

Topological aspects of generalized gravitational entropy

Felix M. Haehl,^a Thomas Hartman,^b Donald Marolf,^c Henry Maxfield^a and Mukund Rangamani^a

^a*Centre for Particle Theory & Department of Mathematical Sciences, Durham University, South Road, Durham DH1 3LE, U.K.*

^b*Department of Physics, Cornell University, Ithaca, New York 14853, U.S.A.*

^c*Department of Physics, University of California, Santa Barbara, Santa Barbara, CA 93106, U.S.A.*

E-mail: f.m.haehl@gmail.com, hartman@cornell.edu,
marolf@physics.ucsb.edu, h.d.maxfield@durham.ac.uk,
mukund.rangamani@durham.ac.uk

ABSTRACT: The holographic prescription for computing entanglement entropy requires that the bulk extremal surface, whose area encodes the amount of entanglement, satisfies a homology constraint. Usually this is stated as the requirement of a (spacelike) interpolating surface that connects the region of interest and the extremal surface. We investigate to what extent this constraint is upheld by the generalized gravitational entropy argument, which relies on constructing replica symmetric q -fold covering spaces of the bulk, branched at the extremal surface. We prove (at the level of topology) that the putative extremal surface satisfies the homology constraint if and only if the corresponding branched cover can be constructed for every positive integer q . We give simple examples to show that homology can be violated if the cover exists for some values q but not others, along with some other issues.

KEYWORDS: Gauge-gravity correspondence, AdS-CFT Correspondence

ARXIV EPRINT: [1412.7561](https://arxiv.org/abs/1412.7561)

Contents

1	Introduction	1
2	Review of holographic entanglement	3
2.1	The RT and LM constructions	4
2.2	A question of homology	6
3	Exemplifying replica symmetric homology violation	8
3.1	A torus with a crosscap	8
3.2	Implications for the homology constraint	11
3.3	Further examples	11
4	A topological condition on q-Rényi saddles	13
4.1	Boundary conditions for branched covers	13
4.2	Stronger criterion from the cosmic brane construction	14
4.3	Another look at the boundary condition	14
5	Relation between topological consistency and homology constraint	15
5.1	General strategy	15
5.2	Example: BTZ black hole	17
5.3	Topological consistency is equivalent to homology constraint	18
5.4	Non-orientable manifolds	21
6	Discussion	22
A	Equivalence of topological consistency conditions	24
B	Review of algebraic topology	26

1 Introduction

Holography provides an intriguing connection between quantum information and geometry, in part inspired by the geometrization of quantum entanglement by the Ryu-Takayanagi (RT) proposal [1, 2] and its covariant generalization [3] (HRT). These proposals identify a particular bulk codimension-2 surface as the geometric encoder of the entanglement structure in the dual field theory. Whilst the original proposals owed their origins to analogies with black hole entropy and covariant entropy bounds, we now have a derivation of the RT formula (in situations with time reflection symmetry) from a gravitational path integral courtesy of Lewkowycz and Maldacena (LM) [4].

This may at first seem like a complete story for the time-reflection symmetric states of a holographic QFT. Although it requires certain assumptions, the LM construction gives a first principles derivation of how the minimal surface of the RT proposal comes about from a gravitational path integral. But while the LM construction captures the dynamical part of the RT conjecture, to our knowledge it has yet to be established whether the quantum gravity path integral employed by LM is cognizant of the topological constraints that must be imposed on the RT surface. This is the primary question that will concern us in this paper.

To appreciate the issues involved, recall that RT (HRT) proposes that the holographic entanglement entropy of a given region \mathcal{A} of the field theory is given by the area of a minimal (extremal) bulk surface \mathcal{E} anchored on the boundary $\partial\mathcal{A}$ of \mathcal{A} , with \mathcal{A} , $\partial\mathcal{A}$ considered to lie in the boundary of the bulk spacetime. The extremal surface \mathcal{E} is required to be homologous to the region \mathcal{A} in question [5, 6]. The cleanest phrasing of this statement to our knowledge appears in [7, 8] for the RT and HRT proposals respectively. One requires that there be a spacelike codimension-1 interpolating homology surface $\mathcal{R}_{\mathcal{A}}$ whose only boundaries are \mathcal{E} and \mathcal{A} . Without this constraint, the holographic formula would be at odds with known quantum mechanical properties of entanglement entropy, such as strong subadditivity; the topological prescription has been motivated as a natural, but *ad hoc* way of ensuring consistency. The simplest rationale for the homology constraint is that, in its absence, a subsystem \mathcal{A} and its complement \mathcal{A}^c can end up having the same entanglement even when the state of the system is impure (e.g., in a black hole geometry).

One can intuitively motivate this picture by realizing that one must introduce a cut for fields along \mathcal{A} when computing matrix elements of the reduced density matrix $\rho_{\mathcal{A}}$ via the path integral. In particular, the replica construction for computing Rényi entropies requires that the operators of the QFT are cyclically permuted when one crosses the cut. Since local QFT operators are the boundary values of bulk fields, one expects the cut to extend into the bulk as well. The boundary of the bulk branching surface $\mathcal{R}_{\mathcal{A}}$ is formed from \mathcal{E} and \mathcal{A} so that the homology constraint is satisfied; see figure 1 for an illustration. Existence of the homology surface $\mathcal{R}_{\mathcal{A}}$ ensures that fields are appropriately branched and one can treat the bulk geometry itself as a branched cover over some fundamental domain.

Based on the above arguments, one might imagine that the converse holds true trivially at the level of topology. That is, given a replica-symmetric bulk saddle with the correct boundary conditions, might we be guaranteed to find a codimension-2 defect and a codimension-1 interpolating surface implementing the homology constraint? Surprisingly, this turns out not to be true. One can construct bulk geometries with the requisite replica symmetry but which nevertheless do not admit an appropriate branching surface $\mathcal{R}_{\mathcal{A}}$. While it is plausible that such geometries are never dominant in the bulk path integral, their presence begs the question “under what conditions does the LM construction give rise to the homology constraint?” We will address this question in some detail below. A simplified version of our final statement is that, as long as one has a family of geometries parameterized by a real parameter q which for all positive integer values of q admits a branched cover description, then we recover the homology constraint from the LM construction.

Before beginning, we pause to dispel a potential confusion. The reader may well ask whether the above condition coincides with the assumptions actually made by Lewkowycz and Maldacena in their original paper [4]. The potential confusion arises from the fact that the phrasing of the assumptions in [4] is subject to at least two different interpretations, which are naturally termed ‘global’ and ‘local’. Indeed, the present authors do not agree among themselves as to which interpretation best fits the words written in [4]. Under the global interpretation, LM assumes that the replica geometries may be analytically continued to $q \sim 1$ with global topology outside the conical singularity given by a trivial \mathbf{S}^1 bundle $\mathbf{S}^1 \times X$ for some X . This global assumption coincides with the intuitive discussion above and immediately implies the homology constraint, as the homology surface is given by any global section of this trivial bundle.

However, the global assumption is rather stronger than one might like. As mentioned above, there is no a-priori reason why a (q -fold quotient of a) given replica-symmetric solution at integer q should have this structure. Furthermore, as may be readily seen by noting that there is no obstacle to constructing bulk solutions with small conical singularities on homology-violating surfaces, the global product structure is not required by the assumption that the $q \rightarrow 1$ limit of the bulk solutions be described by small conical singularities. Such considerations motivates us to focus on the alternate and weaker local interpretation of the LM assumptions, which does not determine the global structure and imposes the $\mathbf{S}^1 \times X$ structure only locally along the conical singularity. The implications are two-fold. First, as is necessary to allow LM to work in topologically nontrivial examples (such as computing the entropy of an interval on a torus), the \mathbf{S}^1 fiber need not be defined far from the singularity. Second, even the region near-but-outside the singularity is allowed to be a non-trivial \mathbf{S}^1 bundle over some X . Our main result can now be restated as saying that this local interpretation of the LM assumptions *also* implies the homology constraint so long as the family of geometries exists at all $q \geq 1$ and is an appropriate \mathbb{Z}_q quotient of a smooth q -fold replica at all integer q .

The outline of the paper is as follows: in section 2 we give a rapid overview of the concepts we need from QFT and holography vis-à-vis entanglement considerations. We also formulate the above issue concerning the homology constraint in detail. In section 3 we present counterexamples to the naïve intuition that the homology constraint is an automatic consequence of existence of a branched cover. In section 4 we formulate an essential topological consistency condition for branched covers as constructed by LM. This condition is a statement about consistency of the field theory replica trick with the existence of bulk branched covers at all integer values of q . We then illustrate and prove in section 5 that this topological consistency condition is in fact equivalent to the homology constraint on the RT prescription. We conclude with a discussion in section 6. Further technical details and a review of algebraic topology relevant for our proof can be found in the appendices.

2 Review of holographic entanglement

To set the stage for our discussion, let us quickly review the salient features of the RT/HRT and LM constructions that will play a role below. Consider a holographic d -dimensional QFT with a bulk gravity dual in asymptotically $d + 1$ dimensional AdS spacetime. We will

Boundary regions		
Symbol	Description	Dimension
\mathcal{B}	full boundary manifold	d
Σ_t	fixed-time slice	$d - 1$
\mathcal{A}	subregion of Σ_t	$d - 1$
$\partial\mathcal{A}$	entangling surface	$d - 2$
$\tilde{\mathcal{B}}_q$	q -fold branched cover of \mathcal{B} used in the replica trick	d
Bulk regions		
Symbol	Description	Dimension
\mathcal{M}	full bulk manifold (with $\partial\mathcal{M} = \mathcal{B}$)	$d + 1$
\mathcal{E}	RT minimal surface (with $\partial\mathcal{E} = \partial\mathcal{A}$)	$d - 1$
$\mathcal{R}_{\mathcal{A}}$	homology surface interpolating between \mathcal{E} and \mathcal{A}	d
$\tilde{\mathcal{M}}_q$	smooth bulk replica manifold (with singular $\partial\tilde{\mathcal{M}}_q = \tilde{\mathcal{B}}_q$)	$d + 1$
\mathcal{M}_q	fundamental domain of $\tilde{\mathcal{M}}_q$ when $\tilde{\mathcal{M}}_q$ is a branched cover (i.e. $\tilde{\mathcal{M}}_q/\mathbb{Z}_q$)	$d + 1$
\mathbf{e}_q	branching surface of the branched cover $\tilde{\mathcal{M}}_q \rightarrow \mathcal{M}_q$	varies
\mathbf{r}_q	homology surface interpolating between \mathbf{e}_q and \mathcal{A} (if it exists)	d

Table 1. Definition of boundary and bulk regions that we consider in the course of our discussion.

limit our discussion to situations where the QFT is planar, strongly coupled, and has a suitable gap in its spectrum, so that we can regard the bulk as a two derivative effective field theory which we take to be Einstein-Hilbert gravity coupled to matter degrees of freedom.¹

2.1 The RT and LM constructions

The QFT resides on a background geometry \mathcal{B} foliated by Cauchy surfaces Σ_t .² The dual gravitational background to this field theory is \mathcal{M} and has \mathcal{B} as its conformal boundary. Specifically, we will take \mathcal{M} (and all other bulk regions) to denote the conformal compactification, so that \mathcal{M} is a manifold with boundary: $\partial\mathcal{M} = \mathcal{B}$. See table 1 for an overview of our notation.

We are interested in computing the entanglement entropy in such a holographic QFT for a region $\mathcal{A} \subset \Sigma_t$ of the boundary geometry. The HRT prescription for computing holographic entanglement entropy requires us to consider a bulk codimension-2 extremal surface $\mathcal{E} \subset \mathcal{M}$ anchored on $\partial\mathcal{A}$ and takes the area of \mathcal{E} to give the boundary entanglement, i.e.,

$$S_{\mathcal{A}} = \frac{\text{Area}(\mathcal{E})}{4G_N}, \quad \partial\mathcal{E} = \mathcal{E} \cap \mathcal{B} = \partial\mathcal{A}. \quad (2.1)$$

¹The discussion generalizes to other planar QFTs dual to (classical) higher derivative gravitational theories following [9, 10], for which we would be evaluating some other local functional on a codimension-2 surface \mathcal{E} . As we primarily focus in on topological aspects, much of what we describe later will go through unmodified.

²We often use language appropriate for Lorentzian spacetimes despite focusing on Euclidean geometries. We will be interested in surfaces $\Sigma_{t=0}$ which sit at a moment of \mathbb{Z}_2 time-reflection symmetry (allowing thereby translation between the two cases by a suitable Wick rotation).

Crucially the bulk extremal surface is required to satisfy a *homology constraint* originally motivated in [6]. Usually this is stated as the requirement that the extremal surface must be homologous to the region \mathcal{A} . More precisely, following e.g., [7, 8] we will take this to mean that there exists a bulk codimension-1 spacelike surface $\mathcal{R}_{\mathcal{A}}$ which interpolates between the extremal surface and the boundary region of interest. To wit,

$$\exists \mathcal{R}_{\mathcal{A}} \subset \mathcal{M} : \quad \partial \mathcal{R}_{\mathcal{A}} = \mathcal{E} \cup \mathcal{A}. \tag{2.2}$$

In the RT/HRT constructions, taking \mathcal{A} to be spacelike is an additional restriction on the allowed minimal/extremal surfaces, though it is naturally incorporated in the maximin proposal of [11]. We will have much more to say about this constraint below.

We need one more ingredient to make contact with the LM path integral derivation. This ingredient is the replica trick for computing powers of the reduced density matrix $\rho_{\mathcal{A}}$ whose von Neumann entropy is the entanglement entropy under discussion. We define

$$\rho_{\mathcal{A}} = \text{Tr}_{\mathcal{A}^c} \rho, \quad S_{\mathcal{A}} = -\text{Tr}(\rho_{\mathcal{A}} \log \rho_{\mathcal{A}}), \tag{2.3}$$

where ρ is the total density matrix on $\Sigma_t = \mathcal{A} \cup \mathcal{A}^c$. After setting $t = 0$ as in footnote 2 and passing to Euclidean signature, ρ can be viewed as a state prepared by some path integral over \mathcal{B} (now denoting a Euclidean boundary) cut along $\Sigma_{t=0}$ where boundary conditions are imposed on fields at $\Sigma_{t=0}$ to compute particular matrix elements of ρ . Up to a normalizing factor, the trace over \mathcal{A}^c to obtain $\rho_{\mathcal{A}}$ is implemented by sewing up the part of the cut along \mathcal{A}^c , leaving \mathcal{B} with a cut only along \mathcal{A} . The replica construction in the QFT then proceeds by sewing q copies together cyclically along the cuts at \mathcal{A} to construct a singular manifold $\tilde{\mathcal{B}}_q$ whose partition function computes $\text{Tr}(\rho_{\mathcal{A}}^q)$. This then allows us to recover the q^{th} Rényi entropy of the QFT via:

$$S_{\mathcal{A}}^{(q)} = \frac{1}{1-q} \log \text{Tr}(\rho_{\mathcal{A}}^q) = \frac{1}{1-q} \log \frac{Z_q}{Z_1^q}, \tag{2.4}$$

where Z_q is the partition function of the QFT on $\tilde{\mathcal{B}}_q$ and Z_1 that on $\mathcal{B}_1 \equiv \mathcal{B}$. The entanglement entropy $S_{\mathcal{A}}$ of (2.3) is recovered in the limit $q \rightarrow 1$.

The LM construction first implements this computation of Rényi entropies holographically by extending the replica trick into the bulk. It then extracts the entanglement entropy as above by giving a geometric implementation of the continuation to non-integer q .

To compute Z_q , one proceeds by obtaining a bulk manifold $\tilde{\mathcal{M}}_q$ with boundary $\tilde{\mathcal{B}}_q$ (for some explicit examples see [12, 13]); as always, the partition function is simply given by the on-shell gravitational action computed on this geometry. This bulk computation follows the usual rules of Euclidean quantum gravity and the LM saddle point analysis remains valid when we analytically continue $q \in \mathbb{Z}_+ \mapsto q \in \mathbb{R}$, as long as $(q-1) \ell_{\text{AdS}}/\ell_{\text{Planck}} \gg 1$.

To get from the replica spacetimes $\tilde{\mathcal{M}}_q$ of LM to the RT minimal surface one proceeds as follows. Since the boundary geometry $\tilde{\mathcal{B}}_q$ is a q -fold cover over \mathcal{B} branched at $\partial \mathcal{A}$ with cyclic \mathbb{Z}_q symmetry, we can restrict attention to a single ‘fundamental domain’ by focusing on the quotient spacetime $\tilde{\mathcal{B}}_q/\mathbb{Z}_q$; this is just a copy of \mathcal{B} itself. Assuming the

bulk saddle point geometry $\tilde{\mathcal{M}}_q$ to respect replica symmetry,³ we may similarly consider the bulk quotient $\mathcal{M}_q = \tilde{\mathcal{M}}_q/\mathbb{Z}_q$. LM focus on the case where the action of \mathbb{Z}_q on $\tilde{\mathcal{M}}_q$ has a codimension-2 fixed point set \mathbf{e}_q with boundary $\partial\mathcal{A}$. This \mathbf{e}_q is to be identified as the progenitor of the extremal surface \mathcal{E} . It is assumed to result in a conical defect of angle $\frac{2\pi}{q}$ in \mathcal{M}_q . The desired QFT partition function Z_q on $\tilde{\mathcal{B}}_q$ is then q times the bulk action on \mathcal{M}_q , computed without a contribution from the conical defect.

The point of considering this quotient is that it allows continuation to arbitrary real values of q . The protocol is to find a geometry \mathcal{M}_q with boundary \mathcal{B} and from which a ‘singular’ codimension-2 surface \mathbf{e}_q ending at $\partial\mathcal{A}$ has been excised. One then imposes as a further boundary condition that \mathbf{e}_q is a conical defect of opening angle $\frac{2\pi}{q}$. The geometry is fixed by minimizing the action subject to this requirement, with no contribution to the action from the singularity. In the $q \rightarrow 1$ limit we require $\mathcal{M}_q \rightarrow \mathcal{M}$. The defect surface \mathbf{e}_q then becomes the minimal area surface \mathcal{E} in the Euclidean geometry. One can furthermore argue that the contribution to Z_q localizes on this surface, giving a correction to the action proportional to the area, in such a way that the area of the extremal surface computes the entanglement entropy.

2.2 A question of homology

The LM construction shows that the computation of Rényi entropies for arbitrary positive real q can be performed by finding geometries \mathcal{M}_q with boundary \mathcal{B} , but with a conical defect of angle $\frac{2\pi}{q}$. At integer values of q , via $\mathcal{M}_q = \tilde{\mathcal{M}}_q/\mathbb{Z}_q$ this should be related to a nonsingular replicated bulk $\tilde{\mathcal{M}}_q$ whose boundary is the replica $\tilde{\mathcal{B}}_q$ (on which this \mathbb{Z}_q acts as the replica symmetry). When this is the case, we say that \mathcal{M}_q lifts to a q -fold branched cover. Of primary interest to us is the relationship between the homology constraint on the one hand, and this lifting of the singular bulk geometry to a nonsingular replicated bulk on the other.

The bulk conical defect, coming from the fixed point set of the replica symmetry, is the codimension-2 surface \mathbf{e}_q anchored on the boundary at $\partial\mathcal{A}$. On the boundary $\tilde{\mathcal{B}}_q$, upon traversing a small loop around $\partial\mathcal{A}$, one passes through \mathcal{A} , and goes from one copy of \mathcal{B} to the next. It is tempting to imagine a natural picture of the full geometry that arises from continuing this reasoning into the bulk as follows: at the level of topology, $\tilde{\mathcal{M}}_q$ is formed by sewing together q copies of \mathcal{M}_q along some \mathbf{e}_q such that traversing a small loop around \mathbf{e}_q also results in a change of sheet in the cover $\tilde{\mathcal{M}}_q$ of \mathcal{M}_q . We shall investigate the correctness of this picture below.

This picture is straightforward in cases where there is a codimension-1 interpolating surface \mathbf{r}_q bounded by $\mathbf{e}_q \cup \mathcal{A}$: in other words, if the homology constraint is satisfied by the conical defect. Cutting along \mathbf{r}_q , and gluing together q copies cyclically, just as on the boundary, builds the covering space $\tilde{\mathcal{M}}_q$ at the level of topology. Passing through \mathbf{r}_q

³The \mathbb{Z}_2 time-reflection symmetry about $t = 0$ of the state ρ intertwines with the cyclic \mathbb{Z}_q symmetry of the replica construction, to give a larger dihedral symmetry group \mathbb{D}_q ; see [14] for its relevance in computing Rényi entropies. We refrain from utilizing the full dihedral symmetry, allowing for the possibility that the LM construction gives a surface that does not lie at $t = 0$ in the bulk. Therefore, in what follows, replica symmetry will always refer to the cyclic \mathbb{Z}_q group.




Figure 1. Replica construction in the boundary and bulk for $q = 3$. The replica symmetric q copies of the field theory on \mathcal{B} , form a q -fold branched cover $\tilde{\mathcal{B}}_q$ which fixes the asymptotic data for the bulk problem. The bulk covering spacetime $\tilde{\mathcal{M}}_q$ has a \mathbb{Z}_q symmetry with fixed point locus \mathbf{e}_q (shown as the wavy lines) anchored on $\partial\mathcal{A}$. Typically one also encounters via this construction a bulk interpolating surface \mathbf{r}_q (the light blue branching surface) in the bulk whose boundaries are \mathbf{e}_q and \mathcal{A} . Conventional intuition dictates that the bulk spacetimes are all covers over a single fundamental domain (one of the components in the picture) branched over the codimension-1 surface \mathbf{r}_q . Passing through this surface cycles through the sheets of the bulk in a fashion identical to passage through \mathcal{A} . The homology condition posits that such an \mathbf{r}_q exists. We argue that this picture is accurate as long as we are suitably careful with the notion of allowed branched covers. As $q \rightarrow 1$, $\mathbf{r}_q \rightarrow \mathcal{R}_{\mathcal{A}}$ and $\mathbf{e}_q \rightarrow \mathcal{E}$.

will cause a change of sheet in the cover, just as passing through \mathcal{A} changes sheets on the boundary. So any \mathcal{E} obeying the homology constraint⁴ naturally leads to a family of bulk saddles $\tilde{\mathcal{M}}_q$ obeying the correct boundary conditions. We may say that \mathcal{M} lifts to $\tilde{\mathcal{M}}_q$, with \mathcal{E} lifting to \mathbf{e}_q .⁵ See figure 1 for an illustration of this scenario.

However, the converse, or at least the strongest converse one might propose, fails to be true. More specifically, the following three statements will be demonstrated in sections 3, 4, and 5.

- (i). The existence of a two-sided interpolating codimension-1 surface \mathbf{r}_q implies that \mathcal{M}_q lifts to a branched cover (homology \implies lift). This is what we have informally argued above.

⁴The natural homology constraint to assume is the one in Lorentz signature. We will take the Lorentz signature spacetime to be time-orientable. It then admits a nowhere-vanishing vector field along which $\mathcal{R}_{\mathcal{A}}$ can be deformed until it lies in the moment of time symmetry. It follows that the homology constraint also holds in Euclidean signature. For future reference, we also note that Lorentz-signature time-orientability makes $\mathcal{R}_{\mathcal{A}}$ two-sided, i.e., it has a continuous unit normal. In case of an orientable bulk, two-sidedness is equivalent to orientability of $\mathcal{R}_{\mathcal{A}}$.

⁵When all surfaces of interest lie on the $t = 0$ slice, $\mathcal{R}_{\mathcal{A}}$ is uniquely determined by \mathcal{E} and \mathcal{A} as long as that slice has no closed, boundaryless components. More generally, $\mathcal{R}_{\mathcal{A}}$ is unique if $H_d(\mathcal{M}) = 0$, which we expect to hold in most physical situations (see, for example [15]). If $H_d(\mathcal{M}) \neq 0$ then there is an ambiguity, though it will not make a difference to the entanglement entropy in the classical limit (but it will matter for Rényi entropies and for quantum corrections).

- (ii). There are branched covers with the correct boundary conditions which cannot be realized in this way. That is, for given $q \in \mathbb{Z}_+$, there can exist an \mathcal{M}_q formed from a quotient of a branched cover which does not admit an interpolating surface \mathbf{r}_q between \mathcal{A} and the fixed point set \mathbf{e}_q (lift $\not\Rightarrow$ homology).
- (iii). However, given a continuous family of bulk geometries \mathcal{M}_q parameterized by real q which for every $q \in \mathbb{Z}_+$ lifts to a q -fold branched cover, we will show that each \mathcal{M}_q admits an interpolating surface \mathbf{r}_q . Taking $q \rightarrow 1$ then shows that \mathcal{E} satisfies the homology constraint as desired (lift $\forall q \implies$ homology).

We also note that, on top of this, the fixed point set arising from a \mathbb{Z}_q quotient may give rise to something other than a $\frac{2\pi}{q}$ conical defect. It is possible to generate fixed point sets with the wrong codimension, and also to engineer situations wherein the codimension-2 fixed point set has an incorrect defect angle. In section 3 we give examples where both these scenarios can be realized.

Below we generally confine ourselves to topological arguments. In particular, we refrain from employing dynamical information from the path integral to constrain \mathcal{M}_q . This is in part due to the fact that classification of all replica invariant saddles with given boundary geometry is a notoriously hard problem (even for $d = 2$).⁶ In the same spirit, the notion of replica symmetry should always be understood in a topological sense. Formally, the bulk will be a multi-sheeted surface — we refrain from specifying the detailed geometry on the sheets, but do keep track of what the boundary conditions of the gravitational path integral imply for moving across the sheets (see section 4).

3 Exemplifying replica symmetric homology violation

We begin by considering some simple examples, focusing only on the topology, which illustrates some of the unexpected features that may be encountered.

3.1 A torus with a crosscap

Suppose we wish to compute the entropy of the thermal state of a two-dimensional field theory living on a circle (or equivalently the entanglement entropy between the two halves of the thermofield double). The state is defined by a path integral over a cylinder of length $\beta = 1/T$. The q^{th} Rényi entropy can be computed by the partition function on q copies of the cylinder glued cyclically, which gives a torus $\tilde{\mathcal{B}}_q$ of length $q\beta$. The replica symmetry is implemented by rotating through β in the Euclidean time direction. The region \mathcal{A} of interest becomes a spatial circle at fixed Euclidean time, with empty boundary.

To compute the partition function holographically, one must choose how the torus is to be ‘filled in’ to obtain a bulk $\tilde{\mathcal{M}}_q$. One option is to fill the spatial circle with a disk, like thermal AdS, in which case there are no fixed points under the replica symmetry, so the

⁶The kinematical aspect of our analysis is reminiscent of the first attempt to prove the RT proposal in [5] where the homology condition was introduced. But while [5] argued that the bulk should be a branched cover satisfying homology, we would instead like to ascertain the conditions under which the homology condition becomes automatic. We thank Matt Headrick for a discussion on this issue.

quotient \mathcal{M}_q is smooth, and the entanglement entropy obtained would vanish. The region \mathcal{A} is contractible in the bulk, so this is consistent with the homology constraint.

If, on the other hand, the Euclidean time circle is filled in with a disk, as in the BTZ black hole geometry, \mathcal{A} is not contractible. But the centre of the disk is fixed under the rotation implementing replica symmetry. So it gives rise to a circle e_q of fixed points and a conical defect in the quotient \mathcal{M}_q . This defect becomes the bifurcation circle of the event horizon in the $q \rightarrow 1$ limit. There is an obvious interpolating surface joining it to \mathcal{A} on the boundary, so again the homology constraint is obeyed.

However, ignoring for now the equations of motion, there are more topologies that might in principle be allowed. One is to fill the Euclidean time circle not with a disk, but using a cross-cap. (In what follows, the spatial circle will play little role, so we focus on some 2-dimensional slice corresponding to a point on this circle.) This means that we first fill it with an annulus, taking the outside edge of the annulus to be the boundary, but we then identify antipodal points of the inside edge so that passing through the inside edge causes one to jump to the point directly opposite. The replica symmetry can be implemented by rotation of the cross-cap in the obvious manner. This gives a space $\tilde{\mathcal{M}}_q$ with the topology of a Möbius strip times the spatial circle. The (single) edge of the Möbius strip is the boundary Euclidean time circle.

There are now two cases, depending on whether q is odd or even, both of which are illustrated in figure 2. We begin with q odd. In this case, there are no fixed points of the replica symmetry, and the quotient \mathcal{M}_q is a smooth geometry with the same topology as $\tilde{\mathcal{M}}_q$. But \mathcal{A} has non-trivial homology in \mathcal{M}_q : there is no surface in the bulk whose only boundary is a spatial circle. In this way, the homology constraint is violated because the fixed point set of the replica symmetry (being empty here) is not homologous to the boundary region.

The second case occurs for even q . While no surface is fixed by every element of the replica group, the $q/2$ replica symmetry now rotates half way round the time circle. On the inner edge of the annulus where points are identified with their antipodes, this symmetry thus takes every point to itself. In the quotient \mathcal{M}_q , the resulting singular set is the inside circle of the annulus times the spatial circle. This set is not codimension-2, but instead codimension-1; it is a \mathbb{Z}_2 orbifold plane. Thus, for even q there at least exists a fixed point set under a \mathbb{Z}_2 subgroup of \mathbb{Z}_q , despite the fixed points of the full replica group \mathbb{Z}_q being empty for $q > 2$.

There are several objections that one might raise to this example. The most obvious is that a metric is never constructed, and that there can never be any genuine saddle points with the given topology. The clearest refutation of this is to give an example sharing all the same qualitative features, but with a metric. One can in fact give such a construction for pure three dimensional gravity. To explain the idea, consider the example given instead by starting from a geometry with two disjoint torus asymptotic boundaries, and then taking a \mathbb{Z}_2 quotient by swapping the two tori and simultaneously rotating half way round the time circle. This gives a cross-cap in the time circle as described here. Now generalize it to start with not tori, but higher genus Riemann surfaces with negative Euler characteristic. There is an easy way to put a constant negative curvature metric on these geometries, as in the

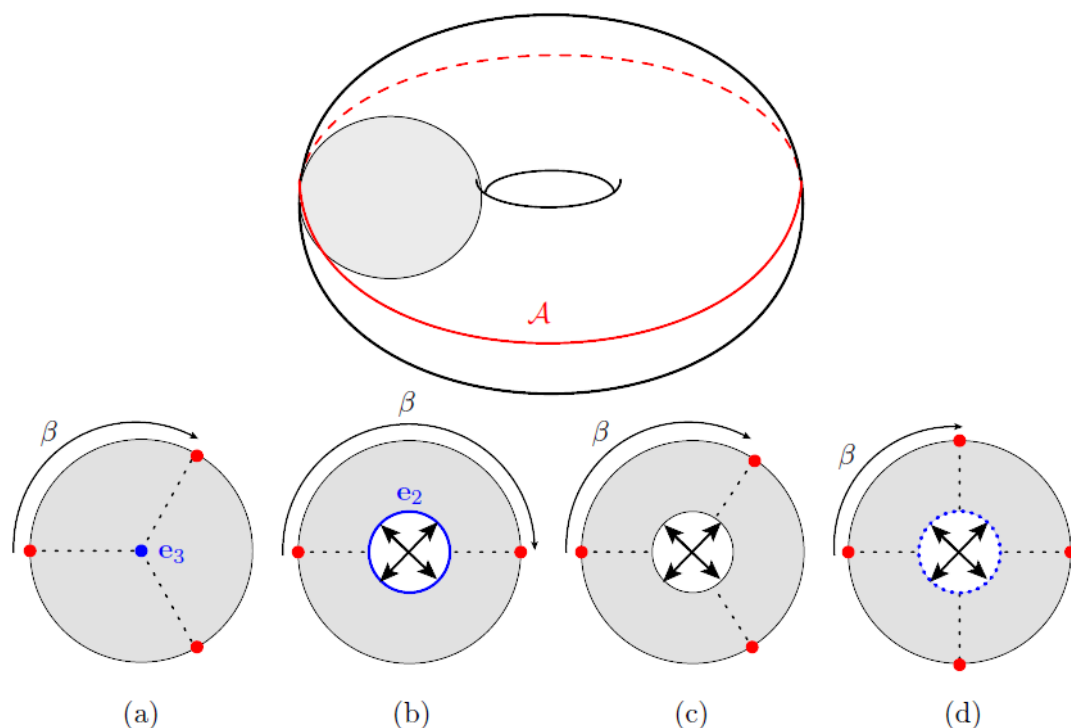


Figure 2. Different ways of filling the boundary torus. The replica \mathbb{Z}_q symmetry acts as a rotation along the Euclidean time direction by β . Case (a) shows a slice of \mathcal{M}_3 after filling the boundary time circle with a disk; the replica fixed point set e_3 (blue) is a codimension-2 surface in the centre of the circle. Cases (b), (c) and (d) show slices of $\tilde{\mathcal{M}}_q$ for $q = 2, 3, 4$, respectively, after filling the boundary torus with a cross cap. For $q = 2$ the fixed point set e_2 is a codimension-1 orbifold plane wrapping the cross cap. For $q \geq 3$ there is no fixed point set under \mathbb{Z}_q (i.e. e_q is empty) and the homology constraint is violated. However, we indicate in (d) that for even q the cross cap itself is still a fixed point set under the subgroup \mathbb{Z}_2 of rotations by $\frac{q}{2}\beta$, resulting in an orbifold plane under the quotient.

Maldacena-Maoz wormhole [16]. If the Riemann surface has a fixed-point free involutive isometry, the \mathbb{Z}_2 quotient of this combined with swapping the two boundaries gives the relevant example.

A second objection is that this fails to satisfy the homology constraint in a very particular way, by having no fixed points under the replica symmetry. This is impossible if the region \mathcal{A} has a nonempty boundary, since there must be a set of fixed points extending from $\partial\mathcal{A}$ into the bulk. We will show later in this section that it is possible to circumvent this objection and construct an example with \mathcal{E} nonempty, but not homologous to \mathcal{A} .

Finally, these geometries might never be dominant saddles in the path integral. While this is plausible, it is very difficult to see how their dominance can be ruled out in general. For this reason, we would like to make arguments that apply using only the topology.

3.2 Implications for the homology constraint

We believe this example has a real lesson to teach us: the homology constraint follows from the holographic replica trick only when the defect arises from a \mathbb{Z}_q quotient *at every positive integer* q . In the context of the LM argument, one may start with a bulk geometry \mathcal{M} with boundary \mathcal{B} , and pick some extremal surface \mathcal{E} ending on $\partial\mathcal{A}$. Now introduce a small conical defect along \mathcal{E} , and adjust the geometry so it remains on-shell away from the defect, which we now call \mathbf{e}_q . The angle of the conical singularity can be dialed to $\frac{2\pi}{q}$ for real positive values of q , to give a family of singular bulks \mathcal{M}_q , returning to \mathcal{M} when $q \rightarrow 1$. We assume that this can be done without changing the topology of the bulk, or the defect within it. When q hits an integer, we can make a connection with the replica trick, but only if \mathcal{M}_q can be lifted to a branched cover. That is, we require that there exists some $\tilde{\mathcal{M}}_q$, with boundary $\tilde{\mathcal{B}}_q$, and with a \mathbb{Z}_q symmetry such that the quotient yields \mathcal{M}_q . Under this condition the action evaluated on $\tilde{\mathcal{M}}_q$ will legitimately give the q^{th} Rényi entropy (assuming $\tilde{\mathcal{M}}_q$ is the dominant saddle).

The examples tell us that this may happen at some q despite \mathcal{E} violating the homology constraint. But taking \mathcal{M} as the cross-cap geometry above, with \mathcal{E} chosen to be empty, this lift to $\tilde{\mathcal{M}}_q$ can be constructed only at odd q , and does not exist for even q . This is because in the lift we must choose which copy of the boundary to end up on after passing through the cross-cap, say by rotating through k copies. But going through the cross-cap twice is homotopic to passing round the boundary circle, so must take us through one copy. This means we require $1 = 2k \pmod q$, which has a solution if and only if q is odd, so no choice of k gives the correct cover on the boundary. We argue below that demanding the stronger condition that the lift must exist for all q so that \mathcal{M}_q can be used to compute $S_{\mathcal{A}}^{(q)}$ at every positive integer, is equivalent to the homology constraint on \mathcal{E} .

This example generalizes to allow nonempty \mathcal{E} as follows. Take the bulk \mathcal{M} to be the same geometry as above, a solid torus with a neighborhood of a loop round the torus cut out and replaced with a Möbius strip times a circle. Consider now the region \mathcal{A} as not a whole boundary circle, but an interval. Choose the surface \mathcal{E} to be a curve joining the endpoints of \mathcal{A} , but passing round the non-contractible spatial circle on the side of the torus opposite of \mathcal{A} , and hence not homologous to it. Nonetheless, again at odd values of q this lifts to a smooth replica-symmetric bulk $\tilde{\mathcal{M}}_q$ in much the same way. When going round the surface \mathcal{E} , move from sheet to sheet as usual, but when passing through the crosscap pass to the $(\frac{q+1}{2})^{\text{th}}$ sheet relative to where you begin. In essence, the crosscap acts to replace the closed geodesic that would otherwise be required, but without the associated fixed point set or conical defect.

3.3 Further examples

There are several aspects of the above examples that may appear to be required to get the sort of homology violation we see: for example, a non-orientable spacetime, a distinction only between odd and even q , or the specific way a topological feature replaces a part of the spacetime where a fixed point ‘should be’. We now briefly describe an example to show that none of these are necessary. Take \mathcal{M} as a 3-dimensional ball, with boundary $\mathcal{B} = S^2$,

cut out a ball at the center, and then identify antipodal points of the spherical edge of the resulting hole. This results in an \mathbb{RP}^3 with a boundary, an orientable spacetime. Take \mathcal{A} as an interval on the boundary sphere, running between the poles perhaps, and choose \mathcal{E} to pass through the nontrivial topology we have introduced in the centre. This can be lifted to a replica-symmetric cover⁷ to compute Rényi entropies for odd q as above, but not for even q . A generalization identifies the points on the sphere bounding the cut out region as in the construction of a Lens space $L(p; p')$ from identifications of the boundary of a 3-ball (the special case of antipodal points being $p = 2$). The result of this is that a lift will exist for q co-prime to p and not otherwise.

We conclude this section with a different sort of example, which shows that it is possible to obtain a codimension-2 defect with deficit angle different from $\frac{2\pi}{q}$ from the quotient construction. The simplest example generalizes the computation of thermal entropy from BTZ; again we will not explicitly mention the trivial dependence on the spatial direction and focus on a constant spatial slice. The usual geometry then picks \mathcal{E} as a point in the centre of this slice, and passing anticlockwise round this point takes you up one sheet in the cover. To generalize this, choose \mathcal{E} as a collection of several points, and for each point choose some integer number of sheets to change by when passing round them anticlockwise. To be consistent with the boundary covering space, the only requirement is that these integers sum to one. But if the integer chosen for a point, say n is anything other than ± 1 , the quotient of the covers thus constructed will result in a variety of conical deficit angles, depending on the greatest common factor r of n and q . Specifically, the covering space will contain r copies of the point, related by the \mathbb{Z}_r subgroup of the \mathbb{Z}_q replica symmetry, and the resulting defect angle after the quotient will be $\frac{2\pi r}{q}$.

For a specific example, take two points, choosing to increase by two sheets when passing round the first, and to decrease by one when passing round the second. The $q = 4$ covering space has the topology of a torus with a single boundary (being the usual four glued copies of the boundary circle). This can be understood as a regular octagon, with top and bottom edges identified, as well as left and right edges; the remaining diagonal edges form the boundary circle. The replica symmetry acts by rotations by $\frac{\pi}{2}$. This leaves the obvious fixed point at the centre, which becomes the usual $\frac{\pi}{2}$ deficit after the quotient. The replica symmetry rotating by π has an additional two fixed points, being the centre of the top or bottom edge (identified with one another), and the centre of the left or right edge. Under the quotient, these two points become identified, and the result is a conical defect of angle π .

These sorts of examples can be thought of as many degenerate extremal surfaces lying on top of one another, which is very physically natural when viewing the defects as cosmic strings, and can be recovered as a limit of several of the usual defects coalescing. We would not expect these configurations to dominate in simple examples, particularly in the $q \rightarrow 1$ limit, but we can not rule out the possibility that these surfaces with multiplicity may be favourable enough to dominate in some more complicated geometry.

⁷An easy way to see this is to notice that $\mathcal{M} - \mathcal{E}$ deformation retracts to a circle, and going twice round this circle (or p times in the generalization) is homotopic to a boundary loop passing through \mathcal{A} once.

4 A topological condition on q -Rényi saddles

In light of the example described in section 3, we now wish to formulate a refined requirement regarding existence of covering spaces consistent with replica symmetry, which will be sufficiently strong to impose the expected topological constraint on the branching surface.

4.1 Boundary conditions for branched covers

The replica trick on the boundary involves defining the CFT on a branched cover $\tilde{\mathcal{B}}_q$. To compute the semiclassical Rényi entropy, one should evaluate the gravitational action on every bulk manifold $\tilde{\mathcal{M}}_q$ satisfying the equations of motion and the boundary condition $\partial\tilde{\mathcal{M}}_q = \tilde{\mathcal{B}}_q$, and choose the dominant saddle. This is prohibitively difficult, so we ignore the question of which saddle dominates, and furthermore we restrict to $\tilde{\mathcal{M}}_q$ which can be realized as a branched cover over the original space \mathcal{M} . That is, we assume that $\mathcal{M}_q = \tilde{\mathcal{M}}_q/\mathbb{Z}_q$ is homeomorphic to the original manifold \mathcal{M} . The boundary condition $\partial\tilde{\mathcal{M}}_q = \tilde{\mathcal{B}}_q$ imposes a topological condition on $\mathcal{M}_q - \mathbf{e}_q$, which we have seen above is *not* the homology constraint. First we need to understand this condition on $\mathcal{M}_q - \mathbf{e}_q$ more precisely.

Let \mathcal{B} be the boundary, and \mathcal{A} the region whose entanglement entropy we want to compute. Consider arbitrary closed loops in $\mathcal{B} - \partial\mathcal{A}$, i.e., loops which may intersect \mathcal{A} , but not $\partial\mathcal{A}$. Then there is a homomorphism $\phi : \pi_1(\mathcal{B} - \partial\mathcal{A}) \rightarrow \mathbb{Z}$ which counts the number of times the loop passes through \mathcal{A} . This map is defined to take into account the orientation of the loop relative to \mathcal{A} , so it computes the signed intersection number. Hence for any given q there is a map $\phi_q : \pi_1(\mathcal{B} - \partial\mathcal{A}) \rightarrow \mathbb{Z}_q$ which is just the previous map modulo q . The replicated boundary can be found by taking the cover $\tilde{\mathcal{B}}_q - \partial\mathcal{A}$ of $\mathcal{B} - \partial\mathcal{A}$, defined so that loops in $\ker \phi_q$ lift to closed loops in the covering space. This is just saying that ϕ_q counts which sheet we are on in the replicated CFT, and loops which intersect q times with \mathcal{A} are to be identified with closed loops in the q -fold covering space in accordance with the replica trick.

Replica symmetry and the boundary condition $\partial\tilde{\mathcal{M}}_q = \tilde{\mathcal{B}}_q$ imply a similar construction for the bulk. First, it implies that there exists a map ψ_q keeping track of movements between sheets of the cover in the bulk. Second, it requires that this restricts to ϕ_q on the boundary $\partial(\mathcal{M}_q - \mathbf{e}_q) = \mathcal{B} - \partial\mathcal{A}$ to give the correct boundary conditions. This can be summarized as the existence of a bulk sheet-counting homomorphism ψ_q , factoring ϕ_q so that the following diagram commutes:

$$\begin{array}{ccc}
 \pi_1(\mathcal{B} - \partial\mathcal{A}) & \xrightarrow{\phi_q} & \mathbb{Z}_q \\
 \downarrow i_* & \nearrow \psi_q & \\
 \pi_1(\mathcal{M}_q - \mathbf{e}_q) & &
 \end{array}
 \tag{4.1}$$

Here i_* is the pushforward induced by the inclusion i of the boundary into the bulk. Below, we show more formally that the existence of such a ψ_q is equivalent to the existence of a replica-symmetric covering space of the bulk.

4.2 Stronger criterion from the cosmic brane construction

The condition (4.1) is true of any branched cover satisfying the boundary conditions required by AdS/CFT. The example in section 3 illustrates that with this condition alone, the singularity \mathbf{e}_q is not necessarily homologous to \mathcal{A} .

But now, let us we restrict to branched covers of the original manifold \mathcal{M} constructed following LM: choose a codimension-2 surface \mathcal{E} , and introduce a conical defect $\frac{2\pi}{q}$ at this surface. Can this be lifted to a branched cover $\tilde{\mathcal{M}}_q$ obeying the correct boundary conditions? In general, this is possible only if the boundary condition (4.1) holds for $\mathcal{M} - \mathcal{E}$. This criterion must be applied at each q separately, as in general it may be possible to find ψ_q for some values of q but not for others, as in the crosscap example in section 3. If ψ_q exists lifting ϕ_q for all q , then there is a lift⁸ $\psi : \pi_1(\mathcal{M} - \mathcal{E}) \rightarrow \mathbb{Z}$ of ϕ , meaning that the following diagram commutes:

$$\begin{array}{ccc}
 \pi_1(\mathcal{B} - \partial\mathcal{A}) & \xrightarrow{\phi} & \mathbb{Z} \\
 \downarrow i_* & \nearrow \psi & \\
 \pi_1(\mathcal{M} - \mathcal{E}) & &
 \end{array}
 \tag{4.2}$$

We will prove that (4.2) is equivalent to the homology condition for the RT surface.

To illustrate this, let us return to the crosscap example of section 3 in this language. In this case, the starting point \mathcal{M} is the Möbius strip times a circle. We choose the defect \mathcal{E} to be the empty set, and try to construct branched covers at integer q . This can be done in a way satisfying the boundary condition (4.1) for odd q , as illustrated in figure 2, but for even q , there is no way to construct the branched cover. This can be seen very explicitly: we have $\phi_q(\Gamma) = 1$, where Γ is the loop going once around the boundary time circle. Assume there exists a lift ψ_q and require $\psi_q(\Gamma) = 1$. Now try to find a consistent way of choosing a value under ψ_q for the bulk loop γ that goes once through the cross-cap: since Γ is homotopic to 2γ , we require $1 \pmod q = \psi_q(\Gamma) = \psi_q(2\gamma) = 2\psi_q(\gamma)$ which has a solution in \mathbb{Z}_q if and only if q is odd. At the level of the maps ϕ and ψ of (4.2), this contradiction for even values of q manifests itself as follows: both fundamental groups in (4.2) are \mathbb{Z} (ignoring the spatial circle), but one of them is generated by Γ , the other is generated by γ . This means ϕ is the identity (in particular $\phi(\Gamma) = 1$) but i_* is multiplication by 2, so clearly ψ does not exist.

4.3 Another look at the boundary condition

Above we motivated (4.1) from the AdS/CFT boundary conditions. We now discuss more formally how to show that, for a codimension-2 defect, this condition is equivalent to the existence of a topologically replica-symmetric covering space of the bulk.

We take here the perspective that a branched cover is defined as a covering space in the usual sense, where any point in the base space has a neighborhood lifting to q

⁸This is guaranteed if the relevant groups are finitely generated, which holds for compact manifolds with boundary.

homeomorphic copies of itself, but after removing the branching surface.⁹ Given this perspective, a replica-symmetric q -fold cover of \mathcal{M}_q , branched at \mathbf{e}_q , restricting correctly on the boundary, is defined by the following diagrams:

$$\begin{array}{ccc}
 \tilde{\mathcal{B}}_q - \partial\mathcal{A} & \xrightarrow{p} & \mathcal{B} - \partial\mathcal{A} & & \pi_1(\tilde{\mathcal{B}}_q - \partial\mathcal{A}) & \xrightarrow{p_*} & \pi_1(\mathcal{B} - \partial\mathcal{A}) & \xrightarrow{\phi_q} & \mathbb{Z}_q & \longrightarrow & 1 \\
 \downarrow \tilde{i} & & \downarrow i & & \downarrow \tilde{i}_* & & \downarrow i_* & & & & \\
 \tilde{\mathcal{M}}_q - \mathbf{e}_q & \xrightarrow{P} & \mathcal{M}_q - \mathbf{e}_q & & \pi_1(\tilde{\mathcal{M}}_q - \mathbf{e}_q) & \xrightarrow{P_*} & \pi_1(\mathcal{M}_q - \mathbf{e}_q) & \xrightarrow{\psi_q} & \mathbb{Z}_q & \longrightarrow & 1
 \end{array} \tag{4.3}$$

The left hand diagram gives the covering space for boundary (top row) and bulk (bottom row), with respective covering maps p and P . Then i and \tilde{i} are the injective inclusion maps of boundary into bulk, and the diagram commutes to implement the inclusions consistently. This implies immediately existence and commutativity of the left square of the diagram on the right. The additional maps are required to make the rows exact; the existence of ϕ_q and ψ_q is then equivalent to replica symmetry. This is because we require the q -fold cover to possess the replica symmetry, acting via a \mathbb{Z}_q deck transformation group, acting cyclically. This implies that the covering space is *normal*, so that $P_*(\pi_1(\tilde{\mathcal{M}}_q - \mathbf{e}_q))$ is a normal subgroup of $\pi_1(\mathcal{M}_q - \mathbf{e}_q)$, and taking the quotient by that subgroup gives the \mathbb{Z}_q deck transformation group. This shows that the map ψ exists, being the quotient map. The same argument also applies on the boundary. In particular $\ker(\phi_q) = \text{im}(p_*)$ and $\ker(\psi_q) = \text{im}(P_*)$. This shows that exactness is just a formal restatement of replica symmetry (in a topological sense) both at the boundary and in the bulk.

We prove in appendix A that the holographically natural consistency condition (4.1) holds if and only if $\mathcal{M}_q - \mathbf{e}_q$ is such that the diagrams (4.3) exist as described: existence of the map ψ_q to factor ϕ_q as $\psi_q \circ i_*$ implies that a cover can be constructed, and conversely existence of a cover implies that there is such a ψ_q .

5 Relation between topological consistency and homology constraint

Having dealt with the topological consistency condition required for the LM construction, we now turn to the central thesis of this work: “when is the homology condition satisfied?” We will give an overview of our general strategy first and exemplify it with the BTZ spacetime. In section 5.3 we prove Theorem 1 which posits that we are guaranteed the homology constraint provided the topological consistency condition is met (and vice-versa).

5.1 General strategy

The extension of the replica trick into the bulk [4] constructs a bulk geometry with boundary conditions given by the replicated field theory. At integer q this leads to a geometry

⁹This works when the branching surface is codimension-2 as expected, since then there is a uniquely specified way to put back the branching surface in the cover. It fails when the branching surface is codimension-1, since extra information on what happens when passing through the surface must be imposed, but as this can only happen for even q in any case, this will not alter our main conclusions. Similarly, there are examples with fixed point sets of codimension greater than 2, but the local topology of such a quotient depends on q , so that there is no obvious way to fit them into a family at all values of q .

which is a smooth q -fold covering space $\tilde{\mathcal{M}}_q$ of the original bulk geometry. The quotient $\mathcal{M}_q = \tilde{\mathcal{M}}_q/\mathbb{Z}_q$ is then just the original bulk geometry with a conical defect inserted along the codimension-2 fixed point set of \mathbb{Z}_q . By construction \mathcal{M}_q lifts to a q -fold branched cover of \mathcal{M} at every integer value of q and the topological consistency condition (4.2) holds. This topological consistency condition is the essential feature of the LM construction on which we will focus for the remainder of this section. In particular, we will no longer concern ourselves with covering spaces and simply restrict attention to topological properties of \mathcal{M} and \mathcal{E} .

For concreteness, consider the following setup. We start with a bulk geometry \mathcal{M} , a boundary region \mathcal{A} and a candidate extremal surface \mathcal{E} to be used in computing entanglement entropy of \mathcal{A} . The goal of the present section is to first illustrate and then prove the following statement: *The topological consistency condition is satisfied for all q if and only if \mathcal{A} is homologous (in the sense of (2.2)) to the surface \mathcal{E} .*¹⁰ In order to decide whether $\mathcal{M} - \mathcal{E}$ satisfies the topological consistency condition for all q (and thus whether inserting a conical defect $\frac{2\pi}{q}$ along \mathcal{E} would lift to a q -fold branched cover in the sense of LM at all integer q), consider the following maps:

- A *boundary sheet counting map* $\phi \in H^1(\mathcal{B} - \partial\mathcal{A})$: this is the map from (4.2) which counts how many times a boundary loop passes through \mathcal{A} . Although ϕ was originally defined on the homotopy group $\pi_1(\mathcal{B} - \partial\mathcal{A})$, the fact that it maps to an abelian group (\mathbb{Z}) makes its action on the first homology group well-defined. It is an element of the first cohomology group $H^1(\mathcal{B} - \partial\mathcal{A})$ of $\mathcal{B} - \partial\mathcal{A}$ with integer coefficients; such elements are just homomorphisms from boundary loops into \mathbb{Z} .
- A *local intersection map* $u \in H^2(\mathcal{M}, \mathcal{M} - \mathcal{E})$. By $H^2(\mathcal{M}, \mathcal{M} - \mathcal{E})$ we mean the second cohomology group of \mathcal{M} relative to $\mathcal{M} - \mathcal{E}$ with integer coefficients. This u is defined on 2-dimensional surfaces (2-chains) in a neighborhood of \mathcal{E} whose boundary is not a part of \mathcal{E} . Given such a 2-surface \mathcal{D} , the map u counts the (signed) number of intersections of \mathcal{E} with \mathcal{D} .¹¹ The sign of an intersection is given by the relative orientation of \mathcal{D} and \mathcal{E} .

Intuitively, ϕ (being defined by \mathcal{A}) encodes the boundary conditions for bulk geometries that can be lifted to replica symmetric covers. On the other hand, u (being defined by \mathcal{E}) carries topological information about the way such a bulk covering space would be branched. The topological consistency condition relates these two objects: namely, we can translate the topological consistency condition into a certain consistency of ϕ with respect to u . To this end, take any 2-surface \mathcal{D} in the bulk, which is anchored outside of $\partial\mathcal{A}$ on the boundary, i.e., $\partial\mathcal{D} \subset \mathcal{B} - \partial\mathcal{A}$. Now compare the following two properties of \mathcal{D} :

- Use ϕ to count how many times the loop $\partial\mathcal{D}$ passes through \mathcal{A} on the boundary (taking into account orientations).

¹⁰We postpone subtleties concerning orientation to section 5.4.

¹¹Note that this is only true in the absence of torsion cycles. If the bulk spacetime has torsion (in the topological sense), then the second cohomology group is not isomorphic to intersection counting homomorphisms. To illustrate the essence of our argument, we refrain from considering these subtleties at the present. We will however account for torsion in our proof of Theorem 1 in section 5.3.

- On the other hand, consider the parts of \mathcal{D} which lie in a tubular neighborhood of \mathcal{E} and compute their (signed) intersection number with \mathcal{E} using u .

Our main statements can be summarized as follows:

1. The topological consistency condition means that these two countings have to agree for any 2-surface \mathcal{D} anchored at the boundary.
2. The two countings agree for all 2-surfaces \mathcal{D} anchored at the boundary if and only if \mathcal{A} and \mathcal{E} are homologous.

Illustration and proof of these points will be our main task. The rationale behind the first point, is roughly the following; it is a reformulation of the consistency condition at the level of the unreplicated bulk geometry. In fact, the topological consistency condition as formulated in (4.2) means exactly that ϕ can be lifted to a global sheet counting map $\psi \in H^1(\mathcal{M} - \mathcal{E})$ for bulk loops in a way that is consistent with u on loops that bound 2-surfaces which intersect \mathcal{E} . Roughly speaking, 2-surfaces intersecting \mathcal{E} encode the same information about local branching near the intersection point as do small loops going around the intersection point. We thus require that u is inherited from the global sheet counting map ψ ; technically we need $u = \delta\psi$ where δ is the coboundary map as defined in appendix B. This idea will be made precise in section 5.3.

In the rest of this section we proceed as follows. We first illustrate the above ideas using the example of the BTZ black hole. Here we refrain from rigor and just give qualitative arguments for the validity of the argument. section 5.3 then presents a rigorous and general proof that the topological consistency condition is satisfied if and only if \mathcal{A} is homologous to \mathcal{E} in an appropriate sense.

5.2 Example: BTZ black hole

Let us consider the single interval entanglement entropy of a CFT_2 on a spatial circle and at finite temperature, in the deconfined phase above the Hawking-Page transition. The Euclidean bulk \mathcal{M} is topologically a solid torus. For illustration, consider the various extremal surfaces sketched in figure 3. In each case we want to study the relation between our topological consistency condition (4.2) and the homological properties of \mathcal{E} compared to \mathcal{A} . The geodesics $\mathcal{E}^{(a)}$ and $\mathcal{E}^{(c)}$ are both homologous to \mathcal{A} (i.e., in each case there exists an interpolating surface whose only boundaries are \mathcal{A} and $\mathcal{E}^{(a,c)}$) and should therefore be allowed saddles in the RT formula. As we will see, both of them satisfy the topological consistency condition for all q and are thus consistent with the LM construction. The extremal surface $\mathcal{E}^{(b)}$, on the other hand, is not homologous to \mathcal{A} which will manifest itself as a breakdown of the topological consistency condition (meaning that $\mathcal{E}^{(b)}$ can never arise as the fixed point set of a LM-type argument).

By inspection, it is evident from the examples in figure 3 how homology of \mathcal{A} and \mathcal{E} are tied to consistency of the boundary sheet counting map ϕ and the bulk intersection map u . For instance, in case (a) we have drawn two choices of surfaces \mathcal{D} with boundary on $\mathcal{B} - \partial\mathcal{A}$ and clearly for both of them the number of intersections of $\partial\mathcal{D}$ with \mathcal{A} matches the




Figure 3. Three different candidates for bulk geodesics (blue) whose length may be considered to compute the entanglement entropy of the spatial boundary interval \mathcal{A} (red). The green surfaces \mathcal{D} illustrate the connection between boundary sheet counting (intersections with \mathcal{A}) and bulk intersections. Thick dots illustrate intersections. The bulk surfaces $\mathcal{E}^{(a)}$ and $\mathcal{E}^{(c)}$ are homologous to \mathcal{A} ; which one of them is dominant depends on the length of \mathcal{A} . The surface $\mathcal{E}^{(b)}$ is not homologous to \mathcal{A} and is also forbidden by our topological consistency condition. In this latter case, boundary intersections do not match bulk intersection numbers.

local intersection number of \mathcal{D} with $\mathcal{E}^{(a)}$ in the bulk. The reader can easily verify that this matching indeed holds for any choice of 2-surface \mathcal{D} anchored on $\mathcal{B} - \partial\mathcal{A}$. This consistency of ϕ and u allows for a global sheet counting map $\psi \in H^1(\mathcal{M} - \mathcal{E})$ which reduces to ϕ at the boundary.

Similar reasoning holds for case (c): the only novelty here is the fact that orientations need to be taken into account to properly count intersections. For instance, the intersection number of the rear disk with $\mathcal{E}^{(c)}$ is zero due to opposite orientations of $\mathcal{E}^{(c)}$ relative to \mathcal{D} at the two intersection points. This matches the fact that the boundary of the rear disk does not intersect with \mathcal{A} .

Now consider case (b) where $\mathcal{E}^{(b)}$ indeed violates the homology constraint. As we can see, there now exist 2-surfaces \mathcal{D} on which ϕ (intersections of $\partial\mathcal{D}$ with \mathcal{A}) takes a different value than u (intersection number of \mathcal{D} with $\mathcal{E}^{(b)}$). Therefore the boundary sheet counting map ϕ does not lift to a global sheet counting map ψ consistent with u on loops which are boundaries of 2-surfaces that intersect $\mathcal{E}^{(b)}$. Such a lift not being possible means that covers of the bulk branched along $\mathcal{E}^{(b)}$ do not have the correct boundary conditions. This situation can therefore never arise from a LM construction.

At a pictorial level, this demonstrates our claims for the BTZ geometry: the homology constraint on \mathcal{A} and \mathcal{E} comes in conjunction with consistency of bulk intersection number with \mathcal{E} on the one hand, and boundary intersection number with \mathcal{A} on the other hand. This consistency is equivalent to the topological consistency condition (4.2). We will now turn to a rigorous proof of this idea.

5.3 Topological consistency is equivalent to homology constraint

We begin with a reminder of the notation. The bulk \mathcal{M} is a $d + 1$ dimensional orientable manifold with boundary \mathcal{B} , \mathcal{A} is a $d - 1$ dimensional submanifold of \mathcal{B} with boundary $\partial\mathcal{A}$, and \mathcal{E} is a $d - 1$ dimensional submanifold of \mathcal{M} , also with boundary $\partial\mathcal{A}$. It should be

borne in mind that either or both of \mathcal{A} and \mathcal{E} may be disconnected. For this proof, we will focus on the case when the bulk is orientable; we will return to the extension to the non-orientable case in the next subsection. A review of the algebraic topology required for this section can be found in appendix B.

To construct the cover on the boundary, we have a map from $\pi_1(\mathcal{B} - \partial\mathcal{A})$ to \mathbb{Z} , counting signed intersections with \mathcal{A} . This is equivalent to a map $H_1(\mathcal{B} - \partial\mathcal{A}) \rightarrow \mathbb{Z}$, since this homology group¹² is the abelianization of the fundamental group by the Hurewicz theorem. In turn, by the universal coefficient theorem, this homomorphism is equivalent to a cohomology class $\phi \in H^1(\mathcal{B} - \partial\mathcal{A})$.

This sheet counting map is equivalent to a homological description of the region \mathcal{A} . Regard $\mathcal{B} - \partial\mathcal{A}$ as the boundary minus a tubular open neighbourhood of $\partial\mathcal{A}$, so it becomes a compact manifold with boundary. We may then use Poincaré-Lefschetz duality (B.6) to send $\phi \in H^1(\mathcal{B} - \partial\mathcal{A})$ to $a \in H_{d-1}(\mathcal{B} - \partial\mathcal{A}, \partial(\mathcal{B} - \partial\mathcal{A})) \approx H_{d-1}(\mathcal{B}, \partial\mathcal{A})$. This is a homology class represented by an appropriately chosen orientation of \mathcal{A} .

The arguments above show that the bulk lifts to an appropriate branched cover at all q whenever there exists some extension of ϕ into the bulk. So we require some $\psi \in H^1(\mathcal{M} - \mathcal{E})$ satisfying $i^*\psi = \phi$, where $i : \mathcal{B} - \partial\mathcal{A} \rightarrow \mathcal{M} - \mathcal{E}$ is the inclusion of boundary into bulk, so that the restriction of ψ to the boundary is ϕ .

In addition to this, we need to know what happens when we traverse a small loop passing round the branching surface \mathcal{E} . We make this precise by considering the normal bundle of \mathcal{E} , which can be embedded in \mathcal{M} as some tubular neighbourhood of \mathcal{E} . The fibres F of the bundle are copies of \mathbb{R}^2 , with \mathcal{E} lying at the origin, so the action of ψ passing round the surface is described by its restriction to $F - \{0\}$ on each fibre, $\psi_F \in H^1(F - \{0\}) \approx \mathbb{Z}$. We would normally like this restriction to each fibre to be a generator of $H^1(F - \{0\})$, which implies that after taking the q -fold cover and the \mathbb{Z}_q quotient there will be a $\frac{2\pi}{q}$ conical defect at every q . This need not hold in general, as exemplified in section 3, and the conical defect could be of a different angle depending on q in a more complicated way. This can be accounted for in the end by counting such a surface with the appropriate multiplicity in the homology computation, and by allowing this multiplicity we may assume that ψ_F is always a generator.

So we would like necessary and sufficient conditions for the existence of some $\psi \in H^1(\mathcal{M} - \mathcal{E})$, such that

1. The restriction $i^*\psi$ to the boundary is $\phi \in H^1(\mathcal{B} - \partial\mathcal{A})$, and
2. The restriction ψ_F to the embedding into \mathcal{M} of the fibres F of the normal bundle N of \mathcal{E} is a generator of $H^1(F - \{0\})$.

The choice of ψ_F will in the course of the proof supply us with an appropriate orientation of \mathcal{E} , and thus an element $e \in H_{d-1}(\mathcal{E}, \partial\mathcal{A})$. We now have enough to state our main result.

¹²All homology and cohomology groups are taken with coefficients in \mathbb{Z} , so we will not explicitly indicate this dependence until we later generalize to the non-orientable case.

Theorem 1. *There exists some $\psi \in H^1(\mathcal{M} - \mathcal{E})$ which restricts to ϕ on the boundary, and which restricts to ψ_F on the fibres of N , if and only if \mathcal{A} is homologous to \mathcal{E} , in the sense that the inclusions of e and a into $H_{d-1}(\mathcal{M}, \partial\mathcal{A})$ are equal.*

Proof. The choice of generators of $H^1(F - \{0\})$ on each fibre is equivalent to an orientation of each fibre, which can be regarded equivalently as a smooth choice of generators of the relative cohomology groups $H^2(F, F - \{0\})$, via a coboundary map. In each connected component of \mathcal{E} , making a choice on one fibre leaves no freedom on the others, which are determined uniquely from continuity. From this we can get a well-known object in the topology of vector bundles, a Thom class $u \in H^2(N, N - \mathcal{E})$, where \mathcal{E} here sits at the zero section. This follows since if \mathcal{E} is connected, then the restrictions $H^2(N, N - \mathcal{E}) \rightarrow H^2(F, F - \{0\})$ are in fact all isomorphisms. This is clear when restricted to a local trivialization of the bundle, and can be extended to the whole of the connected component by a standard procedure of gluing trivializations together one by one using a Mayer-Vietoris sequence. So our choices of ψ_F can be summarized by a choice of generator $u \in H^2(N, N - \mathcal{E})$ for each connected component of \mathcal{E} . See section 4.D of [17] for a discussion of the Thom class.

Now regarding N once again as a tubular neighbourhood¹³ of \mathcal{E} in \mathcal{M} , we can use an excision theorem to give an isomorphism $H^2(N, N - \mathcal{E}) \approx H^2(\mathcal{M}, \mathcal{M} - \mathcal{E})$ (excising $\mathcal{M} - \mathcal{E}$), so we may regard u instead as an element of $H^2(\mathcal{M}, \mathcal{M} - \mathcal{E})$. Roughly speaking, we may think of u as a map from two-dimensional chains with boundaries away from \mathcal{E} , which counts, with signs, the number of intersections with \mathcal{E} .

Now if \mathcal{M} is oriented, the orientation of the normal bundle of \mathcal{E} induces also an orientation of the tangent bundle of \mathcal{E} . This can be characterized by a relative homology class $e \in H_{d-1}(\mathcal{E}, \partial\mathcal{A})$.

Consider now the two cohomological long exact sequences, of the pair $(\mathcal{M} - \mathcal{E}, \mathcal{B} - \partial\mathcal{A})$, and of the triple $(\mathcal{M}, \mathcal{M} - \mathcal{E}, \mathcal{B} - \partial\mathcal{A})$, as explained in (B.3), which fit into a diagram:

$$\begin{array}{ccccccc}
 & & H^1(\mathcal{M} - \mathcal{E}) & \xrightarrow{i^*} & H^1(\mathcal{B} - \partial\mathcal{A}) & & \\
 & \nearrow & \uparrow j^* & & \downarrow \delta & \searrow \delta & \\
 & H^1(\mathcal{M} - \mathcal{E}, \mathcal{B} - \partial\mathcal{A}) & & & & & H^2(\mathcal{M} - \mathcal{E}, \mathcal{B} - \partial\mathcal{A}) \\
 & \nwarrow & \downarrow \delta & & \downarrow \delta & \nearrow i^* & \\
 & & H^2(\mathcal{M}, \mathcal{M} - \mathcal{E}) & \xrightarrow{j^*} & H^2(\mathcal{M}, \mathcal{B} - \partial\mathcal{A}) & &
 \end{array}$$

where the i^* and j^* maps are the relevant restrictions and extensions (dual to inclusion and quotient) respectively, and all the δ s are various relative coboundary maps. The top and bottom rows are the usual exact sequences of relative cohomology, and it is straightforward to show that the diagram commutes at the level of cochains. It is then an easy exercise¹⁴ to show that this induces another long exact sequence, the crucial part of which is

$$H^1(\mathcal{M} - \mathcal{E}) \xrightarrow{(i^*, \delta)} H^1(\mathcal{B} - \partial\mathcal{A}) \oplus H^2(\mathcal{M}, \mathcal{M} - \mathcal{E}) \xrightarrow{\delta - j^*} H^2(\mathcal{M}, \mathcal{B} - \partial\mathcal{A})$$

¹³The existence of such a neighborhood is in fact a technical assumption, though this can be relaxed [18].

¹⁴Exercise 38 in chapter 2.2 of [17].

where the maps are the obvious ones from the diagram above, except that one factor in the second map has a minus sign.

Now an element $\psi \in H^1(\mathcal{M} - \partial\mathcal{A})$ restricting to ϕ on the boundary means that $i^*\psi = \phi$, and restricting to ψ_F on the fibres is equivalent to $\delta\psi = u$. So ψ with the desired properties exists iff $\phi \oplus u$ is in the image of the first map, which by exactness is the kernel of the second, so this is equivalent to $\delta\phi = j^*u$.

Finally, there is a generalization of Poincaré-Lefschetz duality (B.6) that by splitting \mathcal{B} into a tubular neighborhood of $\partial\mathcal{A}$ and its complement gives us an isomorphism $H^2(\mathcal{M}, \mathcal{B} - \partial\mathcal{A}) \approx H_{d-1}(\mathcal{M}, \partial\mathcal{A})$, via cap product with the fundamental class $\mu \in H_{d+1}(\mathcal{M}, \mathcal{B})$. The theorem follows by dualizing the equality $\delta\phi = j^*u$ under this, and showing that the duals of $\delta\phi$ and j^*u are respectively the inclusions of a and e into $H_{d-1}(\mathcal{M}, \mathcal{B} - \partial\mathcal{A})$.

The dual of $\delta\phi$ is $\mu \frown \delta\phi$, which equals the inclusion of $\partial\mu \frown \phi$ into the bulk. This is indeed the correct thing, since $\partial\mu \frown \phi$ is just the boundary Poincaré-Lefschetz duality we originally used to relate a and ϕ .

The final part, computing the dual of j^*u , generalizes a well-known result in the case of closed manifolds [18]. The dual of j^*u is $\mu \frown j^*u = j_*\mu \frown u$, where $j_*\mu \in H_{d+1}(\mathcal{M}, (\mathcal{M} - \partial\mathcal{A}) \cup \mathcal{B})$. By restricting this cap product to a tubular neighborhood N of $\partial\mathcal{A}$, with inclusion map $k : N \rightarrow \mathcal{M}$, it is clear that the result must be in the image of k_* . Choosing N such that it retracts onto \mathcal{E} by a retraction $r : N \rightarrow \mathcal{E}$ homotopic to the identity, it further follows that $k_* = i_*r_*$, so we get a cycle in the image of i_* . Working separately in each connected component of \mathcal{E} , this must be $n i_*e$ for some integer n . A local analysis in a trivialization of N , where the computations are simply in \mathbb{R}^{d+1} , shows that $n = 1$, so the dual of j^*u must in fact equal i_*e as required. □

5.4 Non-orientable manifolds

In the case when the bulk is not orientable, the cohomological parts of the above proof go through unchanged, but the dualities used to make statements in terms of homology are not applicable. In fact, since \mathcal{E} , for instance, need not be orientable, there may not even be a homology class in $H_{d-1}(\mathcal{M}, \partial\mathcal{A}; \mathbb{Z})$ that represents it, so it becomes less clear how to state the result. This difficulty can be overcome by passing to \mathbb{Z}_2 coefficients, but since the correct cohomological argument requires integer coefficients, this gives an insufficiently strong constraint on allowed \mathcal{E} . The correct thing is instead to use homology with local (or twisted) coefficients (see [17], section 3.H).

The generalization to local coefficients is performed by, roughly speaking, allowing the coefficients of the chain complexes to live in some module of the fundamental group of the spacetime. In our case, this will be the orientation class, being the module with action given by ± 1 depending on whether traversing a loop preserves or reverses orientation.

There is a generalized Poincaré-Lefschetz duality that holds for non-orientable spacetimes, between integer cohomology, and homology with coefficients in this module twisted by the orientation class. Since the cohomology we use is ambivalent to the presence or absence of orientation, nothing is altered until the very end, when we take this duality, and as such the homology statement will be in terms of local coefficients.

6 Discussion

The generalized gravitational entropy construction of LM, which is inspired by the replica trick in field theory, provides a derivation of the RT prescription for holographic entanglement entropy. We examined the conditions under which the LM construction guarantees, at the level of topology, that the extremal surface computing entanglement entropy satisfies the homology constraint in the bulk.

Suppose we have a bulk Euclidean geometry \mathcal{M} , a boundary region \mathcal{A} and a minimal surface \mathcal{E} , a candidate for computing the entanglement entropy of \mathcal{A} . The local, dynamical part of the LM argument relates this to the replica trick as follows. Introduce a conical deficit at \mathcal{E} , with defect angle $\frac{2\pi}{q}$, increasing q away from unity, whilst changing the geometry to keep it on shell away from \mathcal{E} . Locally, near \mathcal{E} , we may choose a ‘cigar’ where there are polar coordinates parametrising a distance from \mathcal{E} and an angle around it, with \mathcal{E} itself at the origin. When q is an integer, this cigar with a conical deficit can be unwrapped by allowing the period of the angle to become q times larger, so we obtain a smooth space. The resulting boundary is *locally* the correct space with which to compute the q^{th} Rényi entropy, which makes contact with the replica trick.

It is less obvious that this picture works globally, since there is usually not a globally valid ‘cigar’, with a circle parameterised by a global ‘Euclidean time’ shrinking to zero size at \mathcal{E} . There is no obvious obstruction to introducing a conical defect at any minimal surface to give a continuous family of geometries, but at integer q , when the defects are locally unwrapped, it may not be possible to consistently extend the local pictures to the whole spacetime, or doing so may produce the Rényi entropy for the wrong region. For example, on a torus, the replicated geometries for an interval and its complement are locally the same, but differ globally by what happens on traversing the nontrivial cycles of the torus. This global picture, left implicit in LM, is crucial for understanding which choices for \mathcal{E} give rise to the correct Rényi entropies.

Our main result addresses this global aspect, and can be phrased as follows: *A conical deficit introduced at a given extremal surface \mathcal{E} comes from a quotient of a geometry whose boundary is the replicated space relevant for the q th Rényi entropy of region \mathcal{A} , for every positive integer q , if and only if \mathcal{E} is homologous to \mathcal{A} .* In particular, using the terminology explained in the introduction, the local interpretation of the LM assumptions suffices to enforce the homology constraint so long as it holds at all q and is consistent with the replica trick at integer q .

In this statement it is crucial that the construction of the bulk covering space works at all integer values of q . We gave an explicit example to show that branching along homology violating extremal surfaces can occur when the replica symmetric covering space exists for some value of q , but not for others. For the cosmic brane construction of LM [4] which starts from a singular spacetime \mathcal{M}_q , the homology constraint is implemented in the $q \rightarrow 1$ limit, iff for all integer q , \mathcal{M}_q lifts to a non-singular q -fold replica symmetric branched cover (asymptoting to the replicated boundary geometry). We emphasize that — since we have worked only at the level of topology — these \mathcal{M}_q need not be actual solutions to the theory but merely smooth manifolds satisfying appropriate boundary conditions at

the cosmic brane and at infinity. The dynamical part of the argument, which requires the geometries to be on-shell (at least in a neighbourhood of $q = 1$), is essentially independent of our considerations.

As such our analysis requires that the entanglement entropy is computed by an extremal surface which arises within a family of cosmic brane configurations as a limiting case. We do not require that the Rényi entropies themselves for integer values of q are computed by the elements of the same family. They could be computed (at large central charge) by other families, which would allow for Rényi phase transitions as in [19], or even isolated configurations. This also implies that the logarithmic negativity, which is obtained for bi-partitioning of a pure state by analytic continuation of the even- q Rényi entropies [20], could arise from a completely different family, for we only need a family of cosmic brane solutions for even integral q .

We emphasize once more that the arguments were purely topological, with no recourse to dynamical information of the bulk gravitational theory. In particular, they hold for any choice of the boundary region \mathcal{A} including disconnected ones. The statements regarding topological consistency, covering spaces, and the proof of Theorem 1 are applicable to such situations. Of course, being agnostic of the dynamics has drawbacks, in that we will be unable to decide which of the homology respecting extremal surfaces actually computes the entanglement entropy. On the other hand, this gives us the advantage of being able to allow breaking replica symmetry away from integral values of q , as perhaps may be necessary for non Einstein-Hilbert gravitational dynamics [21].

We have assumed the bulk spacetime to be orientable in our discussion. For non-orientable spacetimes as indicated in section 5.4 the natural generalization involves homology groups with twisted coefficients as opposed to integral homology used in our proof. This is relevant, for example, in the case of the \mathbb{RP}^2 -geon spacetime [22], which is a pure state of a two-dimensional CFT (on a Klein bottle). In the bulk the spacetime has a horizon but a single asymptotic boundary (hence single-exterior black hole). In this case the homology condition is necessary to argue that the entanglement entropy for a region and its complement (on a spatial circle of the boundary) are the same as required by purity; see [23] for an analysis of the entanglement structure in the \mathbb{RP}^2 -geon (on the $t = 0$ slice the \mathbb{Z}_2 -valued homology group will suffice). It would be interesting to flesh out the details of the twisted homology constraint for more general physical examples.

Finally, let us also note that our analysis takes seriously the Euclidean LM construction. We only require the boundary region \mathcal{A} to lie on the time-reflection symmetric surface. We are not a-priori guaranteed then that the extremal surface \mathcal{E} , and consequentially the homology surface $\mathcal{R}_{\mathcal{A}}$, lie at the moment of time-reflection symmetry in the bulk (see [23] for an example in thermal AdS where a sub-dominant extremal surface of this kind was described). If such surfaces give the dominant contribution to the holographic entanglement entropy, then various results in the holographic entanglement entropy literature such as strong-subadditivity [6, 11] and causality [8, 11] would have to be revisited.

Acknowledgments

We would like to thank Matthew Headrick, Veronika Hubeny and Aitor Lewkowycz for illuminating discussions. DM and MR would like to acknowledge the hospitality of the Issac Newton Institute, Cambridge where this project was initiated. TH, DM and MR would also like to thank the Aspen Center for Physics (supported by the National Science Foundation under Grant 1066293) for their hospitality.

FH is supported by a Durham Doctoral Fellowship. DM was supported in part by the U.S. National Science Foundation under Grant No PHY11-25915, and by funds from the University of California. HM is supported by a STFC studentship. MR was supported in part by the FQXi grant “Measures of Holographic Information” (FQXi-RFP3-1334), the STFC Consolidated Grant ST/L000407/1, and by the ERC Consolidator Grant ERC-2013-CoG-615443: SPiN (Symmetry Principles in Nature).

A Equivalence of topological consistency conditions

We here prove two theorems which will make precise the equivalence between the existence of a replica symmetric covering space as in (4.3), and the existence of a sheet counting map in the bulk factoring the boundary sheet counting map by the inclusion (4.1).

To make the structure of the arguments clearer, we will prove something more general. We will take X to be a subspace of Y , and look for covering spaces \tilde{X} and \tilde{Y} , with an inclusion map consistent with the inclusion of the base spaces. We additionally require the covers to possess certain symmetries, specifically that they are normal with deck transformation groups G and H respectively. The existence of such a cover will be shown to be equivalent to the existence of a commutative square of homomorphisms between $\pi_1(X)$, $\pi_1(Y)$, G and H . See section 1.3 of [17] for the necessary mathematical material.

To apply this in the context we require it, identify X and Y with boundary and bulk geometries respectively, both with the branching surfaces removed. Then G and H are both taken as \mathbb{Z}_q , so that we are requiring the covers to be replica symmetric. The square of homomorphisms will then be equivalent to (4.1).

The first part of the proof, Theorem 2, constructs the covering spaces given the appropriate homomorphisms. Theorem 3 shows the converse, that the appropriate homomorphisms exist starting from the covering spaces.

Theorem 2. *Let X and Y be path-connected, locally path-connected, semi-locally simply-connected spaces, and $i : X \rightarrow Y$ an injective inclusion map. Let $\phi : \pi_1(X) \rightarrow G$ and $\psi : \pi_1(Y) \rightarrow H$ be surjective homomorphisms, and $\rho : G \rightarrow H$ an injective homomorphism such that*

$$\begin{array}{ccc}
 \pi_1(X) & \xrightarrow{\phi} & G \\
 \downarrow i_* & & \downarrow \rho \\
 \pi_1(Y) & \xrightarrow{\psi} & H
 \end{array}$$

commutes. Then there are covering spaces $p : \tilde{X} \rightarrow X$, and $P : \tilde{Y} \rightarrow Y$, with deck transformation groups G, H and an injective inclusion map $\tilde{i} : \tilde{X} \rightarrow \tilde{Y}$ such that

$$\begin{array}{ccc} \tilde{X} & \xrightarrow{p} & X \\ \downarrow \tilde{i} & & \downarrow i \\ \tilde{Y} & \xrightarrow{P} & Y \end{array}$$

commutes.

In the case when $G = H (= \mathbb{Z}_q, \text{ say})$, and the groups are finite, ρ is an isomorphism so we can identify the two groups. This is the way we use the theorem above. A similar remark applies for the second theorem.

Proof. Denote homotopy equivalence classes of loops by $[\cdot]$, and reversal of curves by $\bar{\cdot}$.

We construct \tilde{X} as the set of equivalence classes $[\gamma]_{\tilde{X}}$ of curves $\gamma : [0, 1] \rightarrow X$ starting at the basepoint in X , where $\gamma \sim_{\tilde{X}} \eta$ if (i) $\gamma(1) = \eta(1)$ and (ii) $[\gamma\bar{\eta}] \in \ker(\phi)$. The projection map is $p : [\gamma]_{\tilde{X}} \mapsto \gamma(1)$, which is clearly well-defined. This is the standard construction of the covering space with fundamental group $\ker(\phi)$. The deck transformation group is $\pi_1(X)/\ker(\phi) \cong G$, since ϕ is surjective. Similar statements apply to construct \tilde{Y} .

It remains only to construct the inclusion map \tilde{i} , by $\tilde{i} : [\gamma]_{\tilde{X}} \mapsto [i\gamma]_{\tilde{Y}}$. This is well defined, since if $\gamma \sim_{\tilde{X}} \eta$, then (i) $i\gamma(1) = i\eta(1)$, and (ii) $\psi[i\gamma\bar{i\eta}] = \psi i_*[\gamma\bar{\eta}] = \rho\phi[\gamma\bar{\eta}] = 1$ since $[\gamma\bar{\eta}] \in \ker \phi$, so $[i\gamma\bar{i\eta}] \in \ker \psi$. Finally, \tilde{i} is injective, since if $[i\gamma]_{\tilde{Y}} = [i\eta]_{\tilde{Y}}$, then (i) $i\gamma(1) = i\eta(1) \Rightarrow \gamma(1) = \eta(1)$ by injectivity of i , and (ii) $1 = \psi[i\gamma\bar{i\eta}] = \psi i_*[\gamma\bar{\eta}] = \rho\phi[\gamma\bar{\eta}] \Rightarrow [\gamma\bar{\eta}] \in \ker(\phi)$, where we have used injectivity of ρ in the last step. \square

Theorem 3. *Suppose that there are normal covering spaces $p : \tilde{X} \rightarrow X$, and $P : \tilde{Y} \rightarrow Y$, with deck transformation groups G, H and injective inclusion maps $i : X \rightarrow Y$ and $\tilde{i} : \tilde{X} \rightarrow \tilde{Y}$ such that*

$$\begin{array}{ccc} \tilde{X} & \xrightarrow{p} & X \\ \downarrow \tilde{i} & & \downarrow i \\ \tilde{Y} & \xrightarrow{P} & Y \end{array}$$

commutes. Then an injective homomorphism $\rho : G \rightarrow H$ exists, such that the following diagram commutes:

$$\begin{array}{ccccccc} \pi_1(\tilde{X}) & \xrightarrow{p_*} & \pi_1(X) & \xrightarrow{\phi} & G & \longrightarrow & 1 \\ \downarrow \tilde{i}_* & & \downarrow i_* & & \downarrow \rho & & \\ \pi_1(\tilde{Y}) & \xrightarrow{P_*} & \pi_1(Y) & \xrightarrow{\psi} & H & \longrightarrow & 1 \end{array}$$

The maps ϕ and ψ here are the quotient maps by the fundamental groups of the covering spaces (using the fact that the coverings are normal), which implies that the rows are exact.

Proof. Let $g \in G$. Since ϕ is surjective, there is some loop γ in X so that $\phi([\gamma]) = g$. Define $\rho(g) = \psi i_*[\gamma]$. We must check that this is well defined, so suppose that η is another loop in X with $\phi([\eta]) = g$. Then $\phi([\gamma\bar{\eta}]) = \phi([\gamma][\eta]^{-1}) = 1$. This means that $\gamma\bar{\eta}$ lifts to a loop

$\tilde{\gamma}$ in \tilde{X} , by exactness of the top row. Now $i_*[\gamma\bar{\eta}] = i_*p_*[\tilde{\gamma}] = P_*\tilde{i}_*[\tilde{\gamma}] \in \ker(\psi)$ by exactness of the bottom row, so $\psi i_*[\gamma] = \psi i_*[\eta]$. It's also clear from the definition that $\psi i_* = \rho\phi$.

Finally, we check injectivity of ρ , so let $g \in G$ be distinct from the identity. Pick a loop γ with $\phi[\gamma] = g$, and lift it to a curve $\tilde{\gamma}$ in \tilde{X} . Since g is not the identity, this is not a loop, and since \tilde{i} is injective, the curve $\tilde{i}\tilde{\gamma}$ in \tilde{Y} is also not a loop. By the commutativity, this is the same curve as the lift of $i\gamma$, which implies that $[i\gamma] \notin \text{im}(P_*) = \ker(\psi)$. So $\rho(g) = \psi([i\gamma])$ is not the identity, and ρ is injective. \square

B Review of algebraic topology

This appendix reviews the algebraic topology required for the main result of the paper. See [17] for more detailed discussions.

The space of singular n -chains $C_n(X)$ of a topological space X is the free abelian group with basis singular n -simplices, which are maps from the standard n -simplex $\Delta^n = \{(t_0, \dots, t_n) \in \mathbb{R}^{n+1} \mid \sum_i t_i = 1, t_i \geq 0\}$ to X . The boundary maps $\partial_n : C_n(X) \rightarrow C_{n-1}(X)$ act on chains σ by a sum of restrictions of σ to its $n+1$ faces, with appropriate signs. A cycle is defined as a chain with zero boundary, in the kernel of ∂ , and a boundary is a chain in the image of ∂ . Since $\partial_n\partial_{n+1} = 0$, all boundaries are cycles, and this means we can define the singular homology groups as cycles modulo boundaries: $H_n(X) = \ker \partial_n / \text{im } \partial_{n+1}$.

A generalization of this is a relative homology group, where we ignore what goes on in some part of the space. Let A be a subset of the space X ; the space of relative chains is the space of chains in X modulo chains in A , $C_n(X, A) = C_n(X)/C_n(A)$. The usual boundary map continues to be well-defined between relative chains, so we may again take cycles modulo boundaries to get the relative homology groups $H_n(X, A)$. Cycles here may therefore have a boundary in the subspace A , and it is trivial if it is homologous to a cycle completely within A . Often the space A will be the complement of some set, $A = X - B$, when $H_n(X, X - B)$ is sometimes referred to as ‘homology at B ’, since it ignores what happens to chains except at the subspace B .

Since the relative homology ignores what goes on in some subspace, nothing is lost by removing some part of that subspace. This is made precise by the excision theorem, which states that if $Z \subseteq A \subseteq X$, and the closure of Z is contained in the interior of A , then $H_n(X, A)$ is isomorphic to $H_n(X - Z, A - Z)$.

Of central importance is the fact that relative homology groups fit into a long exact sequence:

$$\dots \xrightarrow{\partial} H_n(A) \xrightarrow{i_*} H_n(X) \xrightarrow{j_*} H_n(X, A) \xrightarrow{\partial} H_{n-1}(A) \xrightarrow{i_*} \dots \tag{B.1}$$

The maps i_* here are inclusions of cycles in A into X , and the maps j_* are the quotients by chains in A to get to relative homology. Finally, the maps ∂ are boundary maps: any relative n -cycle must have boundary contained in A by definition, and this gives a homology class in $H_{n-1}(A)$. Exactness is geometrically very intuitive: for example, a cycle in $H_n(A)$ gives zero when it is included into $H_n(X)$ if and only if it is the boundary of some chain in X , which is the case if and only if it is in the image of the boundary map ∂ from the relative homology group. So $\ker i_* = \text{im } \partial$.

This can be slightly generalized to the long exact sequence of the triple (X, A, B) , where $B \subseteq A \subseteq X$, by doing everything relative to the smallest subspace B :

$$\dots \xrightarrow{\partial} H_n(A, B) \xrightarrow{i_*} H_n(X, B) \xrightarrow{j_*} H_n(X, A) \xrightarrow{\partial} H_{n-1}(A, B) \xrightarrow{i_*} \dots \quad (\text{B.2})$$

Now most of the work we will do will be not in terms of homology, but cohomology, to which all of the above carries over. To pass to cohomology, all the relevant spaces and maps should be dualized: we consider spaces of cochains $C^n(X) = \text{hom}(C_n(X), \mathbb{Z})$, joined by coboundary maps $\delta_n : C^n(X) \rightarrow C^{n+1}(X)$ dual to the boundary maps, so $\delta f(x) = f(\partial x)$. The kernel of δ gives the cochains, and the image the coboundaries; the cohomology groups are cochains mod coboundaries $H^n(X) = \ker \delta_n / \text{im } \delta_{n-1}$.

The constructions for relative homology groups carry over to cohomology, by dualizing the relative chain complex. In particular the excision theorem is the same in cohomology, and the long exact sequences of relative homology groups dualize, for example for the triple (X, A, B) we have

$$\dots \xrightarrow{i^*} H^{n-1}(A, B) \xrightarrow{\delta} H^n(X, A) \xrightarrow{j^*} H^n(X, B) \xrightarrow{i^*} H^n(A, B) \xrightarrow{\delta} \dots \quad (\text{B.3})$$

While most of the argument is phrased in terms of cohomology, in the end we want to translate the result into homology language. The crucial tool to do this is a generalization of Poincaré duality. Geometrically, these dualities can be roughly thought of as taking a submanifold to a function on submanifolds of complementary dimension, counting with signs the number of intersections between the submanifolds.

This requires the cap product between chains and cochains, defined by

$$\begin{aligned} \frown : C_k(X) \times C^l(X) &\longrightarrow C_{k-l}(X) \\ \sigma \frown \varphi &= \varphi(\sigma_{[v_0, \dots, v_l]}) \cdot \sigma_{[v_l, \dots, v_k]} \end{aligned} \quad (\text{B.4})$$

where $k \geq l$ and $\sigma_{[v_0, \dots, v_l]}$ denotes the restriction of σ to the simplex spanned by the vertices v_0, \dots, v_l . This map induces a cap product between homology and cohomology classes.

A connected closed orientable n -manifold X has top homology group $H_n(X)$ isomorphic to \mathbb{Z} , and an orientation of X is equivalent to a generator of this group, a *fundamental class* $[X]$. The Poincaré duality map is given by the cap product with the fundamental class $[X] \frown : H^k(X) \rightarrow H_{n-k}(X)$, which is in fact an isomorphism between complementary homology and cohomology groups.

This duality generalizes to manifolds with boundary. Crucial to this is a relative version of the cap product. Using the same definitions as before, it can be checked that the product on chains and cochains induces a product on relative (co)homology groups

$$\frown : H_k(X, A \cup B) \times H^l(X, A) \longrightarrow H_{k-l}(X, B) \quad (\text{B.5})$$

where A and B are open sets in X . The duality theorem we need applies this in the case of an oriented n -manifold X with boundary ∂X , where the boundary is decomposed as $\partial X = A \cup B$, where A and B are $(n - 1)$ -dimensional manifolds with common boundary

$A \cap B = \partial A = \partial B$. An orientation of X is defined in this case by a generator of $H_n(X, \partial X)$, which as before gives us isomorphisms

$$[X] \frown: H^k(X, A) \xrightarrow{\cong} H_{n-k}(X, B). \quad (\text{B.6})$$

This includes the special cases where $A = \emptyset$ and $B = \partial X$, or vice-versa, known as Poincaré-Lefschetz duality.

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

References

- [1] S. Ryu and T. Takayanagi, *Holographic derivation of entanglement entropy from AdS/CFT*, *Phys. Rev. Lett.* **96** (2006) 181602 [[hep-th/0603001](https://arxiv.org/abs/hep-th/0603001)] [[INSPIRE](#)].
- [2] S. Ryu and T. Takayanagi, *Aspects of holographic entanglement entropy*, *JHEP* **08** (2006) 045 [[hep-th/0605073](https://arxiv.org/abs/hep-th/0605073)] [[INSPIRE](#)].
- [3] V.E. Hubeny, M. Rangamani and T. Takayanagi, *A covariant holographic entanglement entropy proposal*, *JHEP* **07** (2007) 062 [[arXiv:0705.0016](https://arxiv.org/abs/0705.0016)] [[INSPIRE](#)].
- [4] A. Lewkowycz and J. Maldacena, *Generalized gravitational entropy*, *JHEP* **08** (2013) 090 [[arXiv:1304.4926](https://arxiv.org/abs/1304.4926)] [[INSPIRE](#)].
- [5] D.V. Fursaev, *Proof of the holographic formula for entanglement entropy*, *JHEP* **09** (2006) 018 [[hep-th/0606184](https://arxiv.org/abs/hep-th/0606184)] [[INSPIRE](#)].
- [6] M. Headrick and T. Takayanagi, *A holographic proof of the strong subadditivity of entanglement entropy*, *Phys. Rev. D* **76** (2007) 106013 [[arXiv:0704.3719](https://arxiv.org/abs/0704.3719)] [[INSPIRE](#)].
- [7] M. Headrick, *General properties of holographic entanglement entropy*, *JHEP* **03** (2014) 085 [[arXiv:1312.6717](https://arxiv.org/abs/1312.6717)] [[INSPIRE](#)].
- [8] M. Headrick, V.E. Hubeny, A. Lawrence and M. Rangamani, *Causality & holographic entanglement entropy*, *JHEP* **12** (2014) 162 [[arXiv:1408.6300](https://arxiv.org/abs/1408.6300)] [[INSPIRE](#)].
- [9] X. Dong, *Holographic entanglement entropy for general higher derivative gravity*, *JHEP* **01** (2014) 044 [[arXiv:1310.5713](https://arxiv.org/abs/1310.5713)] [[INSPIRE](#)].
- [10] J. Camps, *Generalized entropy and higher derivative gravity*, *JHEP* **03** (2014) 070 [[arXiv:1310.6659](https://arxiv.org/abs/1310.6659)] [[INSPIRE](#)].
- [11] A.C. Wall, *Maximin Surfaces and the Strong Subadditivity of the Covariant Holographic Entanglement Entropy*, *Class. Quant. Grav.* **31** (2014) 225007 [[arXiv:1211.3494](https://arxiv.org/abs/1211.3494)] [[INSPIRE](#)].
- [12] T. Faulkner, *The Entanglement Renyi Entropies of Disjoint Intervals in AdS/CFT*, [arXiv:1303.7221](https://arxiv.org/abs/1303.7221) [[INSPIRE](#)].
- [13] T. Barrella, X. Dong, S.A. Hartnoll and V.L. Martin, *Holographic entanglement beyond classical gravity*, *JHEP* **09** (2013) 109 [[arXiv:1306.4682](https://arxiv.org/abs/1306.4682)] [[INSPIRE](#)].
- [14] M. Headrick, A. Lawrence and M. Roberts, *Bose-Fermi duality and entanglement entropies*, *J. Stat. Mech.* **1302** (2013) P02022 [[arXiv:1209.2428](https://arxiv.org/abs/1209.2428)] [[INSPIRE](#)].

- [15] E. Witten and S.-T. Yau, *Connectedness of the boundary in the AdS/CFT correspondence*, *Adv. Theor. Math. Phys.* **3** (1999) 1635 [[hep-th/9910245](#)] [[INSPIRE](#)].
- [16] J.M. Maldacena and L. Maoz, *Wormholes in AdS*, *JHEP* **02** (2004) 053 [[hep-th/0401024](#)] [[INSPIRE](#)].
- [17] A. Hatcher, *Algebraic Topology*, Cambridge University Press, Cambridge U.K. (2002).
- [18] H. Samelson, *On the Thom class of a submanifold*, *Michigan Math. J.* **12** (1965) 257.
- [19] A. Belin, A. Maloney and S. Matsuura, *Holographic phases of Renyi entropies*, *JHEP* **12** (2013) 050 [[arXiv:1306.2640](#)] [[INSPIRE](#)].
- [20] P. Calabrese, J. Cardy and E. Tonni, *Entanglement negativity in extended systems: a field theoretical approach*, *J. Stat. Mech.* (2013) P02008 [[arXiv:1210.5359](#)] [[INSPIRE](#)].
- [21] J. Camps and W.R. Kelly, *Generalized gravitational entropy without replica symmetry*, *JHEP* **03** (2015) 061 [[arXiv:1412.4093](#)] [[INSPIRE](#)].
- [22] J. Louko and D. Marolf, *Single exterior black holes and the AdS/CFT conjecture*, *Phys. Rev. D* **59** (1999) 066002 [[hep-th/9808081](#)] [[INSPIRE](#)].
- [23] H. Maxfield, *Entanglement entropy in three dimensional gravity*, *JHEP* **04** (2015) 031 [[arXiv:1412.0687](#)] [[INSPIRE](#)].