extremal surface Y contained in $W_\sigma = W_{X_1} \cap W_{X_2}$. The existence of this surface means that X_1 is not exposed, giving our desired contradiction. □

As an aside, note that extending the last part of this argument to quantum extremal surfaces, under assumption of the quantum focusing conjecture requires a small amount of extra work because the definition of quantum expansion is nonlocal. However strong

⁵This σ may have kinks in which case the expansions are technically not well-defined. However, we expect that all of our results in this section generalizes easily to such surfaces using standard geometric techniques. Intuitively, by slightly smoothing out the kinks we can get a new infinitesimally nearby surface with very large expansions of the desired signs around the kinks.

subadditivity is enough to ensure that the quantum expansion of σ satisfies the desired inequalities.

It is now easy to define the outermost extremal wedge:

Definition 1. *Let X be the outermost extremal surface for a connected component of $\mathcal{I}$. The outermost extremal wedge W_X is the outer wedge of X .*

Next we would like to define the simple wedge. Conceptually, the simple wedge is the largest bulk region that can be reconstructed from the near boundary state of the bulk fields using exclusively the bulk equations of motion. Consider some boundary state ρ whose bulk dual (M, g) we would like to reconstruct and evolve it to the far past and future with some Hamiltonian. In the classical regime, following Eq. (1.1), HKLL prescribes the values of the bulk fields in $\mathcal{W}_C$ from the set of one-point functions of their corresponding local boundary operators on $\mathcal{I}$. In fact, the bulk equations of motions are sufficient to reconstruct the maximal Cauchy development of $\mathcal{W}_C$ – which we can denote by W_C .

Cosmic censorship in general [39, 40] and causal wedge inclusion in particular [31, 41, 42] guarantees that the causal wedge contains no extremal surfaces in its interior: $W_C \subseteq W_X$. It is therefore impossible to reconstruct the region behind X causally. HKLL alone, however, appears to *prima facie* fail at an earlier stage: the non-standard Cauchy evolution appears to stop short of recovering the gap region between W_C and W_X , which is generically non-empty.

What if we evolve ρ using a different Hamiltonian? Consider turning on a set of CFT operators at various times during the evolution. This would “extract” a new set of one-point functions from ρ and therefore has the potential to expand the reconstructible region. In keeping with our philosophy of the simple wedge, we must restrict to sources that have a (semi-)classical bulk dual. Therefore, following [24, 25], we refer to boundary sources as simple if the bulk fields that they produce propagate causally into the bulk from the boundary insertion point – and restrict to such sources henceforth. The change in time evolution when such simple sources are applied within some boundary time interval $[t_i, t_f]$ is given by the following time-ordered operator:

$$E = \mathcal{T}exp \left[-i \int_{t_i}^{t_f} dt' J(t') \mathcal{O}_J(t') \right], \quad (2.1)$$

where $J(t)$ is a simple source and $\mathcal{O}_J$ is its corresponding simple operator. Note that $\mathcal{O}_J$ might involve spatial integrals of local boundary operators $\mathcal{O}(t', x')$. An example of a simple operator is a spatial integral of a single-trace operator of the boundary gauge theory.

Adding E to the evolution, say in a future-directed timefold, changes the spacetime from M to some M' . By causality, $M - J^+[t_i]^6 \subset M'$. In particular, the perturbation to the spacetime is localized away from the past event horizon $\mathcal{H}^-$. However, sources like Eq. (2.1) will typically change where the new future event horizon $\mathcal{H}^+$ intersects $\mathcal{H}^-$. In particular, suppose that we find simple sources that “expand” the causal wedge, i.e. place the new causal surface C' in the future of C on $\mathcal{H}^-$. Said in the CFT language, the new set of one-point data reconstructs a $\mathcal{W}'_C$ that contains $\mathcal{H}^-_C$. Furthermore, knowing the bulk equations of motion and the original Hamiltonian, we can reconstruct the Cauchy development of $\mathcal{H}^-_C$, a wedge in the *original* spacetime that contains W_C as a proper subset.

It is natural to define the simple wedge according to the maximal success of this procedure:

Definition 2. *The simple wedge is the maximal bulk region that can be reconstructed from simple operators acting on the dual CFT state, with the inclusion of simple sources and timefolds.*

Although we have defined the simple wedge in the context of classical field theory in the bulk, it is important to note that HKLL can reconstruct the quantum state of the bulk propagating in the causal wedge at each step. For bulk fields in the $1/N$ expansion, Eq. (1.1) provides the dictionary between local bulk operators and simple CFT operators, realizing this reconstruction in the Heisenberg picture.

Finally, we close this section by relating the causal and outermost extremal wedges. Intuitively, the causal surface should coincide with the extremal surface if and only if there is no focusing whatsoever on the horizons. However, because the extremal wedge is defined as a domain of dependence and the causal wedge is defined in terms of causal horizons, it does not immediately follow that the two must coincide in the absence of focusing. To reassure the reader, we prove the following lemma:

Lemma 4. *Let X be the outermost extremal surface homologous to one or more connected components of $\mathcal{I}$. $W_X = \mathcal{W}_C$ if and only if X is a bifurcation of stationary horizons (and thus $X = C$).*

Proof. If $\mathcal{W}_C = W_X$, then $\partial\mathcal{W}_C = \partial W_X$. If both wedges had been defined in terms of domains of dependence, it would immediately follow that $C = X$. However, since $\mathcal{W}_C$ is not defined as a domain of dependence, we have to work a little harder. The component of $\partial\mathcal{W}_C$ which is spacelike separated to every point in $\text{Int}[\mathcal{W}_C]$ is identical to the component of W_X which is spacelike-separated to every point in $\text{Int}[W_X]$. By

⁶Here t_i stands for a timeslice of $\mathcal{I}$ geometrically.

definition of the causal wedge, the former is C . Since W_X is generated by the domain of dependence of a hypersurface H whose boundary in M is X , X is exactly the set of points that are spacelike separated from every point in $\text{Int}[W_X]$. This immediately shows that $C = X$. Because $\theta_{(n)}[X] = 0$ for all n^a in the normal bundle of X , we find that $\theta_{(\ell)}[C] = 0 = \theta_{(k)}[C]$. By the NCC, a future (past) causal horizon can only have vanishing expansion on a slice if it has vanishing expansion everywhere to the future (past) of that slice. So $\partial\mathcal{W}_C$ and subsequently ∂W_X is generated by stationary horizons, and X is a bifurcate stationary horizon.

If X is a bifurcate stationary horizon, then results of [33, 34] immediately imply that ∂W_X is the union of two truncated stationary horizons $\mathcal{H}^\pm$. Since $\mathcal{H}^+ \cap \partial M = \partial D^+$, $\mathcal{H}^+$ is a past-directed null congruence fired from i^+ . By the theorems of [33], $\partial J^-[\mathcal{I}] = \mathcal{H}^+$ up to geodesic intersections. Since $\mathcal{H}^+$ is stationary, it has no intersections, so $\partial J^-[\mathcal{I}] = \mathcal{H}^+$. Similarly $\partial J^+[\mathcal{I}] = \mathcal{H}^-$. Thus $\partial W_X = \partial\mathcal{W}_C$ and $W_X = \mathcal{W}_C$ by the homology constraint. $\square$

Note that this result remains valid for quantum extremal surfaces assuming the quantum focusing conjecture and a suitable generalization of AdS hyperbolicity.

3 Two Dimensions

Let us illustrate our iterative procedure for removing matter falling across the past and future causal horizons in a simple toy model of JT gravity minimally coupled to a massless scalar field φ (with no direct coupling between the dilaton Φ and φ). The absence of propagating degrees of freedom of the gravitational field as well as the factorization of the bulk matter into left- and right-movers are simplifications that naturally do not generalize to higher dimensions; nevertheless the procedure itself is well-illustrated in this setting, which we include for pedagogical reasons. The additional complications introduced in higher dimensions are resolved in subsequent sections.

Due to focusing resulting from the scalar field φ , the bifurcation surface will generically not be extremal, i.e.

$$\partial_n \Phi|_{C_1} \neq 0 \tag{3.1}$$

where n^a is some vector normal to the causal surface C_1 (in particular, for null n^a future-outwards directed, this would be positive; similarly for a time-reverse). As a consequence of the highly limited number of degrees of freedom in the problem, the extremality failure can only be a result of focusing: the future causal horizon will experience focusing due to the φ left-moving modes and the past causal horizon will experience focusing due to the φ right-moving modes. Our procedure instructs us to




Figure 4. An illustration of the approach to the limit point C_{lim} , where C_A and C_B are infinitesimally close to the limiting point.

first remove the source of focusing of the future horizon by modifying the boundary conditions of φ , which we can easily do by implementing absorbing boundary conditions for the right movers in order to remove all the left-moving modes. This removes focusing from the future causal horizon, which pushes the horizon deeper into the bulk. As a consequence, the new causal surface (which is now marginally trapped) – let us call it C_2 – is pushed further along the past event horizon, which remains unmodified by this procedure. This first step is illustrated in the second panel of Fig. 2.

To turn off focusing of the past horizon, we evolve backwards in time, imposing boundary conditions that remove the right-movers. The past event horizon moves backwards, and the new causal surface C_3 (which is now marginally anti-trapped) is displaced from C_2 along the future event horizon of C_2 . However C_3 is not necessarily extremal since shifting the past causal horizon reveals a part of the spacetime that was previously not included in the causal wedge: in particular, new left-moving modes can now appear in the causal wedge. This is illustrated in the third panel of Fig. 2. This piecewise-null zigzag procedure thus shifts the causal surface deeper into the bulk; we may simply repeat the zigzag iteratively.

In classical gravity, the focusing theorem and cosmic censorship (or strong asymptotic predictability) together guarantee that no extremal surface is ever in causal contact with $\mathcal{I}$: so the zigzag procedure can never modify an extremal surface nor move the causal surface deeper than any extremal surface. Thus the outermost extremal surface is an upper bound on the success of the procedure. Our goal, of course, is to show that this upper bound is in fact attained.

Because the success of this procedure is bounded by the outermost extremal surface and because furthermore the procedure moves the surface monotonically inwards,

a limiting causal surface C_{lim} exists, and the corresponding causal wedge does not intersect any extremal surface. We will now argue that C_{lim} is in fact extremal. Let C_A be a causal surface obtained via iterative zigzags which is infinitesimally close to C_{lim} ; without loss of generality we may take C_A to be in the marginally trapped portion of the zigzag (i.e. the left-movers had just been removed). Let v, u be the affine parameters along the future and past event horizons, respectively, in the spacetime in which C_A is the causal surface. By construction, C_A has no expansion along the future event horizon:

$$\frac{\partial}{\partial v}\Phi|_{C_A} = 0, \quad (3.2)$$

and because by assumption it is not identical to C_{lim} ,

$$\frac{\partial}{\partial u}\Phi|_{C_A} < 0. \quad (3.3)$$

Let C_B be the causal surface obtained by removing the right-movers from the spacetime where C_A is the causal surface. By construction

$$\frac{\partial}{\partial u}\Phi|_{C_B} = 0. \quad (3.4)$$

By the zigzag procedure, C_A and C_B are null-separated along the “old” future event horizon: i.e. along the null congruence that is the future event horizon in the spacetime in which C_A is the causal surface. Let δv be the amount of affine parameter separating C_A and C_B . See Fig. 4 for an illustration. Since the points C_A and C_B must be infinitesimally close to one another (since both are infinitesimally near C_{lim}), the spacetime metric in that neighborhood may be approximated as locally flat instead of AdS_2 ; using u and v as coordinates:

$$ds^2 = -2dudv. \quad (3.5)$$

In these coordinates, it is trivial to relate δv to the change in $\partial_u\Phi$ along v . In particular, we may bound it from below:

$$\delta v \geq \frac{\partial_u\Phi}{\partial_v\partial_u\Phi|_{\text{max}}}, \quad (3.6)$$

where $\partial_v\partial_u\Phi|_{\text{max}}$ is the maximum value of $\partial_v\partial_u\Phi$ on the δv interval. Similarly defining δu as the null separation between C_B and the next causal surface after again removing left movers, we obtain an analogous bound:

$$\delta u \geq \frac{\partial_v\Phi}{\partial_u\partial_v\Phi|_{\text{max}}}. \quad (3.7)$$

Under the assumption that $\partial_u\partial_v\Phi$ is bounded from above in this neighborhood (which we generically expect to be true), δv and δu approach zero no slower than $\partial_u\Phi$ and $\partial_v\Phi$

approach zero: thus C_{lim} must be extremal. Because focusing arguments ensure that the causal wedge is always contained in the outermost extremal wedge, it must be the outermost extremal surface.

4 The Perturbation

Our task in higher dimensional gravity is now clear: we must find a perturbation that removes focusing from the causal horizons (without violating the null energy condition anywhere in the perturbed spacetime), thus shifting the causal wedge deeper in. What kind of perturbation δg would move the causal surface towards rather than away from the appetizer? On a heuristic level, we are looking to open up the bulk lightcones so that more of the bulk is in causal contact with the asymptotic boundary. In searching for such a perturbation, we may build on the intuition of the boundary causality condition [43], which states that the inequality

$$\int \delta g_{ab} k^a k^b \geq 0 \quad (4.1)$$

(where the integral is over a complete null geodesic with generator k^a) is equivalent to demanding that perturbations δg of pure AdS source focusing (as opposed to defocusing). Here we are looking to do the opposite: we are looking to undo focusing, so it makes sense to look for a perturbation that satisfies an opposite inequality, with $\delta g_{kk} < 0$ everywhere on $\mathcal{H}_C^+$. It is a priori not clear that it is possible to find a perturbation that simultaneously satisfies this inequality and also results in a spacetime that solves the Einstein equation. To prove this point, we must show that the requisite δg solves the characteristic constraint equations on the event horizon.

In this section, we will prove that as long as the causal surface is not marginally outer trapped – i.e. as long as $\theta_{(k)}[C] \neq 0$, it is possible to find exactly such a perturbation that (1) satisfies the characteristic constraint equations on the causal horizon and (2) shifts the causal surface deeper into the bulk. The procedure is roughly as follows: we prescribe a δg deformation on $\mathcal{H}^+$; some elements of this δg resemble the “left stretch” construction [44, 45] involving a rescaling of the generators of certain achronal null hypersurfaces – intuitively, this dilutes the infalling content on $\mathcal{H}^+$ and in turn reduces focusing. We then demonstrate that the gravitational constraints on δg along with boundary conditions fix the requisite components of δg in such a way that the perturbed spacetime has a larger causal wedge. In Sec. 5, we will argue how repeating these perturbations pushes in the causal surface up to an apparent horizon in a given timefold, and to the appetizer using several timefolds.




Figure 5. On the left: three-dimensional (left) illustration of the past ($v = 0$) and future ($u = 0$) event horizons, with a slicing of the future event horizon given by the intersection of $u = 0$ with past causal horizons originating from complete slices of the $\mathcal{I}$. On the right: a conformal diagram illustrating the same.

We will call the generators of the future and past event horizons k^a and ℓ^a respectively.⁷ We will extend ℓ^a to the entire spacetime by picking a smooth Cauchy foliation $\{\mathcal{C}_\alpha\}_\alpha$ of $\mathcal{I}$ and defining ℓ^a to be the bulk generators of $\partial I^+[\mathcal{C}_\alpha]$. This defines a null foliation of $J^+[\mathcal{I}]$ by past causal horizons; The past event horizon, which C lies on, is a member of this foliation.

We adopt the coordinate and gauge choices of [25]: first, we introduce double null coordinates (u, v) ⁸, where

$$k^a = \left(\frac{\partial}{\partial v} \right)^a \quad ; \quad \ell^a = \left(\frac{\partial}{\partial u} \right)^a. \quad (4.2)$$

In these coordinates, the causal surface is at $u = v = 0$, the future event horizon $\mathcal{H}^+$ is at $u = 0$ and the past event horizon $\mathcal{H}^-$ is at $v = 0$. See Fig. 5. We can further fix the gauge in a neighborhood of $u = 0$ so that the metric there takes the form:

$$ds^2 = -2dudv + g_{vv}dv^2 + 2g_{vi}dvdy^i + g_{ij}dy^i dy^j \quad (4.3)$$

where i, j denote the transverse direction on $\mathcal{H}^+$. At $u = 0$ exactly, we of course require that $k^a = (\partial/\partial v)^a$ is null, so that:

$$g_{vv}|_{u=0} = 0, \quad (4.4)$$

⁷These will agree with k^a and ℓ^a on C whenever C is C^1 .

⁸Calling these double-null coordinates is a slight abuse of notation, as we will only require $g_{vv} = 0$ at the horizon $u = 0$.

and we may further fix the gauge:

$$g_{vi}|_{u=0} = 0, \quad (4.5)$$

but we cannot require these components to be *identically* zero in a neighborhood of $u = 0$, i.e. the derivatives may not vanish. For instance, the inaffinity of k^a – given by $\kappa_{(v)} = \frac{1}{2}\partial_u g_{vv}|_{u=0}$ – cannot be set to zero in general since we have independently fixed the ℓ^a vector field, and orthogonality to ℓ^a defines the constant (or affine) v slices on $\mathcal{H}^+$.⁹

The extrinsic curvature tensors of constant- v slices are simple in this gauge

$$B_{(v)ij} = \frac{1}{2}\partial_v g_{ij}|_{u=0}, \quad (4.6a)$$

$$B_{(u)ij} = \frac{1}{2}\partial_u g_{ij}|_{u=0}, \quad (4.6b)$$

$$\chi_i = \frac{1}{2}\partial_u g_{vi}|_{u=0}. \quad (4.6c)$$

where $B_{(v)ij}$ and $B_{(u)ij}$ denote the null extrinsic curvatures of constant v slices and χ_i denotes the their twist. Since $\mathcal{H}^+$ is achronal, we can specify a new solution to the Einstein equation via a perturbative deformation of the metric on it, so long as the null constraint equations are satisfied. In particular, we consider the following perturbation on $\mathcal{H}_C^+(u = 0, v \geq 0)$:

$$ds^2 = -2dudv + (g_{vv} + \delta g_{vv})dv^2 + 2(g_{vi} + \delta g_{vi})dvdy^i + (g_{ij} + \delta g_{ij})dy^i dy^j. \quad (4.7)$$

The perturbation components δg_{vv} , δg_{vi} , δg_{ij} and their first u derivatives are initial data that we can freely specify on the characteristic hypersurface $u = 0$ so long as this data satisfies the null constraint equations¹⁰:

$$\delta G_{vv} + \Lambda \delta g_{vv} = 8\pi G \delta T_{vv} \quad (4.8a)$$

$$\delta G_{uv} = 8\pi G \delta T_{uv} \quad (4.8b)$$

$$\delta G_{vi} + \Lambda \delta g_{vi} = 8\pi G \delta T_{vi} \quad (4.8c)$$

$$\delta \hat{G}_{ij} + \Lambda \delta \hat{g}_{ij} = 8\pi G \delta \hat{T}_{ij} \quad (4.8d)$$

where δG_{ab} denotes the linearized perturbation of the Einstein tensor, Λ is the cosmological constant, and the hat in the last equation denotes the traceless part. Note

⁹The covariant definition of the inaffinity of k^a is $k^a \nabla_a k^b = \kappa_{(k)} k^b$.

¹⁰The null constraint equations are the equations of motion containing at most one u derivative of the metric perturbations.

that the corresponding deformations of the stress energy tensor must be sourced by a perturbative modification to the matter fields that itself satisfies the fields' equations of motion on the background geometry.

Thus we will need to prescribe δg as well as a matter source for it. The latter is accomplished by a perturbative “stretch”. A nonperturbative stretch is an exponential rescaling [44, 45]:

$$g'_{ij}(u=0, v, y) = g_{ij}(u=0, ve^{-s}, y) \quad (4.9a)$$

$$T'_{vv}(u=0, v, y) = e^{-2s} T_{vv}(u=0, ve^{-s}, y) \quad (4.9b)$$

$$\kappa'_{(v)}(u=0, v, y) = e^{-s} \kappa_{(v)}(u=0, ve^{-s}, y) \quad (4.9c)$$

where prime denotes the transformed quantities. Our matter source will be obtained in the perturbative limit of this, by setting $e^{-s} = 1 - \epsilon$, where $\epsilon \sim \mathcal{O}(\delta g)$ is the parameter controlling the expansion. Our choice for the full perturbation on $v \geq 0$, metric and matter, is then:

$$\delta g_{vi}(u=0, v, y) = 0 \quad (4.10a)$$

$$\partial_u \delta g_{vv}(u, v, y)|_{u=0} = 2(1 - \epsilon) \kappa_{(v)}(0, v(1 - \epsilon), y) - 2\kappa_{(v)}(u=0, v, y) \quad (4.10b)$$

$$\delta g_{ij}(u=0, v, y) = g_{ij}(u=0, v(1 - \epsilon), y) - g_{ij}(u=0, v, y) \quad (4.10c)$$

$$\delta T_{vv}(u=0, v, y) = (1 - \epsilon)^2 T_{vv}(u=0, v(1 - \epsilon), y) - T_{vv}(u=0, v(1 - \epsilon), y) \quad (4.10d)$$

where in this linearized analysis we will only need to keep track of first order terms in ϵ and δg (i.e. we will drop all terms of order $\epsilon \delta g$). Note that δg_{vv} , $\partial_u \delta g_{ij}$, and $\partial_u \delta g_{vi}$ are allowed to be non-zero. We will see that their values are constrained subject to the above restrictions.

Before we move forward with the analysis, we need to ask whether we can always obtain the stress tensor profile in Eq. (4.10d). This question is difficult to answer in broad generality. Therefore, from now on we restrict our matter sector to consist of a minimally coupled complex scalar field theory ϕ (with an arbitrary potential) coupled to some Maxwell field (or consider either separately), with Lagrangian density

$$\mathcal{L}_{\text{matter}} = -\frac{1}{4} g^{ac} g^{bd} F_{ab} F_{cd} - g^{ab} \bar{\nabla}_a \phi \bar{\nabla}_b \phi^* - V(|\phi|^2) \quad (4.11)$$

where F_{ab} is the field strength and $\bar{\nabla}$ denotes the covariant derivative with respect to the vector potential A^a . Then

$$T_{vv} = 2\bar{\nabla}_v \phi \bar{\nabla}_v \phi^* + F_{iv} F_v^i, \quad (4.12)$$

Since both ϕ and A_v are free initial data in the characteristic problem, we can simply generate the desired transformation by setting $\phi'(v, y) = \phi(ve^{-s}, y)$ and $A'_v = e^{-s} A_v(ve^{-s}, y)$, and $A'_i = A_i(ve^{-s}, y)$.

We now proceed to prove that our choice of perturbation solves the null constraint equations with $\delta g_{vv} < 0$ everywhere on $\mathcal{H}_C^+$. Because the unperturbed spacetime satisfies the NEC and no new matter terms are introduced by the perturbation, the perturbed spacetime will likewise satisfy the NEC. For pedagogical clarity, we will focus on the more illuminating vv -constraint here and relegate the remaining constraint equations to Appendix A. Our analysis here will be twofold: we will first analyze the constraint (4.8a) in the absence of δg_{vv} , separately compute the contribution of δg_{vv} , and sum the two together; this is possible so long as we work in the linearized regime.

By the Raychaudhuri equation, R_{vv} depends only on the geometry of the $u = 0$ hypersurface:

$$R_{vv} = -\partial_v \theta_{(v)} - B_{(v)ij} B_{(v)}^{ij} - \kappa_{(v)} \theta_{(v)}. \quad (4.13)$$

Therefore in the absence of δg_{vv} (i.e. implementing only the stretch):

$$\begin{aligned} R'_{vv}(u=0, v, y) &= e^{-2s} R_{vv}(u=0, v e^{-s}, y) - (1 - e^{-s}) \theta_{(v)}(u=0, v=0, y) \delta(v) \\ &= T'_{vv}(u=0, v, y) - (1 - e^{-s}) \theta_{(v)}(u=0, v=0, y) \delta(v) \end{aligned} \quad (4.14)$$

where the delta function term results from the discontinuity in $\theta'_{(v)}$ across $v = 0$. So $R'_{vv} - T'_{vv} = -(1 - e^{-s}) \theta_{(v)}[C] \delta(v)$. We now take the same perturbative limit of this transformation and re-introduce δg_{vv} :

$$\begin{aligned} & -\frac{1}{2} \theta_{(u)} \partial_v \delta g_{vv} - \frac{1}{2} \nabla_\perp^2 \delta g_{vv} + \chi^i \partial_i \delta g_{vv} + \\ & \left(\nabla_\perp \cdot \chi - \partial_v \theta_{(u)} - B_{(v)ij} B_{(u)}^{ij} + 8\pi G(-T_{uv} - \mathcal{L}_{\text{matter}} + F_{uv}^2) \right) \delta g_{vv} \\ & - \epsilon \theta_{(v)}[C] \delta(v) = 0 \end{aligned} \quad (4.15)$$

where all of the quantities multiplying δg_{vv} and its derivatives are background quantities. We offer an alternative derivation of Eq. (4.15) in Appendix B by implementing the “stretch” using an inaffinity shock [44, 45], which directly induces the delta function term in Eq. (4.15).

Since by construction, we are only perturbing the data on $\mathcal{H}_C^+$, $\delta g_{vv}(u=0, v=0^-, y) = 0$, so the delta function term in Eq. (4.15) enforces a jump in δg_{vv} :

$$\delta g_{vv}(u=0, v=0^+, y) = 2\epsilon \frac{\theta_{(v)}[C]}{\theta_{(u)}[C]} \leq 0 \quad (4.16)$$

where the sign comes from the fact that $\theta_{(v)}[C] \geq 0$ and $\theta_{(u)}[C] < 0$.¹¹ Note that by assumption $\theta_{(v)}[C] > 0$, and so $\delta g_{vv} < 0$, at least in a subset of C . This implies that

¹¹In a non-generic spacetime, it is possible that $\theta_{(u)}[C] = 0$, which raises potential concerns about the divergence in Eq. (4.16). We can avoid this issue by shifting the earliest location of the perturbation

the curve generated by $(\partial_v)^a$ is nowhere spacelike and at least timelike on a subset of C . In order to open up the lightcone and move the causal surface deeper into the bulk, it would be sufficient if $\delta g_{vv} \leq 0$ *everywhere* on $\mathcal{H}_C^+$, not just at C .

We will now demonstrate that if $\delta g_{vv}(u = 0, v = 0^+, y) \leq 0$, then $\delta g_{vv}(u = 0, v, y) \leq 0$ for all $v > 0$, by analyzing the constraint that δg_{vv} satisfies on $v > 0$:

$$-\frac{1}{2}\theta_{(u)}\partial_v\delta g_{vv} - \frac{1}{2}\nabla_\perp^2\delta g_{vv} + \chi^i\partial_i\delta g_{vv} + \left(\nabla_\perp\cdot\chi - \partial_v\theta_{(u)} - B_{(v)ij}B_{(u)}^{ij} + 8\pi G(-T_{uv} - \mathcal{L}_{\text{matter}} + F_{uv}^2)\right)\delta g_{vv} = 0 \quad (4.17)$$

which we may view as an “evolution” equation for δg_{vv} on $u = 0$ from which we can derive δg_{vv} on $\mathcal{H}_C^+$ from its value at C .

It is not too difficult to see why $\delta g_{vv}(u = 0, v, y) \leq 0$ for all $v > 0$ starting from $\delta g_{vv}(u = 0, v = 0^+, y) \leq 0$. Suppose $\delta g_{vv} > 0$ at some value of v . Then, assuming that all quantities in (4.17) are continuous, there must exist a “last” constant- v slice σ on which $\delta g_{vv} \leq 0$ everywhere. By continuity, there exists a point $p \in \sigma$ where $\delta g_{vv}|_p = 0$ (and then immediately becomes positive for larger v). By construction, we must have $\partial_i\delta g_{vv}|_p = 0$ and $\nabla_\perp^2\delta g_{vv}|_p \leq 0$. But by (4.17) this implies that $\partial_v\delta g_{vv}|_p \leq 0$, and so δg_{vv} cannot become positive.

This reasoning may seem a bit fast, but it can be made more rigorous (and free of simplifying assumptions) using standard techniques. The operator $\mathcal{L}$ acting on δg_{vv} in (4.17) is parabolic whenever $\theta_{(u)} < 0$; it thus satisfies the weak comparison principle for parabolic operators, which states that if f and h are functions satisfying $\mathcal{L}f \leq 0$ and $\mathcal{L}h \geq 0$ everywhere in the interior of the parabolic domain, and $f \leq h$ on the boundary of the parabolic domain, then $f \leq h$ everywhere in the parabolic domain. Setting $f = \delta g_{vv}$ and $h \equiv 0$, we immediately find that $\mathcal{L}f = 0$ and $\mathcal{L}h = 0$, so the weak comparison principle yields the desired conclusion: $\delta g_{vv} \leq 0$ everywhere on $\mathcal{H}_C^+$. Technically, the weak comparison principle is usually stated for domains in $\mathbb{R}^n$, fortunately, it follows as a fairly direct consequence of the maximum principle for elliptic operators, which does hold for more general manifolds [46]. The functions f and h need only be of Sobolev type $W_0^{1,2}$; that is, only their local weak derivatives are required to exist [47], which is sufficient for our purposes. In fact, a version of the maximum

to some arbitrarily small $v = \delta > 0$ instead of $v = 0$. This new cut lies on a past causal horizon that reaches $\mathcal{I}$ and therefore cannot be stationary. In fact, originating the perturbation at $v = \delta > 0$ makes sense physically since by sending t_i to arbitrarily small values in the source Eq. (2.1) we can affect the region arbitrarily close to C , but not C itself. With this subtlety in mind, we pick $v = 0$ as the origin of the perturbation in the main text because we can get arbitrarily close to C .

principle for elliptical operators on “rough” null hypersurfaces (including caustics and non-local intersections specifically on event horizons) was proved in [48].

To make sure that our δg_{vv} solution exists, we need to also satisfy the other constraints (4.8b), (4.8c), and (4.8d). This is easy to do because they are “evolution” equations for $\partial_u \delta g_{ij}$ and $\partial_u \delta g_{vi}$ on $\mathcal{H}_C^+$ which we can solve no matter what δg_{vv} is. We relegate this discussion to appendix A.

Let us now discuss possible subtleties in our construction due to caustics. Since caustic lines will generically be a measure zero subset of $\mathcal{H}^+$ [37, 49, 50], we believe that they do not pose a fundamental obstruction to our procedure. Caustic lines can intersect C , at which point C will generically be kinked. At the location of the kink, a chunk of transverse directions, associated to the generators that emanate from the caustic line, needs to get inserted in the transverse domain on which we place our boundary data for Eq. (4.17). However, so long as this data satisfies $\delta g_{vv} \leq 0$, Eq. (4.17) still guarantee $\delta g_{vv} \leq 0$ everywhere on $\mathcal{H}_C^+$. In fact, since these new generators do not extend to the past of the caustics by definition, we expect to have even more freedom in specifying this boundary data because we do not have to worry about how this boundary is glued to some past hypersurface.

Let us offer an alternative argument to further ameliorate caustic-related worries. As $\mathcal{H}_C^+$ settles down to Kerr-Newman, there exists an earliest cross section μ_{earliest} lying on a past horizon with no caustics in its future. By setting μ_{earliest} as the origin of our perturbation (the new $v = 0$), we can avoid caustics altogether. Furthermore, each perturbation should make the portion of the horizon to the future of μ_{earliest} more stationary, pushing the new μ_{earliest} further to the past eventually approaching $\mathcal{H}^-$.

Lastly, it is important to show that the perturbation has not shifted the $u = 0$ surface to the point where it is no longer close to the event horizon – especially in the asymptotic region $v \rightarrow \infty$. This is simplest to do if we assume that the *background* horizon settles down to a stationary spherically symmetric configuration at some finite affine parameter, though we expect proper falloffs to hold more generally. The equation for δg_{vv} then simplifies to:

$$-\frac{1}{2}\theta_{(u)}\partial_v\delta g_{vv}-\frac{1}{2}\nabla_{\perp}^2\delta g_{vv}-\partial_v\theta_{(u)}\delta g_{vv}=0 \quad (4.18)$$

On $\mathcal{H}^+$, we have $\theta_{(u)} \sim v$ asymptotically. We can then solve for the asymptotic behavior of δg_{vv} from Eq. (4.18):

$$\delta g_{vv} \sim v^{-2} \quad (4.19)$$

Therefore $u = 0$ asymptotes to a stationary null hypersurface after the perturbation, so it naturally lines up with the new causal horizon in the $v \rightarrow +\infty$ limit. We find

that our proposed perturbation results in a spacetime that solves the Einstein equation and in which the causal horizon is pushed deeper into the bulk unless $\theta_{(k)}[C] = 0$.

5 Zigzag

We will now use the above perturbation to show that the causal surface C can be moved arbitrarily close to the outermost extremal surface X , the appetizer of the lunch, using simple sources only. This requires us to show that (1) the perturbation analyzed in the previous section can be engineered from simple sources on the boundary, (2) the perturbation can be iteratively repeated both in the past and future, resulting in the approach of C to X , without incurring high complexity, and (3) that this procedure does not change the geometry of the lunch, nor does the causal surface breach the lunch region in the process. We will begin our discussion by assuming for simplicity¹² that the causal surface and the appetizer have the same topology; topological differences between the two surfaces are discussed at the end.

To show (1), we simply employ our assumption of the validity of HKLL, discussed in Sec. 1, and evolve the data on $\mathcal{H}_C^+ \cup \mathcal{H}_C^-$ “sideways” towards $\mathcal{I}$ to find appropriate boundary conditions at $\mathcal{I}$, which will be smeared local sources.¹³ This sideways evolution was also used in [25] to prove the simple entropy conjecture in the case where the horizon was only perturbatively non-stationary.

We will shortly demonstrate (2) in detail. The process is similar in spirit to the zigzag process in Sec. 3, but instead of removing infalling chiral modes, we apply our perturbation in Sec. 4. First, we will discuss the consequences of repeated iterations of our perturbation on the future horizon and then add timefolds into our procedure.

Let us begin by a comparison between the perturbed and unperturbed causal surfaces, denoted C' and C respectively, as a result of one instance of our perturbation on $\mathcal{H}_C^+$. Since the perturbation is localized away from the past event horizon $\mathcal{H}^-$, it is expedient to compare the relative location of C with that of C' using their position on the past event horizon: both are slices of $\mathcal{H}^-$. Note that in the perturbed geometry no special role is played by C . As shown in Sec. 4, the perturbation guarantees that C' is “inwards” (i.e. at larger u) compared with C so long as $\theta_{(k)}[C] \neq 0$. Nothing stops us from then repeating the perturbation above on the *new* future horizon. As long as

¹²No pun intended.

¹³We are aware that some mathematician somewhere must be apoplectic with rage after reading this paragraph. Since we are physicists, we will nevertheless persevere under the usual assumption in holography that HKLL does in fact work out. Concerned readers who do not a priori wish to subscribe to HKLL may take comfort in the following interpretation of our results: we show that reconstruction of the simple wedge is no more complex than that of the causal wedge.

some point on the causal surface satisfies $\theta_{(k)} > 0$, the perturbation pushes the causal wedge further inwards.

The only obstruction in the construction above occurs if $\theta_{(k)}$ vanishes identically on the causal surface. Thus it is clear that the inwards shift of the causal surface obtained via simple sources limits to an outermost marginally outer trapped surface μ . (We can define a rigorous notion of the causal surface approaching arbitrarily close to a surface with $\theta_{(k)} = 0$ by picking an affine parameterization on $\mathcal{H}^-$ and defining proximity of the two surfaces in terms of the maximal elapsed affine parameter between them.)

Let us provide intuition for the existence of μ . Say on $\mathcal{H}^-$ we can identify two cuts μ_1 and μ_2 such that μ_2 encloses μ_1 and $\theta_{(k)}[\mu_1] \leq 0$, $\theta_{(k)}[\mu_2] \geq 0$. Then we expect an outermost marginally outer trapped surface μ in-between μ_1 and μ_2 .¹⁴ On $\mathcal{H}^-$, C plays the role of μ_2 . For μ_1 , we can pick $\partial J^+[X] \cap \mathcal{H}^-$ – whenever it is a full cross section of $\mathcal{H}^-$ – which satisfies $\theta_{(k)} \leq 0$ by the focusing theorem. However, note that $\partial J^+[X] \cap \mathcal{H}^-$ might be empty if $\mathcal{H}^-$ falls into a singularity before intersecting $\partial J^+[X]$. Even so, at least for Kasner-like singularities, we can find cross sections of $\mathcal{H}^-$ in a neighborhood of the singularity which are trapped [31]. Note also that generically, our choices for μ_1 and μ_2 satisfy $\theta_{(k)}[\mu_1] < 0$, $\theta_{(k)}[\mu_2] > 0$. We will then have candidates for both μ_1 and μ_2 , so μ exists.

Prima facie the procedure at this point appears to have failed! The causal surface will generically stop well away from null separation with X , and even further away from coincidence with X . However, this is only true on the particular timefold under consideration. To proceed to close the gap further, we reverse the arrow of time. We can then repeat the procedure above in time reverse to shrink the discrepancy between the causal wedge and W_X . We iterate this procedure via forward and reverse timefolds; each step brings the causal surface and X closer. Just like for the JT gravity case in Section 3, the causal surface should limit to the outermost extremal surface after sufficiently many timefolds. Importantly, since the bulk physics involved is entirely classical, the number of timefolds required for the causal surface to approach the outermost extremal surface, within a given precision, should be independent of the Planck scale. This means that the complexity of the process cannot diverge in the classical limit $N \rightarrow \infty$. We therefore conclude that for X and C of identical topology, the simple wedge and the outermost extremal wedge coincide.

Finally, we address (3): throughout this construction, the geometry of the lunch is left unaltered: the perturbation is localized in the causal complement of the lunch, guaranteeing that the lunch remains undisturbed.

¹⁴This is proved when μ_1 and μ_2 satisfy $\theta_{(k)}[\mu_1] < 0$, $\theta_{(k)}[\mu_2] > 0$ and lie on smooth spacelike slices in $D = 4$ in [51], but since we can approach $\mathcal{H}^-$ with such spacelike slices we expect the same to hold for $\mathcal{H}^-$, and the technology used in the extant proof is expected to generalize to higher dimensions.

What about the case where the topologies are different? This could for example be the case in a time-symmetric and spherically symmetric null shell collapse in AdS where C is a sphere and $X = \emptyset$ [52]. If such topology difference exists, we would expect it to present itself between C and the outermost marginally trapped surface on $\mathcal{H}^-$ in a timefold (or several timefolds) of our procedure above. Furthermore, because each iteration of our $\mathcal{H}_C^+$ perturbation pushes the causal surface inwards on $\mathcal{H}^-$ a bit, at some point the jump from C to C' would have to involve a topology change. As there is nothing explicit in our construction that constrains the topology of C' according to that of C , we do not see a fundamental obstruction against such topology changes arising from our iterative procedure. A rigorous treatment of such cases might be interesting but is left to future work.

6 Simplicity Killed the Python

Having established that the simple wedge is in fact reconstructible using exclusively simple experiments, we now explore the implications of our results beyond the converse to the Python’s lunch: what is the dual to the area of the outermost extremal surface? What is the field theory interpretation of our results, and in particular what is the “simple state” dual to the simple wedge?

6.1 Time-Folded Simple Entropy

As noted in the introduction, the simple entropy of [24, 25] is a coarse-graining over high complexity measurements conducted after a fixed boundary time t_{bdy} on a single timefold:

$$S_{\text{simple}}[t_{\text{bdy}}, \rho_{\text{bdy}}] = \max_{\rho \in \mathcal{B}} S_{\text{vN}}[\rho] \quad (6.1)$$

where ρ_{bdy} is the actual state of the CFT, t_{bdy} is a choice of boundary time slice, and $\mathcal{B}$ is the set of all CFT states (density matrices) that have the same one-point functions as ρ_{bdy} under any simple sources turned on after the time t_{bdy} (and with some very late time cutoff to avoid recurrences). That is, $\mathcal{B}$ consists of the set of CFT states ρ such that

$$\langle E \mathcal{O} E^\dagger \rangle_{\rho_{\text{bdy}}} = \langle E \mathcal{O} E^\dagger \rangle_{\rho}, \quad (6.2)$$

for all possible E defined as in Eq. 2.1. The simple entropy at a given boundary time is thus a coarse-graining over high complexity data that preserves all of the simple data to the future (or past) of that time.

With these restrictions to a particular subset of boundary time and a fixed timefold, the simple entropy was proposed as a dual to the outer entropy, which is a bulk-defined quantity that coarse-grains over the exterior of an apparent horizon (a surface which is by definition always marginally trapped¹⁵). The outer entropy coarse-grains over all possible spacetimes that look identical outside of a given apparent horizon to find the spacetime with the largest HRT surface, and thus the largest von Neumann entropy in the CFT:

$$S_{\text{outer}}[\mu] = \max_{\mathcal{X}} \frac{\text{Area}[X]}{4G\hbar} = \max_{\rho \in \mathcal{H}} S_{\text{vN}}[\rho] = \frac{\text{Area}[\mu]}{4G\hbar}, \quad (6.3)$$

where $\mathcal{X}$ consists of the HRT surfaces of all classical holographic spacetimes containing $O_W[\mu]$, and $\mathcal{H}$ is the corresponding set of CFT states; the final equality is proved in [24, 25]. This is done by discarding the spacetime behind μ , constructing a spacetime with an HRT surface X_μ whose area is identical to the area of μ , and then CPT-conjugating the spacetime around X_μ . By construction, $O_W[\mu]$ is left unaltered.

The proposal that the simple and outer entropies are identical says that there is a particular definition of black hole entropy which is a consequence of coarse-graining over the highly complex physics that we expect describes the interior:

$$S_{\text{outer}}[\mu(t_{\text{bdy}})] = S_{\text{simple}}[t_{\text{bdy}}] \quad (6.4)$$

where $t_{\text{bdy}} = \partial J^-[\mu] \cap \mathcal{I}$. Our construction in Section 4 establishes this conjecture for apparent horizons: in a given timefold, it is possible to push the event horizon all the way up to $\mu(t_{\text{bdy}})$ without accessing any high complexity data for $t > t_{\text{bdy}}$.

Our construction is of course more general, as it applies to timefolds. Extending the simple entropy proposal to include timefolds immediately yields the holographic dual to the area of the outermost extremal surface X :

$$\frac{\text{Area}[X]}{4G\hbar} = S_{\text{outer}}[X] = S_{\text{simple}} \quad (6.5)$$

where S_{simple} is obtained from $S_{\text{simple}}[t_{\text{bdy}}]$ by taking $t_{\text{bdy}} \rightarrow -\infty$ and including arbitrary timefolds. The inclusion of timefolds removes the need for a reference apparent horizon, and the coarse-grained spacetime (in which the outermost extremal surface is in fact the HRT surface) is obtained by CPT-conjugating around the outermost extremal surface X ; see Fig. 3. Crucially, note that the coarse-graining procedure leaves the outermost extremal wedge untouched and coarse-grains only over the lunch. Standard entanglement wedge reconstruction via quantum error correction [53–55] applies

¹⁵A refinement of an apparent horizon that was named a “minimar surface”. Apparent horizons are generically minimar surfaces.

to reconstruction of the outermost extremal wedge, since in the coarse-grained spacetime obtained by CPT conjugation, the outer wedge of X is exactly the entanglement wedge. Since we will argue below that the coarse-grained spacetime has a simple modular Hamiltonian, entanglement wedge reconstruction using, for example, modular flow as in [56] should be much simpler when based on the coarse-grained state rather than the original state. This is consistent with the simplicity of reconstructing the outermost extremal wedge.

6.2 The Simple State

So far we have introduced two manipulations that can be done to our original spacetime. The first was the zigzag procedure, introduced in Section 5 which makes the causal wedge coincide with (or become arbitrarily close to) the outermost extremal wedge. We say that the resulting spacetime is ‘exposed’ because everything in the simple wedge can be directly seen by the boundary.¹⁶ The second is the coarse-graining procedure introduced above where we CPT conjugate about the outermost extremal surface and thereby create a state where the outermost extremal wedge coincides with the entanglement wedge. If we apply *both* manipulations to the spacetime, we can produce a spacetime where all three wedges (approximately) coincide. This is illustrated in Fig. 6.

That is, given any holographic CFT state ρ dual to some entanglement wedge (which will likely have a Python’s lunch), there exists a ρ_{coarse} which is indistinguishable from ρ via simple experiments and has no Python’s lunch. Executing our procedure zigzag on this coarse-grained state yields the coarse-grained exposed spacetime, in which the causal, simple, and entanglement wedges all coincide or come arbitrarily close to coinciding. The dual to this is described by the state obtained via the zigzag procedure together with the set of simple operators that remains from the final timefold. We shall refer to this history that includes the the remaining simple operators the simple history, and in a slight abuse of notation we will denote this entire history of states as ρ_{simple} .

What is the CFT interpretation of ρ_{simple} ? Below we prove that the causal and entanglement wedges coincide exactly if and only if the state dual to the entanglement wedge has a geometric modular flow. The immediate implication is that in the case where the zigzag procedure gives an exact coincidence, the simple state has an exactly local modular Hamiltonian, and rather than being a history it is in fact a single state. Note that this is suggestive of a CFT dual to a gravitational no-hair theorem. To be

¹⁶Note that this definition is unrelated to our definition of an ‘exposed surface’ in the proof of Proposition 1.




Figure 6. A diagram illustrating the relationships between the different states: the original CFT state ρ , which may be either coarse-grained to obtain ρ_{coarse} by forgetting about high complexity operators, or it may be “exposed” by acting on it with simple operators. The two operations modify causally independent and non-intersecting portions of the spacetime, so they commute: after obtaining the coarse-grained state in which X is the HRT surface, we may perform our zigzag procedure to push the causal surface up to X and obtain the simple state in which all three wedges coincide. Or, after obtaining the exposed state in which the causal and outermost extremal wedge coincide, but the entanglement wedge properly contains both, we may coarse-grain to obtain the same simple state.

more precise: this result suggests that the set of stationary holographic black holes is to be identified with the high limited set of states with exactly local modular flow. If coincidence between the causal surface and the appetizer is asymptotic rather than exact, then we come to the conclusion that the modular flow generated by the simple state is very close to local in the sense that only operators with support in the asymptotically

shrinking region between the causal and entanglement wedges are able to definitively tell that the two are not identical. Since that region translates (for simple operators) to access to arbitrarily late or early times, we find that finite-time simple measurements are unable to tell that the modular flow generated by ρ_{simple} at each stage is not local.

The secondary implication is that it is possible to take any holographic state ρ and, via a series of simple operations, render its modular flow (nearly) indistinguishable from a geometric flow via any simple experiments. If the appetizer has sufficient symmetry, then the statement should be true exactly. Let us now prove our theorem, which we do in broad generality for boundary subregions.

Theorem 1. *Let $\mathcal{W}_C[R]$ denote the causal wedge of a boundary subregion R , and let $W_E[R]$ denote the entanglement wedge of the boundary subregion. $\mathcal{W}_C[R] = W_E[R]$ if and only if the boundary modular Hamiltonian on R generates a geometric flow with respect to a boundary Killing vector field on R .*

Proof. Assume that the boundary modular Hamiltonian generates a geometric flow with respect to some Killing vector field ξ^I on ∂M (here I is a boundary spacetime index). Under modular flow, a local operator is mapped to another local operator:

$$\mathcal{O}(x, s) = \rho_R^{-is/2\pi} \mathcal{O}(x) \rho_R^{is/2\pi} = \mathcal{O}(x_\xi(s)) \quad (6.6)$$

where $x_\xi(s)$ is the boost along ξ^I of $x \in D[R]$. Via [56], $W_E[R]$ can be reconstructed by smearing the modular flow of local operators over $D[R]$. Since in this case modular flow is an automorphism on the space of local operators in $D[R]$, we can reconstruct all operators in $W_E[R]$ by smearing local operators:

$$\Phi(x) = \int_{D[R]} f(X|x) \mathcal{O}(X) dx, \quad (6.7)$$

where $f(X|x)$ is a smearing function. If there is a gap between $W_E[R]$ and $\mathcal{W}_C[R]$, then operators that are localized to the gap should commute with all local operators on the boundary (within our code subspace) via the extrapolate dictionary (note that this only works in the large- N limit where we don't have to worry about gravitational dressing). However this is inconsistent with the equation above; so $x \in \mathcal{W}_C[R]$. But this argument holds for all local operators: there exist no local operators in the gap between $W_X[R]$ and $\mathcal{W}_C[R]$. In the large- N limit (without backreaction), this means that there simply is no gap between the two wedges: $W_X[R] = \mathcal{W}_C[R]$.

To prove the other direction, we consider the proof of [56] for the zero-mode formula of entanglement wedge reconstruction (appendix B.1 of [56]). Starting with equation B.71, it is shown that the nonlocality of modular flow on $\partial W_E[R]$ is due to the change in

the intrinsic metric of spatial slices of $\partial W_E[R]$ (in particular, the loss of ultralocality). When $\partial W_E[R]$ is stationary, the metric on codimension-two slices does not change with evolution along the congruence. This means that the modular flow on $\partial W_E[R]$ is local, which in turn implies that the boundary modular flow is local as well. $\square$

7 Discussion

Our primary technical result in this article is the proof of the converse to the Python’s lunch proposal in the strict large- N limit: operators that lie outside of a Python’s lunch are simply reconstructible in the dual CFT, and moreover this reconstruction only relies on the bulk dynamics in the large- N limit, manifestly respecting the causal structure of the background metric.

We emphasize that bulk reconstructions that are causal in this sense cannot work for the interior of the Python’s lunch because no causal horizon can intersect the lunch. The CFT encoding of the Python’s lunch appears to involve highly non-local quantum gravity effects. An example of such non-local dynamics is the ER=EPR conjecture [57, 58] which asserts that the entanglement between an evaporating black hole and its Hawking radiation after the Page time must allow complicated operations on the distant radiation to change the state behind the lunch, drastically violating the naive causal structure dictated by the background metric. It has been speculated that wormhole-like “corrections” to the background geometry connecting the radiation to the black hole interior could explain the “true” causal structure not captured by the background metric. It is natural to speculate that similar dynamics are at play in the Python’s lunch encoding into the boundary.

We will now discuss various generalizations of our main results:

7.1 Boundary-Anchored Surfaces

We have focused here primarily on compact surfaces, but we may pose similar questions for boundary subregions: given the state ρ_R on a boundary subregion R , how complex is the reconstruction of operators behind the event horizon but within the outermost extremal wedge? This requires a treatment of surfaces with a boundary-anchored component rather than surfaces whose components are all compact. Most of our results generalize almost immediately to the boundary-anchored case: Lemma 3 [38] makes no reference the topology of surfaces (beyond the homology constraint); similarly for the proofs of Lemmas 1 and 2. The perturbed initial data prescription at the causal surface also carries over mutatis mutandis. As noted in Section 4, the weak comparison principle operates on the basis of the maximum principle for elliptic operators. The

latter does indeed apply to bounded domains in general, and to boundary-anchored hypersurfaces in AdS particular (see [59] for a discussion in the context of AdS/CFT). The main potential source of difficulty is the falloff: both the causal surface and the outermost extremal surface approach the asymptotic boundary, but not with the same tangent space; the asymptotic falloff of δg must approach zero sufficiently fast so as to not spoil the asymptotics while bridging the gap between the two surfaces. We expect that the appropriate falloff conditions can be satisfied, but this remains a subject for future work.

7.2 Robustness under Quantum Corrections

Fundamentally, the classical calculations done in this paper are only interesting as an approximation to the fully quantum dynamics that actually describe the bulk in AdS/CFT. Do our arguments extend to the semiclassical setting where the background spacetime is still treated classically, but with quantum fields propagating on it? Do they generalize to the regime where perturbative corrections to the geometry, suppressed by powers $G\hbar$, are allowed to contribute? A number of important assumptions break down in this case: Raychaudhuri’s equation is still valid, but the null energy condition will not generally hold, and so light rays emanating from a classical extremal surface can defocus. Fortunately, the quantum focusing conjecture (QFC) states that the generalized entropy of null congruences emanating from QESs is always subject to focusing. In a semiclassical or perturbatively quantum bulk, the appetizer is the outermost quantum, not classical, extremal surface; the QFC ensures that the outermost quantum extremal wedge always contains the causal wedge.¹⁷ The question is whether we can still expand the causal wedge using appropriate sources and timefolds in order to bridge the gap between the causal surface and the appetizer.

This is a much harder question than the classical question discussed here: the class of allowed QFT states is simply much harder to classify and use than classical field theory states. However, in particularly simple examples, for instance where the causal wedge is approximately Rindler-like, a quantum version of the “left stretch” appears to be well defined and gives exactly the right change in energy to reduce focusing and remove the perturbatively small, i.e. $O(1/N)$, distance between the causal and outermost extremal wedges [26]. It is feasible that in the limit of many zigzags, the causal wedge would approach a Rindler-like region, at which point it becomes possible

¹⁷Apparent counterexamples are only possible when the time evolution includes interaction with an auxiliary system; once you include these auxiliary interacting modes when defining quantum extremality then the apparent causality issues go away.

to apply the bulk unitaries discussed in [26] to eliminate the remaining small gap. We leave a detailed study of this to future work.

7.3 Asymptotically Flat Spacetimes

In asymptotically flat spacetimes, asymptotic infinity is lightlike rather than timelike. However, there do not seem to be any major obstructions to adapting the results of this paper to that setting. Instead of the timefolds of the asymptotically AdS problem, in asymptotically flat space one would presumably evolve forwards along future null infinity, then backwards with different boundary conditions that remove focusing at the past event horizon, in order to produce a state where the causal surface is very close to a past apparent horizon. Then one would evolve backwards and forwards along past null infinity in order to produce a state where the causal surface is very close to a future apparent horizon. At each step, the causal wedge increases in size. After sufficiently many such timefolds, the causal surface should approach the outermost extremal surface, as in the asymptotically AdS case. The interpretation of our results in the asymptotically flat case is naturally obfuscated by the relatively inchoate status of flat holography. We may speculate that extremal surfaces are important more generally for defining entropy in gravity; it is also possible that a similar notion of a Python’s lunch applies beyond AdS holography. We do not subscribe to any particular interpretation – here we simply note that the technical aspects of this work are likely not restricted to AdS.

7.4 Complexity Censorship

Let us finish with a few comments on cosmic censorship, a *prima facie* unrelated conjecture about classical General Relativity. (Weak) cosmic censorship [39] is essentially the statement that high curvature physics lies behind event horizons. One of its landmark consequences is that trapped surfaces lie behind event horizons, and that consequently marginally trapped and in particular extremal surfaces lie on or behind event horizons [37]. It is clear from the above discussion that any violation of cosmic censorship would be quite problematic for the Python’s lunch picture: if the nonminimal extremal surface could lie outside of the event horizon (or, in the quantum case, could communicate with $\mathcal{S}$), then operators behind would lie properly within the causal wedge and would thus be reconstructible by HKLL despite being exponentially complex. The Python’s lunch proposal thus appears to depend heavily on the validity of cosmic censorship – which is known to be violated in AdS [60–62]. As matters currently stand, violations of cosmic censorship notwithstanding, it is possible to prove that the holographic entanglement entropy prescription guarantees that trapped surfaces must lie

behind event horizons [42]. We could however have proven the same statement from holographic complexity: marginally trapped surfaces (and therefore, trapped surfaces also, by the reasoning of [42]) must lie behind event horizons, for if they did not, operators behind the Python’s lunch could be reconstructed in a simple procedure. This suggests that in AdS/CFT, aspects of cosmic censorship may be reformulated as “complexity censorship”: that high complexity physics must be causally hidden and thus unable to causally communicate to $\mathcal{I}$.

Acknowledgments

It is a pleasure to thank S. Alexakis, R. Bousso, S. Fischetti, L. Susskind for helpful discussions. NE is supported by NSF grant no. PHY-2011905, by the U.S. Department of Energy under grant no. DE-SC0012567 (High Energy Theory research), and by funds from the MIT physics department. GP is supported by the UC Berkeley physics department, the Simons Foundation through the “It from Qubit” program, the Department of Energy via the GeoFlow consortium and also acknowledges support from a J. Robert Oppenheimer Visiting Professorship. ASM is supported by the National Science Foundation under Award Number 2014215.

A uv and vi constraints

Here we write down the general structure of the perturbative constraints (4.8b), (4.8c), (4.8d):

$$\frac{1}{2}g^{ij}\partial_v\partial_u\delta g_{ij} = (\hat{L}_1)^{ab}\partial_u\delta g_{ab} + (\hat{L}_2)^{ab}\delta g_{ab} + 8\pi G\delta T_{uv} \quad (\text{A.1a})$$

$$\partial_v\partial_u\delta g_{vi} = (\hat{L}'_1)^{ab}_i\partial_u\delta g_{ab} + (\hat{L}'_2)^{ab}_i\delta g_{ab} + 8\pi G\delta T_{vi} \quad (\text{A.1b})$$

$$\partial_v\partial_u\delta \hat{g}_{ij} = (\hat{L}''_1)^{ab}_{ij}\partial_u\delta g_{ab} + (\hat{L}''_2)^{ab}_{ij}\delta g_{ab} + 8\pi G\delta \hat{T}_{ij} \quad (\text{A.1c})$$

where $\hat{L}_n$, $\hat{L}'_n$, and $\hat{L}''_n$ represent linear differential operators that depend only on the background metric and do not have u derivatives. The only significance of Eqs. (A.1a), (A.1b), and (A.1c) for us is that they are consistent with the δg_{vv} solution to Eq. (4.17). Plugging in the δg_{vv} solution reduces Eqs. (A.1a), (A.1b), and (A.1c) to coupled linear differential equations involving $\partial_u\delta g_{ij}|_{u=0}$ and $\partial_u\delta g_{vi}|_{u=0}$ which are first order in ∂_v . Note that the initial conditions at $v = 0$ are $\partial_u g_{ij}(u = 0, v = 0^-, y) = \partial_u g_{vi}(u = 0, v = 0^-, y) = 0$, as enforced by the perturbation being localized away from $\mathcal{H}^-$.

B Alternative derivation of the vv constraint

Here we provide an alternative derivation for the same perturbation discussed in Sec. 4. Instead of the transformations (4.9a), (4.9b), and (4.9c) on $\mathcal{H}_C^+$, we can equivalently insert an inaffinity shock at $v = 0$ [44, 45]:

$$\kappa_{(v)} = (1 - e^{-s})\delta(v) \quad (\text{B.1})$$

and take the $(1 - e^{-s}) \sim \epsilon$ limit. In addition, we want to introduce the following δg_{ab} transformation:

$$ds^2 = -2dudv + (g_{vv} + \delta g_{vv})dv^2 + 2(g_{vi} + \delta g_{vi})dvdy^i + (g_{ij} + \delta g_{ij})dy^i dy^j, \quad (\text{B.2})$$

with

$$\delta g_{vi}|_{u=0} = 0 \quad (\text{B.3})$$

$$\delta g_{ij}|_{u=0} = 0 \quad (\text{B.4})$$

$$\partial_u \delta g_{vv}|_{u=0} = 0 \quad (\text{B.5})$$

We need to apply the combined perturbations in Eqs. (B.1) and (B.2) to the vv constraint:

$$\delta G_{vv} + \Lambda \delta g_{vv} = 8\pi G \delta T_{vv} \quad (\text{B.6})$$

The only contribution in δG_{vv} from Eq. (B.1) is through the $\kappa_{(v)}\theta_{(v)}$ term in Eq. (4.13). Summing this up with the contribution from the contribution from δg_{ab} of Eq. (B.2), we get:

$$\begin{aligned} & -\frac{1}{2}\theta_{(u)}\partial_v\delta g_{vv} - \frac{1}{2}\nabla_\perp^2\delta g_{vv} + \chi^i\partial_i\delta g_{vv} + \\ & \left(\nabla_\perp\cdot\chi - \partial_v\theta_{(u)} - B_{(v)ij}B_{(u)}^{ij} + 8\pi G(-T_{uv} - \mathcal{L}_{\text{matter}} + F_{uv}^2)\right)\delta g_{vv} \\ & - \epsilon\theta_{(v)}[C]\delta(v) = 0 \end{aligned} \quad (\text{B.7})$$

Similarly, the uv and vi constraints could be analyzed resulting in Eqs. (A.1a) and (A.1b).

References

- [1] A. Almheiri, N. Engelhardt, D. Marolf, and H. Maxfield, *The entropy of bulk quantum fields and the entanglement wedge of an evaporating black hole*, *JHEP* **12** (2019) 063, [[arXiv:1905.08762](#)].

- [2] G. Penington, *Entanglement Wedge Reconstruction and the Information Paradox*, [arXiv:1905.08255](#).
- [3] A. Almheiri, R. Mahajan, J. Maldacena, and Y. Zhao, *The Page curve of Hawking radiation from semiclassical geometry*, *JHEP* **03** (2020) 149, [[arXiv:1908.10996](#)].
- [4] A. Almheiri, T. Hartman, J. Maldacena, E. Shaghoulian, and A. Tajdini, *Replica Wormholes and the Entropy of Hawking Radiation*, *JHEP* **05** (2020) 013, [[arXiv:1911.12333](#)].
- [5] G. Penington, S. H. Shenker, D. Stanford, and Z. Yang, *Replica wormholes and the black hole interior*, [arXiv:1911.11977](#).
- [6] 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](#)].
- [7] N. Engelhardt and A. C. Wall, *Quantum Extremal Surfaces: Holographic Entanglement Entropy beyond the Classical Regime*, *JHEP* **01** (2015) 073, [[arXiv:1408.3203](#)].
- [8] D. N. Page, *Information in black hole radiation*, *Phys. Rev. Lett.* **71** (1993) 3743–3746, [[hep-th/9306083](#)].
- [9] P. Hayden and J. Preskill, *Black holes as mirrors: quantum information in random subsystems*, *JHEP* **09** (2007) 120, [[arXiv:0708.4025](#)].
- [10] D. Harlow and P. Hayden, *Quantum Computation vs. Firewalls*, *JHEP* **1306** (2013) 085, [[arXiv:1301.4504](#)].
- [11] R. Bousso, Z. Fisher, S. Leichenauer, and A. C. Wall, *Quantum focusing conjecture*, *Phys. Rev.* **D93** (2016), no. 6 064044, [[arXiv:1506.02669](#)].
- [12] 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.* **D73** (2006) 086003, [[hep-th/0506118](#)].
- [13] A. Hamilton, D. N. Kabat, G. Lifschytz, and D. A. Lowe, *Holographic representation of local bulk operators*, *Phys.Rev.* **D74** (2006) 066009, [[hep-th/0606141](#)].
- [14] 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.* **D75** (2007) 106001, [[hep-th/0612053](#)]. [Erratum: *Phys. Rev.* **D75**, 129902(2007)].
- [15] I. Heemskerk, *Construction of Bulk Fields with Gauge Redundancy*, *JHEP* **09** (2012) 106, [[arXiv:1201.3666](#)].
- [16] D. Kabat, G. Lifschytz, S. Roy, and D. Sarkar, *Holographic representation of bulk fields with spin in AdS/CFT*, *Phys. Rev.* **D86** (2012) 026004, [[arXiv:1204.0126](#)].

- [17] I. Heemskerk, D. Marolf, J. Polchinski, and J. Sully, *Bulk and Transhorizon Measurements in AdS/CFT*, *JHEP* **10** (2012) 165, [[arXiv:1201.3664](#)].
- [18] D. Kabat, G. Lifschytz, and D. A. Lowe, *Constructing local bulk observables in interacting AdS/CFT*, *Phys. Rev. D* **83** (2011) 106009, [[arXiv:1102.2910](#)].
- [19] D. Sarkar and X. Xiao, *Holographic Representation of Higher Spin Gauge Fields*, *Phys. Rev. D* **91** (2015), no. 8 086004, [[arXiv:1411.4657](#)].
- [20] E. Holmgren, *über systeme von linearen partiellen differentialgleichungen*, *Öfversigt af Kongl. Vetenskaps-Academien Förhandlingar*. (1901), no. 58 91–103.
- [21] D. Tataru, *Unique continuation problems for partial differential equations:; between hörmänder’s theorem and holmgren’s theorem*, in *Geometric Methods in Inverse Problems and PDE Control* (C. B. Croke, M. S. Vogelius, G. Uhlmann, and I. Lasiecka, eds.), (New York, NY), pp. 239–255, Springer New York, 2004.
- [22] R. Bousso, B. Freivogel, S. Leichenauer, V. Rosenhaus, and C. Zukowski, *Null Geodesics, Local CFT Operators and AdS/CFT for Subregions*, *Phys.Rev.* **D88** (2013) 064057, [[arXiv:1209.4641](#)].
- [23] A. Almheiri, R. Mahajan, and J. Maldacena, *Islands outside the horizon*, [arXiv:1910.11077](#).
- [24] N. Engelhardt and A. C. Wall, *Decoding the Apparent Horizon: Coarse-Grained Holographic Entropy*, *Phys. Rev. Lett.* **121** (2018), no. 21 211301, [[arXiv:1706.02038](#)].
- [25] N. Engelhardt and A. C. Wall, *Coarse Graining Holographic Black Holes*, *JHEP* **05** (2019) 160, [[arXiv:1806.01281](#)].
- [26] A. Levine, A. Shahbazi-Moghaddam, and R. M. Soni, *Seeing the Entanglement Wedge*, [arXiv:2009.11305](#).
- [27] R. Jackiw, *Lower Dimensional Gravity*, *Nucl. Phys. B* **252** (1985) 343–356.
- [28] C. Teitelboim, *Gravitation and Hamiltonian Structure in Two Space-Time Dimensions*, *Phys. Lett. B* **126** (1983) 41–45.
- [29] S. Dutta and T. Faulkner, *A canonical purification for the entanglement wedge cross-section*, [arXiv:1905.00577](#).
- [30] N. Engelhardt, G. Penington, and A. Shahbazi-Moghaddam, “Finding Pythons in Unexpected Places (to appear).”.
- [31] A. C. Wall, *Maximin Surfaces, and the Strong Subadditivity of the Covariant Holographic Entanglement Entropy*, *Class.Quant.Grav.* **31** (2014), no. 22 225007, [[arXiv:1211.3494](#)].
- [32] R. M. Wald, *General Relativity*. The University of Chicago Press, Chicago, 1984.

- [33] R. Bousso and N. Engelhardt, *Proof of a New Area Law in General Relativity*, *Phys. Rev. D* **92** (2015), no. 4 044031, [[arXiv:1504.07660](#)].
- [34] C. Akers, R. Bousso, I. F. Halpern, and G. N. Remmen, *The Boundary of the Future*, [arXiv:1711.06689](#).
- [35] V. E. Hubeny and M. Rangamani, *Causal Holographic Information*, *JHEP* **1206** (2012) 114, [[arXiv:1204.1698](#)].
- [36] S. W. Hawking, *Gravitational radiation from colliding black holes*, *Phys. Rev. Lett.* **26** (1971) 1344–1346.
- [37] S. W. Hawking and G. F. R. Ellis, *The large scale structure of space-time*. Cambridge University Press, Cambridge, England, 1973.
- [38] R. Bousso and A. Shahbazi-Moghaddam, *Island Finder and Entropy Bound*, [arXiv:2101.11648](#).
- [39] R. Penrose, *Gravitational collapse: The role of general relativity*, *Riv. Nuovo Cim.* **1** (1969) 252–276. [*Gen. Rel. Grav.* **34**, 1141 (2002)].
- [40] R. P. Geroch and G. T. Horowitz, *Global structure of spacetimes*, in *General Relativity: An Einstein Centenary Survey*, pp. 212–293. 1979.
- [41] M. Headrick, V. E. Hubeny, A. Lawrence, and M. Rangamani, *Causality & holographic entanglement entropy*, *JHEP* **12** (2014) 162, [[arXiv:1408.6300](#)].
- [42] N. Engelhardt and r. Folkestad, *Holography Abhors Visible Trapped Surfaces*, [arXiv:2012.11445](#).
- [43] N. Engelhardt and S. Fischetti, *The Gravity Dual of Boundary Causality*, *Class. Quant. Grav.* **33** (2016), no. 17 175004, [[arXiv:1604.03944](#)].
- [44] R. Bousso, V. Chandrasekaran, and A. Shahbazi-Moghaddam, *From black hole entropy to energy-minimizing states in QFT*, *Phys. Rev. D* **101** (2020), no. 4 046001, [[arXiv:1906.05299](#)].
- [45] R. Bousso, V. Chandrasekaran, P. Rath, and A. Shahbazi-Moghaddam, *Gravity dual of Connes cocycle flow*, *Phys. Rev. D* **102** (2020), no. 6 066008, [[arXiv:2007.00230](#)].
- [46] L. Andersson, G. J. Galloway, and R. Howard, *A strong maximum principle for weak solutions of quasi-linear elliptic equations with applications to lorentzian and riemannian geometry*, *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences* **51** (1998), no. 6 581–624.
- [47] D. Gilbarg and N. Trudinger, *Elliptic Partial Differential Equations of Second Order*. Classics in Mathematics. U.S. Government Printing Office, 2001.
- [48] G. J. Galloway, *Maximum principles for null hypersurfaces and null splitting theorems*, *Annales Henri Poincare* **1** (2000) 543–567, [[math/9909158](#)].

- [49] H. Rademacher, *Über partielle und totale differenzierbarkeit von funktionen mehrerer variablen und über die transformation der doppelintegrale*, *Mathematische Annalen* **79** (1919), no. 4 340–359.
- [50] V. Perlick, *Gravitational lensing from a spacetime perspective*, *Living Reviews in Relativity* **7** (Sep, 2004) 9.
- [51] L. Andersson and J. Metzger, *The Area of horizons and the trapped region*, *Commun. Math. Phys.* **290** (2009) 941–972, [[arXiv:0708.4252](#)].
- [52] N. Engelhardt and A. C. Wall, *No Simple Dual to the Causal Holographic Information?*, *JHEP* **04** (2017) 134, [[arXiv:1702.01748](#)].
- [53] A. Almheiri, X. Dong, and D. Harlow, *Bulk Locality and Quantum Error Correction in AdS/CFT*, *JHEP* **04** (2015) 163, [[arXiv:1411.7041](#)].
- [54] X. Dong, D. Harlow, and A. C. Wall, *Reconstruction of Bulk Operators within the Entanglement Wedge in Gauge-Gravity Duality*, *Phys. Rev. Lett.* **117** (2016), no. 2 021601, [[arXiv:1601.05416](#)].
- [55] J. Cotler, P. Hayden, G. Penington, G. Salton, B. Swingle, and M. Walter, *Entanglement Wedge Reconstruction via Universal Recovery Channels*, *Phys. Rev. X* **9** (2019), no. 3 031011, [[arXiv:1704.05839](#)].
- [56] T. Faulkner and A. Lewkowycz, *Bulk locality from modular flow*, *JHEP* **07** (2017) 151, [[arXiv:1704.05464](#)].
- [57] M. Van Raamsdonk, *Evaporating Firewalls*, *JHEP* **11** (2014) 038, [[arXiv:1307.1796](#)].
- [58] J. Maldacena and L. Susskind, *Cool horizons for entangled black holes*, [arXiv:1306.0533](#).
- [59] N. Engelhardt and S. Fischetti, *Surface Theory: the Classical, the Quantum, and the Holographic*, *Class. Quant. Grav.* **36** (2019), no. 20 205002, [[arXiv:1904.08423](#)].
- [60] G. T. Horowitz, J. E. Santos, and B. Way, *Evidence for an Electrifying Violation of Cosmic Censorship*, *Class. Quant. Grav.* **33** (2016), no. 19 195007, [[arXiv:1604.06465](#)].
- [61] T. Crisford and J. E. Santos, *Violating the Weak Cosmic Censorship Conjecture in Four-Dimensional Anti-de Sitter Space*, *Phys. Rev. Lett.* **118** (2017), no. 18 181101, [[arXiv:1702.05490](#)].
- [62] G. T. Horowitz and J. E. Santos, *Further evidence for the weak gravity — cosmic censorship connection*, *JHEP* **06** (2019) 122, [[arXiv:1901.11096](#)].