

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

The Holographic Principle and the Emergence of Spacetime

Permalink

<https://escholarship.org/uc/item/8355c1fr>

Author

Rosenhaus, Vladimir

Publication Date

2014

Peer reviewed|Thesis/dissertation

The Holographic Principle and the Emergence of Spacetime

by

Vladimir Rosenhaus

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Raphael Bousso, Chair

Professor Lawrence Hall

Professor Alex Filippenko

Spring 2014

The Holographic Principle and the Emergence of Spacetime

Copyright 2014

by

Vladimir Rosenhaus

Abstract

The Holographic Principle and the Emergence of Spacetime

by

Vladimir Rosenhaus

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Raphael Bousso, Chair

Results within string theory and quantum gravity suggest that spacetime is not fundamental but rather emergent, with the fundamental degrees of freedom living on a boundary surface of one lower dimension than the bulk. This thesis is devoted to studying the holographic principle and its realization for spacetimes with both negative and positive cosmological constant.

The holographic principle is most explicitly realized in the context of the AdS/CFT correspondence. We examine the extent to which AdS/CFT realizes the holographic principle and study the UV/IR relation. We study aspects of how bulk locality emerges within AdS/CFT. To this effect, we study how to reconstruct the bulk from boundary data. We study how such a reconstruction procedure is sensitive to large changes in the bulk geometry. We study if it is possible to reconstruct a subset of the bulk from a subset of the boundary data. We explore both local and nonlocal CFT quantities as probes of the bulk. One nonlocal quantity is entanglement entropy, and to this effect we construct a framework for computing entanglement entropy within the field theory.

The most ambitious application of the holographic principle would be finding the holographic dual to the multiverse. We investigate properties of this putative duality. We extend the UV/IR relation of AdS/CFT to the multiverse, with the UV cutoff of the theory on future infinity being dual to a late time cutoff (measure) in the bulk. We compare various measure proposals and examine their predictions.

Contents

List of Figures	v
List of Tables	xiv
1 Introduction	1
I Emergent Spacetime with $\Lambda < 0$: AdS/CFT	6
2 Holography for a Small World	9
2.1 Introduction	9
2.2 UV/IR	11
2.2.1 Review of UV/IR	11
2.2.2 Comments on Scale/Radius	11
2.3 Oscillating Particle	12
2.4 Scalar Field Solutions	14
2.4.1 Relativistic Wave Packet	14
2.4.2 Non-relativistic Mode	17
2.4.3 General Solution	18
2.5 Precursors	19
2.6 A Proposal for Finding the States	20
2.6.1 Proposal	21
2.6.2 Ultraboosted states	22
2.7 Conclusion	23
2.8 Equations for Energy Shell	23
3 Light-sheets and AdS/CFT	25
3.1 Introduction	25
3.2 The Causally Connected Region C	29
3.3 The Light-sheet Region L	30
3.3.1 Spacelike Holography vs Light-Sheets	31
3.3.2 Bounds on the Holographic Domain	33
3.4 Proof that $L = C$	36
3.5 Covariant Renormalization Group Flow	39
3.6 Nonlocal Operators and the Bulk Dual	41

3.7	Quantum Gravity Behind Event Horizons	46
3.8	Light-sheets from the conformal boundary of AdS	47
4	Null Geodesics, Local CFT Operators, and AdS/CFT for Subregions	49
4.1	Introduction	49
4.2	The Reconstruction Map	52
4.2.1	General Formulas	52
4.2.2	Global AdS	54
4.2.3	AdS-Rindler	55
4.2.4	Poincare Patch	58
4.2.5	Poincare-Milne	59
4.2.6	AdS-Rindler Revisited	62
4.3	A Simple, General Criterion for Continuous Classical Reconstruction: Cap- turing Null Geodesics	62
4.3.1	Unique Continuation, Null Geodesics, and RG Flow	64
4.3.2	The Diagnostic in Other Situations	65
4.4	Arguments for an AdS-Rindler Subregion Duality	67
4.4.1	Probing the Bulk	67
4.4.2	Hyperbolic Black Holes	68
4.5	Discussion	69
5	AdS Black Holes, the Bulk-Boundary Dictionary, and Smearing Func- tions	71
5.1	Introduction	71
5.2	Smearing functions	73
5.3	Solving the wave equation	77
5.4	Black hole smearing functions and large angular momentum modes	79
5.4.1	Asymptotic behavior of the wave equation	80
5.4.2	Large angular momentum and WKB	81
5.5	Smearing functions for other spacetimes	82
5.5.1	Static spherically symmetric spacetimes	83
5.5.2	Trapped null geodesics	86
5.6	Conclusions	88
6	Scanning Tunneling Macroscopy, Black Holes, and AdS/CFT Bulk Local- ity	90
6.1	Introduction	90
6.2	Bulk reconstruction in Minkowski space	92
6.2.1	Success with propagating modes	92
6.2.2	Reality with evanescent modes	94
6.2.3	Ignoring evanescent modes	94
6.2.4	Dealing with evanescent modes	95
6.2.5	Window function	96
6.2.6	Conclusions	97
6.3	Bulk reconstruction in AdS space	97

6.3.1	Local CFT operators as boundary data	98
6.3.2	Bulk reconstruction near the boundary	99
6.3.3	Bulk reconstruction deeper in	104
6.4	Evanescence in CFT Dual	109
6.5	Discussion	111
6.6	Evanescent Optics	111
7	Entanglement Entropy: A Perturbative Calculation	115
7.1	Introduction	115
7.2	General framework	116
7.2.1	Geometric perturbations	118
7.2.2	Relevant perturbations	121
7.3	Perturbations of a planar entangling surface	122
7.3.1	Calculation	124
7.4	Perturbations of a spherical entangling surface	126
7.4.1	Geometric perturbations	128
7.4.2	Calculation	129
7.5	Notation	131
7.6	Foliation of \mathcal{M} in the vicinity of the entangling surface	132
7.7	Intermediate calculations for Sec. 7.3	134
II	Emergent Spacetime with $\Lambda > 0$	135
8	Boundary Definition of a Multiverse Measure	137
8.1	Introduction	137
8.2	New light-cone time	141
8.2.1	Probabilities from light-cone time	141
8.2.2	A gauge choice for the conformal boundary	142
8.2.3	The multiverse in Ricci gauge	144
8.3	Properties of new light-cone time	147
8.3.1	Rate equation	147
8.3.2	Attractor solution	150
8.3.3	Event counting	150
8.4	The probability measure	152
8.4.1	Probabilities from the new light-cone time cut-off	152
8.4.2	Equivalence to the old light-cone time cut-off	153
8.4.3	Equivalence to the new causal patch measure	154
8.4.4	Equivalence to the old causal patch measure	155
8.5	The general case	156
8.5.1	Perturbative inhomogeneities	156
8.5.2	Singularities in the boundary metric	158
8.5.3	Inhomogeneities	160
8.6	Conformal factor for a homogeneous bubble universe	161

9	Testing measures	166
9.1	Introduction	166
9.2	Counting Observations	169
9.3	Open FRW universes with cosmological constant	171
9.4	The causal patch cut-off	173
9.4.1	Positive cosmological constant	173
9.4.2	Negative cosmological constant	179
9.5	The apparent horizon cutoff	181
9.5.1	Definition	181
9.5.2	Positive cosmological constant	184
9.5.3	Negative cosmological constant	186
9.6	The fat geodesic cutoff	188
9.6.1	Positive cosmological constant	189
9.6.2	Negative cosmological constant	190
10	Future inextendability and non-decoupling of the regulator	193
10.1	Non-decoupling of the regulator	193
10.2	The probability to encounter the cutoff	197
10.3	Objections	199
10.3.1	The cutoff can not be encountered in a late-time limit	199
10.3.2	Hitting the cutoff is an artifact	201
10.4	The Guth-Vanchurin paradox	203
10.5	Discussion	203
10.5.1	Assumptions	204
10.5.2	Observation	205
10.5.3	Interpretation	206
11	Conclusion	209
	Bibliography	211

List of Figures

- 2.1 The interior of the larger cylinder represents AdS. The vertical direction is time and the radial direction is ρ . The CFT lives on the boundary ($\mathcal{R} \times S^{d-1}$) which is at $\rho = \pi/2$. In order for AdS/CFT to fully realize the holographic principle, some theory which has a Hilbert space of dimension $\exp(A/4l_{pl}^{d-1})$, where A is the area of the sphere at $\rho = \pi/2 - \delta$, must be extracted from the CFT. This theory would live on the sphere at $\rho = \pi/2 - \delta$ (inner cylinder) and would need to be able to fully describe the interior: $0 < \rho < \pi/2 - \delta$ 10
- 2.2 (a) A relativistic particle following a radial geodesic inside of AdS. Its closest approach to the boundary is $\rho = \pi/2 - m/E$. (b) The sphere shown is the S^{d-1} of the CFT (it is represented as an S^2 ; in (a) we were only able to draw the boundary as an S^1). On the CFT the particle is represented by a thin shell of $\langle T_{\mu\nu} \rangle$. This shell is shown at multiple instances of time. As the particle starts near the boundary, the shell is small and near the right pole of the sphere. The shell grows as the particle falls towards smaller ρ . As the particle passes through $\rho = 0$ the shell wraps the entire S^{d-1} . As the particle moves out again to larger ρ , the shell contracts. Crucially, the thickness of the shell is m/E 13
- 2.3 (a) The profile of a localized wave packet (2.10) traveling away from the boundary ($q_0 = 10^6$, $\sigma_x = \sigma_z = 10^{-3}$) shown at time $t = 10^{-2}$. It is composed of modes highly oscillatory in the z direction, of wavenumber peaked around q_0 . (b) The CFT image $\langle O(x) \rangle$ at this time (given by (2.20)). 14
- 2.4 A plot of the radial modes $f_{0l}(\rho)$ in (2.21) for fixed l ($l = 100$) and several different choices of Δ ($\Delta = 10, 100, 500$, in colors red, green, blue, respectively). The variable on the horizontal axis is the radial coordinate $r/L = \tan \rho$. The plot shows that increasing Δ leads to stronger confinement to the center of AdS. To violate UV/IR we choose a mode with $\Delta \gg l \gg 1/\delta$ 17
- 2.5 A highly boosted state for which the center of mass energy, and hence back-reaction, is small. In order for the theory on the sphere to only have a finite number of states, we must exclude this state. A plausible reason is that any attempt to probe it would form a black hole that is much larger than the sphere. 22

- 3.1 (a) The Poincaré patch of AdS, with the usual time slicing in the coordinates of Eq. (1.1). (b) Time slices of an arbitrary bulk coordinate system that covers the same near-boundary region as the Poincaré patch but a different region far from the boundary. This illustrates that there is no preferred coordinate system that would uniquely pick out a region described by the boundary, particularly if the bulk is not in the vacuum state. 26
- 3.2 The boundary of AdS; the dashed lines should be identified. Examples of globally hyperbolic subsets b are shown shaded. A causal diamond is a set of the form $I^-(q) \cup I^+(p)$, where q is boundary event in the future of the boundary event p . Let τ be the time along a geodesic from p to q in the Einstein static universe of unit radius ($ds^2 = -dt^2 + d\Omega_{d-1}^2$). With $\tau = 2\pi$, the causal diamond is the boundary of the Poincaré patch. A causal diamond with $\tau < 2\pi$ ($\tau > 2\pi$) is called “small” (“large”). An open interval (t_1, t_2) with $t_2 - t_1 < \pi$ ($t_2 - t_1 > \pi$) is called “short strip” (“tall strip”). 27
- 3.3 The four null hypersurfaces orthogonal to a spherical surface B in Minkowski space. The two cones F_1, F_3 have negative expansion and hence correspond to light-sheets. The covariant entropy bound states that the entropy of the matter on each light-sheet will not exceed the area of B . The other two families of light rays, F_2 and F_4 , generate the skirts drawn in thin outline. Their cross-sectional area is increasing, so they are not light-sheets, and the entropy of matter on them is unrelated to the area of B 30
- 3.4 Penrose diagram of a collapsing star (shaded). At late times, the area of the star’s surface becomes very small (B). The enclosed entropy in the *spatial* region V stays finite, so that the spacelike entropy bound is violated. The covariant entropy bound avoids this difficulty because only future directed light-sheets are allowed by the nonexpansion condition. L is truncated by the future singularity; it does not contain the entire star. 31
- 3.5 (a) A square system in 2+1 dimensions, surrounded by a surface B of almost vanishing length A . The entropy in the enclosed *spatial* volume can exceed A . (b) [Here the time dimension is projected out.] The light-sheet of B intersects only with a negligible (shaded) fraction of the system, so the covariant entropy bound is satisfied. 32
- 3.6 An AdS-Schwarzschild black hole. A sphere on the regulated boundary encloses an infinitely large spatial region that extends all the way to the second, disconnected conformal boundary on the far side of the black hole. 33

- 3.7 The union of all future directed light-sheets, $L^+(b)$, coming off the usual slicing of the boundary Minkowski space (left) covers precisely the Poincaré patch (the wedge-shaped region that lies both in the future of the boundary point A and in the past of D). On the right, we show a a different time slicing of the same boundary region. One of these slices is shown in blue in the bulk Penrose diagram (center); it curves up at B and down at C . The future-directed light-sheet coming off the portion of the slice near C is nearly the same as the future lightcone of C (shown in red/short-dashed), which reaches far beyond the Poincaré patch to the far side of AdS. The bulk region covered by $L^+(b)$ will thus be nearly two Poincaré patches, consisting of the points that lie in the future of A but not in the future of D (long-dashed). 35
- 3.8 Consider an arbitrary boundary region b (enclosed by the solid line), and a point P in the bulk region $C(b)$ (orthogonal to the page). The causal future of P , $J^+(P)$, intersected with b is shown hatched. Roughly, the strategy of the proof is to demonstrate that there exists a time slice $s_{t_*(P)}$ on the boundary that is tangent to the lower boundary of future of P in b . We show that $s_{t_*(P)}$ is the earliest time slice that has any intersection with the future of P , and that P lies on the past light-sheet of $s_{t_*(P)}$ 37
- 3.9 The covariant bulk RG flow presented here reproduces the standard bulk RG flow in certain coordinate systems. Here we illustrate the construction in global coordinates of Anti-de Sitter space. For a given coordinate time cutoff δ , the union over t of the intersection surfaces $s_t^\delta = \ell_{t+\delta/2}^- \cap \ell_{t-\delta/2}^+$ form a timelike hypersurface b^δ in the bulk (left). The cross-sectional area of a given light-sheet will be greater on the surface b^δ than on $b^{\delta'}$ (right). The difference $(A - A')/4$ bounds the entropy on the red light-sheets going from b^δ to $b^{\delta'}$, meaning that the bound applies to the entire darkly shaded wedge between them. The lightly shaded region between the hypersurfaces $b^{\delta'}$ and b^δ is covered by such wedges. 39
- 3.10 A cross-section of Anti-de Sitter space, showing a short strip region \mathcal{S} centered around $\tau = 0$ on the boundary, and the bulk region $S(\mathcal{S})$ spacelike separated from \mathcal{S} . A local operator at the origin of the bulk can be written in terms of local operators on the boundary smeared over the boundary region spacelike-related to the origin, within the green wedges. This region is much larger than \mathcal{S} (red thick line), stretching from $\tau = -\pi/2$ to $\tau = +\pi/2$. . . 42
- 3.11 According to the Hamiltonian on the boundary strip \mathcal{S} , no source acts at the origin in the bulk, so the expectation value of ϕ vanishes everywhere. At the time τ_0 outside the strip, a source term for the nonlocal boundary operator dual to $\phi(0,0)$ can be added to the boundary Hamiltonian. This causes the expectation value of ϕ to be nonzero in the future of $(0,0)$, in contradiction with the earlier conclusion about the same bulk points. Thus, unless we possess information about the exterior of \mathcal{S} on the boundary guaranteeing that such operators do not act, the bulk interpretation of regions outside $\bar{H}(\mathcal{S}) = C(\mathcal{S})$ is potentially ambiguous. 44

3.12	The shaded region shows bulk points spacelike related to a global boundary Cauchy surface σ . The union of all such sets over the collection of boundary Cauchy surfaces which do not intersect \mathcal{S} has an ambiguous bulk interpretation when the boundary Hamiltonian is allowed to vary outside of \mathcal{S} . The unambiguous region, $\bar{H}(\mathcal{S})$, is the complement of this union. In this example, we see that $\bar{H}(\mathcal{S}) = C(\mathcal{S})$	45
4.1	Here we show the AdS-Rindler wedge inside of global AdS, which can be defined as the intersection of the past of point A with the future of point B . The asymptotic boundary is the small causal diamond defined by points A and B . The past lightcone of A and the future lightcone of B intersect along the dashed line, which is a codimension-2 hyperboloid in the bulk. There is a second AdS-Rindler wedge, defined by the points antipodal to A and B , that is bounded by the same hyperboloid in the bulk. We refer to such a pair as the “right” and “left” AdS-Rindler wedges.	50
4.2	Here we depict Poincare-Milne space, together with an AdS-Rindler space that it contains. The bulk of Poincare-Milne can be defined as the intersection of the past of point A with the future of line BE . Clearly this region contains the AdS-Rindler space which is the intersection of the past of A and the future of B . Furthermore, the asymptotic boundary of the Poincare-Milne space and the AdS-Rindler space is identical, being the causal diamond defined by A and B on the boundary.	60
4.3	This is one of many null geodesics which passes through AdS-Rindler space without reaching the AdS-Rindler boundary. The four highlighted points on the trajectory are (bottom to top) its starting point on the near side of the global boundary, its intersection with the past Rindler horizon, its intersection with the future Rindler horizon, and its endpoint on the far side of the global boundary.	63
4.4	On the left we show the effective potential for a null geodesic in a spherical black hole, and on the right the same for a planar black brane. In the case of a spherical black hole, there is a potential barrier which traps some null geodesics in the $r < 3G_N M$ region. Therefore continuous reconstruction from the boundary is not possible for the region $r < 3G_N M$. In the planar case, there are null geodesics reach arbitrarily large finite r without making it to the boundary. Hence there is no bulk region which can be continuously reconstructed from the boundary data.	66
4.5	The geometry defining the Hartle-Hawking state for AdS-Rindler. Half of the Lorentzian geometry, containing the $t > 0$ portion of both the left and right AdS-Rindler spaces, is glued to half of the Euclidean geometry. The left and right sides are linked by the Euclidean geometry, and the result is that the state at $t = 0$ is entangled between the two halves.	68
4.6	The Penrose diagram for a hyperbolic black hole. In the $\mu = 0$ case, regions I and IV become the right and left AdS-Rindler wedges. In this case, the singularity is only a coordinate singularity, so the spacetime can be extended to global AdS.	69

- 5.1 To construct the bulk operator $\Phi(B)$, the CFT operator $O(b')$ is smeared with the smearing function $K(B|b')$ as indicated in (5.7). (a) The support of the pure AdS smearing function $K(B|b')$ is all boundary points b' space-like separated from B (hatched region). (b) Had the AdS-Rindler smearing function existed, it would have only made use of the boundary region that overlaps with $J^+(q) \cap J^-(p)$ (the intersection of the causal future of q and causal past of p), where q and p are chosen so that $J^+(q) \cap J^-(p)$ just barely contains B . Any changes outside this bulk region $J^+(q) \cap J^-(p)$ would have been manifestly irrelevant for computing $\Phi(B)$ 76
- 5.2 The wave equation can be recast as a Schrödinger equation (5.18). We plot the global AdS₄ potential (5.21) for $l = 3$ for a massless field. The plot on the left is in terms of the radial coordinate r appearing in the AdS metric (5.20). The plot on the right is in terms of the tortoise coordinate r_* , and is the one relevant for solving (5.18). The two are related through $r = \tan r_*$. The tortoise coordinate has the effect of compressing the potential at large r , while leaving small r unaffected. The AdS barrier occurs at r_* very close to $\pi/2$; its narrowness allows the modes to decay only as a power law: $\phi \sim r^{-\Delta}$. 78
- 5.3 A plot of the AdS₄-Schwarzschild ($r_0 = 1$) potential (5.23) as a function of the radial coordinate r . The plot on the left is for $l = 4$, and the one on the right for $l = 10$. Unlike for pure AdS, ω is not bounded from below by l ; for a fixed ω , l can be arbitrarily high. The barrier an ω mode must pass through grows as l increases. This results in the $\omega \ll l$ modes having exponential behavior in l . Intuitively, these modes become ever more confined near the horizon with increasing l 79
- 5.4 We are interested in finding for which static spherically spacetimes without horizons a smearing function exists. The smearing function involves the sum (5.31) over modes, which can be grouped into 3 different regimes. Only A poses a threat to the convergence of (5.31). At large r the metric, and consequently the potential (5.19), looks like that of pure AdS (5.38). At smaller r , in regime A the angular momentum l is so large that all terms in the potential except for the centrifugal barrier (5.40) are irrelevant. 83
- 5.5 The wave equation can be recast as a Schrödinger equation (5.18) with a potential $V(r)$ and an energy ω^2 . Here we sketch a possible potential (5.19) for which a smearing function doesn't exist. At large r , $r > R$, the the potential looks like that of pure AdS (the figure has been compressed; the distance between r_2 and R is really much larger). At smaller r the potential, for large l , is approximated by (5.40). If $f(r)/r^2$ ever has positive slope, as shown above, some of the modes ω (dashed line) will have to tunnel through the barrier. Consequently, the sum (5.31) will diverge for $r < r_2$ 84

5.6	The equation for a null geodesic is that of a particle traveling in a 1-d potential (5.44). The potential is plotted for (a) pure AdS and (b) AdS-Schwarzschild ($M = 1$ in AdS ₄). In pure AdS all null geodesics have an endpoint on the boundary, as can be seen from the figure on the left. This is in contrast to spacetime with horizons (right figure) which have some null geodesics which are trapped as a result of the potential U vanishing at the horizon. More generally, whenever there are trapped null geodesics, then there is no smearing function for some points in the bulk.	87
6.1	The setup of bulk reconstruction in $(2+1)$ -dimensional Minkowski half-space $R_+^{2,1}$. We have a bulk field $\phi(x, t, z)$ obeying the wave equation, which we wish to reconstruct from the data $\phi(x, t, 0)$ on a timelike hypersurface at $z = 0$.	92
6.2	The bulk profile of the modes is found from solving the Schrödinger equation in the above effective potential $V(z)$. Modes with $\omega^2 > k^2$ are propagating, whereas those with $\omega^2 < k^2$ are evanescent. In pure AdS space, evanescent modes are forbidden due to their exponential divergence at large z . However, evanescent modes become allowed if there is a change in the effective potential in the large z interior so as to cause the potential to dip below k^2	101
6.3	(a) The effective potential (6.45) for a small AdS black hole. The boundary of AdS space is at tortoise coordinate $r_* = \pi/2$, while the horizon is at $r_* = -\infty$. Evanescent modes arise if the potential ever drops below the value at the local minimum present in the pure AdS space ($\sim l^2$). Here, this occurs because the potential approaches zero at the horizon. (b) The effective potential for pure global AdS space.	105
6.4	The effective potential for a large black as a function of the tortoise coordinate. Note that with tortoise coordinates the region near the boundary (large r) gets compressed near $\pi/2$. Thus, the top corner of the plot is the only portion of the potential that is similar to a portion of the pure AdS potential (Fig. 6.3(b)). Here, the modes are mostly those of type (1) that propagate into the black hole horizon, and type (4) that have low energy and hence are evanescent.	107
6.5	The effective potential for the wave equation in a medium that undergoes a jump in its index of refraction.	112
6.6	Light is shined from the source (z_1), interacts with the sample (z_0), and is received at $z = 0$	113
6.7	The evanescent modes can be converted into propagating modes by placing, for example, a piece of glass near them.	114
7.1	Abstract sketch of the two dimensional transverse space to the entangling surface Σ . \mathcal{C}_\pm are the two sides of the cut \mathcal{C} where the values ϕ_\pm of the field ϕ are imposed.	117
7.2	A sketch of a slightly deformed entangling surface (curved line) in three dimensions. (x_1, x_2) span the transverse space to Σ , while y parametrizes Σ . The foliation (7.9) is designed to capture the geometry of the neighborhood of a given entangling surface Σ	122

7.3	Transverse space to the entangling surface in the analytically continued space-time. Σ is located at the origin. The reduced density matrix is given by a path integral (7.1) with fixed boundary conditions ϕ_+ (ϕ_-) on the upper (lower) dashed blue lines.	123
7.4	We conformally transform between \mathcal{H} (left) and R^d (right). We first map from the $\sigma \equiv u + i\tau$ coordinates of \mathcal{H} to $e^{-\sigma}$ (middle); here the origin is $u = \infty$ and the boundary circle is $u = 0$. We then map via (7.53) to R^d . Dashed lines on the left represent $\tau_E = 0^+, \beta^-$ slices of \mathcal{H} that are mapped through an intermediate step onto $t = 0^\pm$ sides of the cut throughout the interior of the sphere $r = R$ on the right	126
7.5	We show the constant τ_E slices (blue) and constant u slices (red) in the (r, t_E) plane (7.56). The sphere is located at $r/R = 1$, $t_E = 0$ and corresponds to $u \rightarrow \infty$. The vertical line ($r = 0$) corresponds to $u = 0$	127
8.1	Constant light-cone size on the boundary defines a hypersurface of constant “light-cone time” in the bulk. The green horizontal lines show two examples of such hypersurfaces. They constitute a preferred time foliation of the multiverse. In the multiverse, there are infinitely many events of both type 1 and type 2 (say, two different values of the cosmological constant measured by observers). Their relative probability is defined by computing the ratio of the number of occurrences of each event prior to the light-cone time t , in the limit as $t \rightarrow \infty$	138
8.2	Conformal diagram of de Sitter space with a single bubble nucleation. The parent de Sitter space is separated from the daughter universe by a domain wall, here approximated as a light-cone with zero initial radius (dashed line). There is a kink in the diagram where the domain wall meets future infinity. This diagram represents a portion of the Einstein static universe, but the Ricci scalar of the boundary metric is not constant.	139
8.3	(a) The parent de Sitter space with the future of the nucleation event removed and (b) the bubble universe are shown as separate conformal diagrams which are each portions of the Einstein static universe. After an additional conformal transformation of the bubble universe (shaded triangle), the diagrams can be smoothly matched along the domain wall. The resulting diagram (c) has a round S^3 as future infinity but is no longer a portion of the Einstein static universe.	146
8.4	Conformal diagram of de Sitter space containing a bubble universe with $\Lambda = 0$	158

- 9.1 The probability distribution over the timescales of curvature and vacuum domination at fixed observer timescale $\log t_{\text{obs}}$, before the prior distribution over $\log t_c$ and the finiteness of the landscape are taken into account. The arrows indicate directions of increasing probability. For $\Lambda > 0$ (a), the distribution is peaked along the degenerate half-lines forming the boundary between regions I and II and the boundary between regions IV and V. For $\Lambda < 0$ (b), the probability distribution exhibits a runaway toward the small t_c , large t_Λ regime of region II. The shaded region is excluded because $t_{\text{obs}} > t_f = \pi t_\Lambda$ is unphysical. 175
- 9.2 The causal patch can be characterized as the union of all past light-cones (all green lines, including dashed) of the events along a worldline (vertical line). The apparent horizon cutoff makes a further restriction to the portion of each past light-cone which is expanding toward the past (solid green lines). The dot on each light-cone marks the apparent horizon: the cross-section of maximum area, where expansion turns over to contraction. 182
- 9.3 Conformal diagrams showing the apparent horizon cutoff region. The boundary of the causal patch is shown as the past light-cone from a point on the conformal boundary. The domain wall surrounding a bubble universe is shown as the future light-cone of the bubble nucleation event. The region selected by the cutoff is shaded. For $\Lambda > 0$ (a), the boundary of the causal patch is always exterior to the apparent horizon. For $\Lambda < 0$ (b), the apparent horizon diverges at a finite time. Because the apparent horizon cutoff is constructed from light-cones, however, it remains finite. The upper portion of its boundary coincides with that of the causal patch. 183
- 9.4 The probability distribution from the apparent horizon cutoff. The arrows indicate directions of increasing probability. For $\Lambda > 0$ (a), the probability is maximal along the boundary between regions IV and V before a prior distribution over $\log t_c$ is included. Assuming that large values of t_c are disfavored, this leads to the prediction $\log t_\Lambda \sim \log t_c \sim \log t_{\text{obs}}$. For $\Lambda < 0$ (b), the distribution is dominated by a runaway toward small t_c and large t_Λ along the boundary between regions II and III. 185
- 9.5 The probability distribution computed from the scale factor (fat geodesic) cutoff. The arrows indicate directions of increasing probability. For $\Lambda > 0$ (a), the probability distribution is maximal along the boundary between regions IV and V; with a mild prior favoring smaller $\log t_c$, this leads to the prediction of a nearly flat universe with $\log t_c \sim \log t_\Lambda \sim \log t_{\text{obs}}$. For $\Lambda < 0$ (b), the probability distribution diverges as the cosmological constant increases to a value that allows the observer timescale to coincide with the big crunch. . . 190

- 10.1 A multiverse populated by infinitely many observers (vertical line segments) who first see 1 o'clock (at events labeled "1") and then 2 o'clock ("2"). A geometric cutoff selects a finite set of events, whose relative frequency defines probabilities. Events that are not counted are indicated by dashed lines. The left figure shows a global cutoff: all events prior to the time t_0 (curved line) are counted and all later events ignored. (The global time has nothing to do with the observers' clocks, which come into being at a rate dictated by the dynamics of eternal inflation.) The right figure shows the causal patch cutoff, which restricts to the causal past of a point on the future boundary. In both figures, the cutoff region contains observers who see 1 o'clock but not 2 o'clock. Their number, as a fraction of all observers who see 1 o'clock before the cutoff, defines the probability of reaching the end of time between 1 and 2 o'clock. 194

List of Tables

- 8.1 Equivalences between measures. The new light-cone time cut-off is equivalent to probabilities computed from three other measure prescriptions (double arrows). This implies that all four measures shown on this table are equivalent in the approximation described at the beginning of Sec. 2.3. (See Ref. [44] for a more general proof of the equivalence between the old light-cone time cut-off and the old causal patch measure.) 152

Acknowledgments

I am grateful to many for discussions. I would like to especially thank my advisor Raphael Bousso, as well as Ben Freivogel, Stefan Leichenauer, Soo-Jong Rey, and Misha Smolkin.

Chapter 1

Introduction

One of the outstanding problems in theoretical physics has been how to quantize gravity. Early attempts at quantizing gravity tried to quantize the gravitational field using the same procedure that had been successfully used to quantize all other fields. These attempts failed for a number of technical and conceptual reasons. The reason for this failure is now understood at a much more basic level: the gravitational field is not fundamental. Quantizing gravity is like trying to quantize sound waves in a fluid - one is quantizing the wrong thing. What has been learned is that spacetime is an emergent concept. The purpose of this thesis is to further our understanding of how spacetime emerges.

Early Hints of the Holographic Principle

That spacetime is not fundamental but rather emergent is a radical idea and one which emerged after decades of studies of black holes and string theory. In normal contexts, gravity is weak and so our lack of an understanding of quantum gravity is not a hindrance. However, in situations when matter is made sufficiently dense, gravity can become strong. The most dramatic signature of gravity is a black hole. A black hole forms when an object becomes sufficiently dense to induce complete gravitational collapse. Hence, it is the arena of black holes that can be expected to reveal some of quantum gravity's central features.

The study of the quantum properties of black holes developed rapidly in the 1970s. Hawking's area theorem [92] demonstrated that the area of a black hole never decreases, giving a resemblance to the second law of thermodynamics in which the entropy never decreases. Going further, Bardeen, Carter and Hawking [8] developed the four laws of black hole thermodynamics, demonstrating a striking analogy between the behavior of a black hole and a thermodynamic system. At a mathematical level, it appeared that if one identified the black hole's horizon area with entropy and the inverse mass of the black hole as temperature, then Einstein's equations for the black hole's behavior could be recast as the standard equations for the laws of thermodynamics. At this point, as the name suggests, a black hole was believed to be black. Consequently, it was thought it had no temperature and no entropy, and that the analogy with thermodynamics was purely a mathematical one.

At the same time, Bekenstein [10, 11, 12, 13] was disturbed that as a result of a black hole the concept of the second law of thermodynamics would lose its operational meaning. A black hole gave one a mechanism to destroy entropy - one simply drops it into

the black hole. The need to uphold the second law motivated Bekenstein to propose that black holes do in fact have an entropy, and it is given by their area in Planck units. The statement that emerged is one of the generalized second law: the matter entropy plus the black hole entropy can never decrease.

In 1975, Hawking discovered [93, 94] that black holes are not truly black but rather very slowly emit radiation. This discovery set off the study of the quantum properties of black holes. What had merely been a mathematical analogy had now become a reality: black holes are true thermodynamic systems with an entropy and a temperature.

Bekenstein had considered thought experiments in which an entropic object was dropped into a black hole. In 1994 Susskind [166] considered a different thought experiment. Any static object can be converted into a black hole - one simply adds matter to it. Yet, once we have a black hole we know that its entropy is given by its area in Planck units. The second law of thermodynamics forces us to conclude that the object we initially had must have also had an entropy bound by its area in Planck units. What seemed to be emerging then was the following concept: gravity produces fundamental limitations on how densely information can be packed. If we try to squeeze too much entropy into too small of a region, the whole system will collapse into a black hole.

't Hooft and Susskind went further [172, 166]. If the entropy of any object is bound by its area, should it not be the case that the correct fundamental theory is one whose number of degrees of freedom scales as an area rather than an entropy? Without the black hole thought experiments, this seems completely counter-intuitive. The world we observe is local, and degrees of freedom at different locations are independent. Indeed, quantum field theory is a local theory and the number of degrees of freedom scales as a volume. On the other hand, it appears that with gravity present most of these degrees of freedom are fictitious - one can not actually excite them because a black hole forms if one tries. This practical limitation was used as motivation for the much more radical conjecture of 't Hooft and Susskind: the world is a hologram, and the fundamental description lives on the boundary. This holographic principle was a major step towards the demise of the concept of spacetime being fundamental.

Emergent Negative Cosmological Constant Spacetime: the AdS/CFT correspondence

In 1997 the Holographic Principle was realized by Maldacena within string theory for a particular class of spacetimes [127]. The spacetimes for which it was realized are those that are asymptotically Anti-de Sitter (AdS). AdS spacetimes are those with negative cosmological constant ($\Lambda < 0$). AdS space can be thought of as particular kind of confining box: near the center of the space, the cosmological constant is not felt and a probe particle behaves as if it were in flat space. If one gives the particle some energy so that it moves outward, it will eventually begin to feel the cosmological constant and will be attracted back towards the center. In this way, massive particles moving along geodesics in AdS will oscillate about the origin. The boundary of AdS, which is a sphere, is therefore inaccessible to massive particles, although it can be reached by massless particles.

The Anti-de Sitter/ Conformal Field Theory (AdS/CFT) correspondence makes the following statement: the full string theory in AdS space is exactly dual to a particular

field theory that can be thought of as living on the sphere that is at the boundary of AdS. AdS/CFT is perhaps the simplest and cleanest way to realize holography. An important question is to what extent AdS/CFT has fully realized the holographic principle. In particular, the holographic principle states that the theory describing the interior of a sphere lives on the surface of the sphere and also has a Hilbert space dimension set by the area of the sphere. AdS/CFT has clearly realized the first of these conditions because the CFT lives on the boundary. On the other hand, the later condition is more difficult since the boundary of AdS space has infinite area. This is consistent with the CFT having an infinite dimensional Hilbert space, but not very informative for trying to match the number of degrees of freedom. In chapter 2 we address this question of to what extent AdS/CFT has fully realized holography.

Before the AdS/CFT correspondence, string theory provided a quantum theory of gravity, but only in the perturbative regime. To address many of the fundamental questions in quantum gravity, one needs to know the nonperturbative structure of theory. The AdS/CFT correspondence both provided this nonperturbative definition, and at the same time realized and confirmed the Holographic Principle.

Although AdS/CFT was discovered 15 years ago, many aspects of the dictionary relating bulk observables to boundary observables have remained elusive. It will be our goal to make progress in filling in this dictionary. In particular, one of the central mysteries of holography is if spacetime is emergent and not fundamental, how does it do such a good job of emerging so that we perceive a local world?

One way to make progress towards understanding how the bulk emerges is to pose the question: is it the case that the CFT on some subset of the boundary is dual to some subregion of the bulk? This question is intimately tied to the question of how nonlocal holography is. An affirmative answer to our question is not guaranteed, but if it were the case that portions of the boundary could be associated with portions of the bulk in some way, it would not only provide an important ingredient in the bulk-boundary dictionary, but would also give us clues for how to apply holography in other contexts. In chapter 3 we first establish what the candidate bulk and boundary subregions would actually be. Then, in chapter 4 we try to establish to what extent the subregion duality can be realized.

An even more elementary question is how mysterious is holography. In particular, can one give a simple prescription for which CFT quantities encode the various aspects of the bulk? One can ask: if we consider a family of bulk observers in AdS who are near the boundary, can they do experiments to reconstruct the bulk? A priori, it is unclear if the answer should be positive or negative. By the nature of the AdS/CFT dictionary, the extent to which these observers can reconstruct the bulk is the same as the question of the extent to which local CFT operators are enough to reconstruct the bulk. One should recall that a field theory has many quantities which are inherently nonlocal (for instance, entanglement entropies or Wilson loops). Consequently it is important to establish when local CFT operators are sufficient to reconstruct the bulk, and when these more exotic nonlocal objects need to be used. In chapter 5 and 6 we turn to this question. Finally, in contexts that local CFT operators are not sufficient to reconstruct the bulk, one would like to understand which nonlocal objects one should use instead. Entanglement entropy has emerged as a particularly good candidate due to the work of Ryu and Takayanagi [154, 104]

relating boundary entanglement entropy to the area of extremal bulk curves. If the Ryu - Takayanagi conjecture is true, it gives a remarkable way in which to probe the bulk. To test their conjecture one needs to know how to compute entanglement entropy within the CFT, and this is what we turn to in chapter 7.

Emergent Positive Cosmological Constant Spacetime

While AdS/CFT is a beautiful realization of holography, the spacetime we inhabit has a positive cosmological constant rather than a negative one. What we would really like then would be to have a holographic description of a de Sitter spacetime (dS). This is a far more challenging task. Indeed, it is even unclear where the holographic theory should live, as the full de Sitter space is itself spatially closed. Although, a natural possibility is the hologram lives on the spatial sphere at future infinity.

The main difficulty, however, to understanding emergent $\Lambda > 0$ spacetimes is the difficulty of realizing them in string theory. In particular, the same arguments leading to the AdS/CFT correspondence are not applicable to a putative dS/CFT correspondence [163, 164, 186, 131]. Nevertheless, one could conjecture there exists a CFT dual to de Sitter space and then try to understand this correspondence without knowing the details of what the actual boundary theory is. Unfortunately, it is unclear if such a correspondence can be realized [168].

The potential absence of a holographic duality for pure stable de Sitter space may mean we must directly find the holographic dual of the full multiverse. It is believed that the effective potential within string theory is one with an enormous number of metastable minima. One therefore has not just one universe, but an infinite number of them. One can start with a region that is de Sitter, yet tunneling events will cause it to decay to other universes which can have either positive or negative Λ . The structure of the theory on future infinity is therefore a fractal rather than just a sphere. One is forced to define a theory on this fractal. In chapter 8 we take some first steps towards doing this.

Unlike the case of AdS/CFT where we know the theory that is dual to string theory in AdS space, we do not yet know what the theory is that is dual to the multiverse. At this stage, one of the hopes is to extract pieces of the bulk-boundary mapping in AdS/CFT and apply them to the multiverse context. Concretely, in AdS/CFT (as discussed in chapter 2) when comparing bulk quantities with CFT quantities one must typically impose an IR regulator in the bulk and correspondingly a UV regulator on the boundary. The AdS/CFT dictionary relates these regulators. In the context of a multiverse duality, a UV regulator on the theory on future infinity would correspond to a late time regulator in the bulk.

The most basic question that needs to be answered is how to relate this boundary UV regulator to a bulk late time regulator. In the context of the multiverse, the question of the late time regulator is known as the measure problem, and predates holography. While in the context of the multiverse we do not have a theory telling us what the correct regulator is, we instead have experimental observations. The hope then is to work in the reverse direction - we guess a bulk measure and see what predictions it gives us. The requirement that predictions match observation then gives us clues about what the correct analog of UV/IR is for the multiverse. In chapter 9 we test various measure proposals to see how well their predictions match observation. Finally, in chapter 10 we point out a serious difficulty

that is encountered with measures in the multiverse - the bulk late time regulator (measure) never decouples. As a field theory statement for the putative theory on future infinity, this means the UV cutoff never decouples.

Part I

Emergent Spacetime with $\Lambda < 0$: AdS/CFT

In this first part we study the holographic principle as realized in the AdS/CFT correspondence. This part is devoted to studying to what extent AdS/CFT realizes holography, what the dictionary is between the bulk theory and the boundary theory, and in what way bulk locality emerges. This part is organized as follows.

In chapter 2, based on [152], we study to what extent AdS/CFT has realized the holographic principle. The holographic principle asserts that the complete description of the interior of a sphere is a theory which not only lives on the surface of the sphere, but also has $A/4$ binary degrees of freedom. In this context we revisit the question of UV/IR that was initially studied by Susskind and Witten [165]. We construct states which are localized deep in the interior yet are encoded on short scales in the CFT, seemingly in conflict with the UV/IR prescription. We make a proposal to address the more basic question of which CFT states are sufficient to describe physics within a certain region of the bulk.

In chapter 3, based on [40], we ask whether the CFT restricted to a subset b of the AdS boundary has a well-defined dual restricted to a subset $H(b)$ of the bulk geometry. The Poincaré patch is an example, but more general choices of b can be considered. We propose a geometric construction of H . We argue that H should contain the set C of causal curves with both endpoints on b . Yet H should not reach so far from the boundary that the CFT has insufficient degrees of freedom to describe it. This can be guaranteed by constructing a superset L of H from light-sheets off boundary slices and invoking the covariant entropy bound in the bulk. The simplest covariant choice is $L = L^+ \cap L^-$, where L^+ (L^-) is the union of all future-directed (past-directed) light-sheets. We prove that $C = L$, so the holographic domain is completely determined by our assumptions: $H = C = L$. In situations where local bulk operators can be constructed on b , H is closely related to the set of bulk points where this construction remains unambiguous under modifications of the CFT Hamiltonian outside of b . Our construction leads to a covariant geometric RG flow. We comment on the description of black hole interiors and cosmological regions via AdS/CFT.

In chapter 4, based on [31], we investigate the nature of the AdS/CFT duality between a subregion of the bulk and its boundary. In global AdS/CFT in the classical $G_N = 0$ limit, the duality reduces to a boundary value problem that can be solved by restricting to one-point functions of local operators in the CFT. We show that the solution of this boundary value problem depends continuously on the CFT data. In contrast, the AdS-Rindler subregion cannot be continuously reconstructed from local CFT data restricted to the associated boundary region. Motivated by related results in the mathematics literature, we posit that a continuous bulk reconstruction is only possible when every null geodesic in a given bulk subregion has an endpoint on the associated boundary subregion. This suggests that a subregion duality for AdS-Rindler, if it exists, must involve nonlocal CFT operators in an essential way.

In chapter 5, based on [118], we study more generally the bulk - boundary mapping. In Lorentzian AdS/CFT there exists a mapping between local bulk operators and nonlocal CFT operators. In global AdS this mapping can be found through use of bulk equations of motion and allows the nonlocal CFT operator to be expressed as a local operator smeared over a range of positions and times. We argue that such a construction is not possible if there are bulk normal modes with exponentially small near boundary imprint. We show that the AdS-Schwarzschild background is such a case, with the horizon intro-

ducing modes with angular momentum much larger than frequency, causing them to be trapped by the centrifugal barrier. More generally, we argue that any barrier in the radial effective potential which prevents null geodesics from reaching the boundary will lead to modes with vanishingly small near boundary imprint, thereby obstructing the existence of a smearing function. While one may have thought the bulk-boundary dictionary for low curvature regions, such as the exterior of a black hole, should be as in empty AdS, our results demonstrate otherwise.

In chapter 6, based on [149], we further study the bulk -boundary dictionary and bulk reconstruction from boundary data. We establish resolution bounds on reconstructing a bulk field from boundary data on a timelike hypersurface. If the bulk only supports propagating modes, reconstruction is complete. If the bulk supports evanescent modes, local reconstruction is not achievable unless one has exponential precision in knowledge of the boundary data. Without exponential precision, for a Minkowski bulk, one can reconstruct a spatially coarse-grained bulk field, but only out to a depth set by the coarse-graining scale. For an asymptotically AdS bulk, reconstruction is limited to a spatial coarse-graining proper distance set by the AdS scale. AdS black holes admit evanescent modes. We study the resolution bound in the large AdS black hole background and provide a dual CFT interpretation. Our results demonstrate that, if there is a black hole of any size in the bulk, then sub-AdS bulk locality is no longer well-encoded in boundary data in terms of local CFT operators. Specifically, in order to probe the bulk on sub-AdS scales using only boundary data in terms of local operators, one must either have such data to exponential precision or make further assumptions about the bulk state.

In chapter 7, based on [153], we would like to start using the entanglement entropy as a probe of bulk geometry. We therefore develop a framework for computing entanglement entropy for the CFT. We provide a framework for a perturbative evaluation of the reduced density matrix. The method is based on a path integral in the analytically continued space-time. It suggests an alternative to the holographic and ‘standard’ replica trick calculations of entanglement entropy. We implement this method within solvable field theory examples to evaluate leading order corrections induced by small perturbations in the geometry of the background and entangling surface. Our findings are in accord with Solodukhin’s formula for the universal term of entanglement entropy for four dimensional CFTs.

Chapter 2

Holography for a Small World

2.1 Introduction


The spherical entropy bound [14, 15] gives a remarkable bound on the number of states contained within a sphere of area A . The holographic principle [172, 166, 20, 23] builds on it to make the far more extraordinary assertion that there exists a theory that lives on the sphere and describes all of physics within the sphere, and accomplishes this feat while only having a Hilbert space of dimension $\exp(A/4l_{pl}^2)$.

AdS/CFT has only partially realized the holographic principle. The CFT does live on the boundary of AdS, but its Hilbert space is infinite-dimensional. The boundary of the region described has an area which is also infinite, and comparing two infinities is not meaningful.

To confirm the holographic principle in AdS we must extract from the CFT a theory of the appropriate dimension that is capable of fully describing the innermost region of AdS, out to a sphere of area A (see Fig. 2.1). The CFT on a lattice is a natural candidate, and was the basis of the UV/IR proposal of Susskind and Witten [165]. However, the UV/IR proposal is known to have some limitations, motivating us to analyze its validity in a range of contexts.

In the first part of this chapter we will construct certain states which explicitly violate UV/IR; the bulk states will be within the sphere of area A , yet their boundary image will have features on scales far smaller than the lattice spacing. Our examples will demonstrate UV/IR is violated for both particles which stay inside the sphere and those which enter and leave; for particles that are both relativistic and non-relativistic; for both static states and dynamic states.

In the second part of the chapter we start afresh in looking for the correct theory describing the interior of the sphere of area A . We make a proposal for an answer to the question: what CFT states should this theory contain? This is a much easier question than constructing a full theory; a full theory requires having not only the states, but also the observables and a way of doing time evolution. It is, however, a more difficult question than finding just the states confined within the sphere that the Bekenstein bound specializes to. If the interior of the sphere is truly a holographic image, the hologram must be able to describe relativistic particles which enter and leave the sphere. We will discuss how our



(a)

Figure 2.1: The interior of the larger cylinder represents AdS. The vertical direction is time and the radial direction is ρ . The CFT lives on the boundary ($\mathcal{R} \times S^{d-1}$) which is at $\rho = \pi/2$. In order for AdS/CFT to fully realize the holographic principle, some theory which has a Hilbert space of dimension $\exp(A/4l_{pl}^{d-1})$, where A is the area of the sphere at $\rho = \pi/2 - \delta$, must be extracted from the CFT. This theory would live on the sphere at $\rho = \pi/2 - \delta$ (inner cylinder) and would need to be able to fully describe the interior: $0 < \rho < \pi/2 - \delta$.

proposal avoids the difficulties UV/IR encountered. We also comment on the necessity of excluding ultraboosted states from the description.

The chapter is organized as follows. In Sec. 2.2.1 we review the UV/IR proposal. In Sec. 7.2 we review the well-known example of a relativistic particle oscillating inside of AdS. The boundary image of the bulk gravitational field induced by the particle is the energy-momentum tensor concentrated to a thin shell, and consequently strongly violating UV/IR. In Sec. 2.4.1 we begin our study of scalar field wave packets. We construct a well localized and highly relativistic wave packet which goes inside the bulk IR cutoff, yet on the boundary the expectation value of the CFT operator dual to the scalar field is localized to a region well below the lattice spacing. In Sec. 7.4 we consider a scalar field with large mass (in AdS units), and consider a mode with angular momentum much larger than the inverse of the lattice spacing but much smaller than the mass. The mode is localized within the central AdS radius, yet the boundary image is on scales below the lattice spacing. In Sec. 2.4.3 we consider a general solution of the Klein-Gordon equation and show that in order for the CFT to not lose information about it when placed on a lattice would require a lattice spacing far smaller than the one prescribed by UV/IR. In Sec. 2.5 we discuss the possibility that the information contained in local CFT operators which went missing when a UV cutoff was placed is in fact retained in some other “precursor” operators. In Sec. 2.6 we present our proposal for which CFT states are sufficient to fully describe the interior of the sphere of area A .

The question of the validity of the “scale/radius” relation is related but somewhat

different from the question of the validity of the UV/IR proposal. As a side note, in Sec. 2.2.2 we comment that scale/radius does not follow from the rescaling isometry of the Poincare patch metric, nor is it generally valid. Our example in Sec. 2.4.1 explicitly violates scale/radius.

2.2 UV/IR

2.2.1 Review of UV/IR

In this section we review the UV/IR proposal [165]. Consider AdS in global coordinates

$$ds^2 = \frac{L^2}{\cos^2 \rho} (-d\tau^2 + d\rho^2 + \sin^2 \rho d\Omega_{d-1}^2). \quad (2.1)$$

The CFT lives on the boundary of AdS, at $\rho = \pi/2$, on a sphere that has been conformally rescaled to have radius 1. The UV/IR prescription seeks to provide a theory that can describe the interior of AdS for all $\rho < \pi/2 - \delta$, where $\delta \ll 1$. The full CFT is of course capable of doing this, however it has an infinite-dimensional Hilbert space and the holographic principle tells us a finite-dimensional Hilbert space should suffice.

In many computations in AdS/CFT, a bulk quantity that is IR divergent is dual to a CFT quantity that is UV divergent. For instance, the divergence of the length of a string ending on the boundary is dual to the divergent self-energy of a quark. This observation motivated [165] to propose that the theory we are looking for is the CFT placed on a lattice. Since the CFT is an $SU(N)$ gauge theory, [165] wanted to count N^2 degrees of freedom per lattice site. The lattice spacing is then fixed by having the number of CFT degrees of freedom match the area in Planck units of the sphere at $\rho = \pi/2 - \delta$:

$$\frac{1}{l_{pl}^{d-1}} \left(\frac{L}{\delta} \right)^{d-1}. \quad (2.2)$$

Using the relation $N^2 = (L/l_{pl})^{d-1}$, we see that the lattice size is fixed to be δ .

Since each of these degrees of freedom has, like a harmonic oscillator, an infinite-dimensional Hilbert space, [165] needed to further impose that each oscillator can only be excited to the first few energy levels. This amounts to imposing an energy density cutoff of $(1/\delta)^d$. The energy density cutoff will not be important for us since UV/IR will face difficulties already at the stage of the spatial lattice.

We should note that the terminology ‘‘UV/IR’’ is used in a range of contexts. For us, UV/IR will mean the specific proposal reviewed above of how to truncate the CFT and still be able to describe the portion of the interior, $0 < \rho < \pi/2 - \delta$.

2.2.2 Comments on Scale/Radius

In this section we include a few comments on the ‘‘scale/radius’’ relation and how it relates to UV/IR. The scale/radius relation is the statement that an object close to the boundary should be dual to a CFT state with small spatial extent, whereas an object deep in the bulk should be dual to a CFT state of large spatial extent.

The degree to which the scale/radius relation is generally valid is somewhat orthogonal to the one of UV/IR we are interested in. In particular, that a CFT state may have a certain spatial extent does little in terms of telling us on what scale the state has features, and hence what kind of lattice on the CFT would be sufficient to accurately describe this state.

In fact, we see no basis for even scale/radius being valid in any dynamical context. The scale/radius relation is motivated by the isometry of AdS under the rescaling of Poincare coordinates: $(t, x, z) \rightarrow (\lambda t, \lambda x, \lambda z)$. Consider some bulk solution; for instance a solution $\Phi_0(x, t, z)$ of the scalar wave equation. This configuration will have a boundary imprint $\phi_0(x, t)$ which is obtained by extracting the coefficient of the decaying tail of Φ_0 near the boundary. The AdS/CFT dictionary tells us $\phi_0(x, t)$ is equal to $\langle O(x, t) \rangle$ on the CFT. The isometry of AdS means one can construct a one-parameter family of bulk solutions $\Phi_0(\lambda x, \lambda z, \lambda t)$ and these will have a boundary imprint $\phi_0(\lambda x, \lambda t)$. Thus as λ is increased, Φ_0 will be peaked deeper in the bulk and the boundary imprint grows in spatial extent by a factor λ .

However, for the scale/radius relation to be relevant in explaining the emergent radial direction as an energy scale in the CFT, it would need to be true as a dynamical statement. AdS/CFT is a complete duality; not only can a bulk state be mapped to a boundary state, but also the equivalence must be maintained under time evolution. So it is not sufficient to show that processes characteristically closer to the the boundary have a smaller boundary size at some point in their evolution. One must show that throughout the evolution of a localized bulk object, its radial location is correlated with the spatial extent of the CFT image. The isometry of AdS under rescalings only implies the former and not the later.

Our example in Sec. 7.2 of an oscillating particle (and the growing/contracting shell to which it is dual) is generally regarded as consistent with scale/radius. However, this is true in a trivial way: the oscillating particle is obtained from the static one at $\rho = 0$ through a scale transformation in Poincare coordinates. Our example in Sec. 2.4.1 will explicitly violate scale/radius.

2.3 Oscillating Particle

In this section we review the example of a relativistic particle oscillating inside of AdS (see Fig. 2.2). The backreaction of the particle changes the metric and induces a nonzero $\langle T_{\mu\nu} \rangle$ on the boundary. As the particle passes through the center, $\langle T_{\mu\nu} \rangle$ remains concentrated on a thin shell. It was pointed out in [148] that this example is in tension with UV/IR.

It will be convenient to consider AdS in coordinates

$$ds^2 = - \left(1 + \frac{r^2}{L^2} \right) d\tau^2 + \frac{dr^2}{1 + \frac{r^2}{L^2}} + r^2 d\Omega_{d-1}^2. \quad (2.3)$$

These coordinates are related to the (2.1) coordinates by $r = L \tan \rho$. The boundary of AdS




Figure 2.2: (a) A relativistic particle following a radial geodesic inside of AdS. Its closest approach to the boundary is $\rho = \pi/2 - m/E$. (b) The sphere shown is the S^{d-1} of the CFT (it is represented as an S^2 ; in (a) we were only able to draw the boundary as an S^1). On the CFT the particle is represented by a thin shell of $\langle T_{\mu\nu} \rangle$. This shell is shown at multiple instances of time. As the particle starts near the boundary, the shell is small and near the right pole of the sphere. The shell grows as the particle falls towards smaller ρ . As the particle passes through $\rho = 0$ the shell wraps the entire S^{d-1} . As the particle moves out again to larger ρ , the shell contracts. Crucially, the thickness of the shell is m/E .

is at $r \rightarrow \infty$. In this limit the metric (2.3) asymptotes to

$$ds^2 = r^2 \left(-\frac{d\tau^2}{L^2} + d\Omega_{d-1}^2 \right). \quad (2.4)$$

The metric of the $S^{d-1} \times \mathcal{R}$ on which the CFT lives is obtained by conformally rescaling (2.4) by a factor of $1/r^2$, giving a sphere of radius equal to 1.

A particle of mass m oscillating in AdS (Fig. 2.2a) satisfies the geodesic equation

$$\dot{r}^2 = \left(\frac{E}{m} \right)^2 - 1 - \left(\frac{r}{L} \right)^2, \quad (2.5)$$

where E is the energy of the particle with respect to the timelike Killing field,

$$E/m = (1 + r^2/L^2)\dot{\tau}. \quad (2.6)$$

The proper energy of the particle at the center of AdS is equal to E , and the CFT energy of this state is EL . The largest r the particle reaches is $r_{max} \approx LE/m$, where we have assumed $E \gg m$. In terms of coordinates (2.1), $\rho_{max} = \pi/2 - \alpha$, where we defined $\alpha \equiv m/E$. Since the particle is relativistic, $\alpha \ll 1$.

The computation of $\langle T_{\mu\nu} \rangle$ for this state was done by Horowitz and Itzhaki [99], and we collect their results in Appendix 2.8. In Fig. 2.2b we have sketched how $\langle T_{\mu\nu} \rangle$ evolves. The energy of the CFT state is concentrated on a shell of thickness α . As the particle falls into the bulk the shell expands, reaching a maximum size when the particle reaches $r = 0$. The shell then contracts as the particle moves out towards larger r . Crucially, the thickness of the shell is equal to α .




Figure 2.3: (a) The profile of a localized wave packet (2.10) traveling away from the boundary ($q_0 = 10^6$, $\sigma_x = \sigma_z = 10^{-3}$) shown at time $t = 10^{-2}$. It is composed of modes highly oscillatory in the z direction, of wavenumber peaked around q_0 . (b) The CFT image $\langle O(x) \rangle$ at this time (given by (2.20)).

UV/IR tells us that the CFT on a lattice with spacing δ should describe the bulk out to $\rho = \pi/2 - \delta$. Thus, it must describe the oscillating particle which goes through ρ smaller than the cutoff. Yet, if we choose $\alpha \ll \delta$, then the shell has a width far smaller than the lattice spacing. The extent to which the cutoff CFT fails to describe the particle can be made arbitrarily large by making α small. The extreme case would be a massless particle traveling through the bulk. Its boundary dual is a shell that is completely localized on the lightcone $\theta = \tau$.

2.4 Scalar Field Solutions

The example in Sec. 7.2 of an oscillating particle presents a constraint on imposing any kind of lattice on the CFT. However, this example is a bit special and we would like to have a larger set of examples to test UV/IR. That is what we do in this section. Instead of working with the gravitational field, we will consider a free scalar field ϕ ,

$$(\square + m^2)\phi = 0. \quad (2.7)$$

The CFT operator O is dual to ϕ . Throughout this section we will consider some solutions of (2.7) and look at $\langle O \rangle$ for these states. We will be interested in solutions of (2.7) which either at some time, or for all time, are contained within the bulk region $\rho < \pi/2 - \delta$. In our examples we will construct states that do this and also have $\langle O \rangle$ that is concentrated on scales much less than δ .

2.4.1 Relativistic Wave Packet

In this section we would like to construct a wave packet that travels from near the boundary of AdS into the bulk. We would like this packet to remain well concentrated and

have negligible spread as it propagates into the bulk. It is familiar from wave physics that packets with large momentum in one direction have, for a long period of time, negligible spreading in the transverse directions. In AdS we can construct packets with a similar property (shown in Fig. 2.3a). For these packets we will want to find how $\langle O \rangle$ evolves with time. Near the boundary ($\rho \approx \pi/2$), the field value will decay as $(\cos \rho)^\Delta$. The AdS/CFT dictionary tell us that the coefficient in front of this term determines $\langle O \rangle$. So to find $\langle O \rangle$ we need to find the behavior of the tail of the wave packet at $\rho \approx \pi/2$. It should be pointed out that just because the peak of the packet in the bulk has negligible transverse spread does not yet imply the tail will also have negligible spread with time; although we will see this is what occurs in AdS.

The packets we will construct can be used to violate UV/IR by arbitrarily large amounts. For some given lattice spacing δ , we make a packet with transverse spread much less than δ . We then make it sufficiently energetic so that it doesn't spread for a long time. On the CFT, $\langle O \rangle$ remains concentrated to a region much less than δ , even when the packet is at $\rho < \pi/2 - \delta$.

To construct the packets it will be convenient to use Poincare coordinates, which are a good approximation to global coordinates near the boundary and for small angular spread. Inserting $\rho = \pi/2 - z$ into (2.1) gives

$$ds^2 = \frac{L^2}{z^2} (-dt^2 + dz^2 + dx^2), \quad (2.8)$$

where the angular coordinate is now x and the time coordinate has been relabeled $t = \tau$. In these coordinates the mode solutions to (2.7) are

$$\varphi_{qk}(x, t, z) = z^{d/2} J_\nu(qz) e^{ikx - i\omega_{qk}t}, \quad (2.9)$$

where $0 < q < \infty$, the energy is $\omega_{qk}^2 = q^2 + k^2$, and $\nu = \Delta - d/2$. The conformal dimension Δ is taken to be of order 1, and is related to the mass m through $\Delta(\Delta - d) = m^2 L^2$.

We now construct our packet out of the modes (2.8) peaked around a large $q = q_0$ with spread in q of σ_z^{-1} , and peaked around $k = 0$ with spread in k of σ_x^{-1} :

$$\Phi(x, t, z) = \int dq dk \sqrt{q} \varphi_{qk}(t, x, z) e^{-k^2 \sigma_x^2/4} e^{-(q-q_0)^2 \sigma_z^2/4}, \quad (2.10)$$

where $q_0 \gg \sigma_z^{-1}, \sigma_x^{-1} \gg 1$. To verify the packet has the behavior we desire, we make some simplifications to allow us to evaluate the integrals in (2.10). First we notice that particles in AdS only behave differently from those in Minkowski space when they get close enough to the boundary such that the gravitational potential energy becomes comparable to their kinetic energy. In terms of the modes (2.9) this is reflected in the the z component:

$$J_\nu(qz) \sim \frac{1}{\sqrt{qz}} e^{iqz}, \quad \text{for } qz \gg \nu. \quad (2.11)$$

This means we can think of q as a kind of radial momentum. From the form of (2.10) we see that q is peaked around q_0 with spread σ_z^{-1} which is much less than q_0 . Thus for $z \gg \nu/q_0$ we can approximate (2.10) by

$$\Phi(x, t, z) \approx z^{(d-1)/2} \int dq dk e^{iqz} e^{-k^2 \sigma_x^2/4} e^{ikx} e^{-(q-q_0)^2 \sigma_z^2/4} e^{-i\omega_{qk}t}. \quad (2.12)$$

Aside from the uninteresting power of z in front, (2.12) is of the same form as a packet propagating in the z direction in Minkowski space. Since in (2.12), $k \lesssim \sigma_x^{-1} \ll q_0$, we approximate

$$\omega_{qk} = \sqrt{q^2 + k^2} \approx q + \frac{k^2}{2q_0}. \quad (2.13)$$

This allows us to separate the q and k integrals in (2.12), and easily evaluate the integral over q ,

$$\Phi(x, t, z) = z^{(d-1)/2} \psi(x, t) e^{-(z-t)^2/\sigma_z^2} e^{iq_0(z-t)}, \quad (2.14)$$

where

$$\psi(x, t) = \int dk e^{ikx} e^{-k^2\sigma_x^2/4} e^{-i\frac{k^2 t}{2q_0}}. \quad (2.15)$$

We have dropped constants and have labeled the transverse spread as ψ because, as can be seen from the energy (2.13), it satisfies the non-relativistic Schrodinger equation for a particle of mass q_0 . Evaluating (2.15) thus gives the familiar answer,

$$\psi(x, t) = \frac{1}{\sigma_x \sqrt{1 + \frac{2it}{\sigma_x^2 q_0}}} \exp\left(-\frac{x^2}{\sigma_x^2 \left(1 + \frac{2it}{\sigma_x^2 q_0}\right)}\right), \quad (2.16)$$

showing the packet's spread with time is the expected,

$$\sigma_x(t) = \sigma_x \sqrt{1 + \left(\frac{2t}{q_0 \sigma_x^2}\right)^2}. \quad (2.17)$$

This shows that the transverse spread is negligible for times $t \lesssim q_0 \sigma_x^2$.

We would now like to evaluate $\langle O(x, t) \rangle$. Noting that for small z the Bessel function can be approximated as

$$J_\nu(qz) \sim (qz)^\nu \quad \text{for } qz < \nu, \quad (2.18)$$

we obtain $\langle O(x, t) \rangle$ from the $z \rightarrow 0$ limit of $z^{-\Delta} \Phi(x, t, z)$. Using (2.10) we find,

$$\langle O(x, t) \rangle = \int dq dk q^{\nu+1/2} e^{-k^2\sigma_x^2/4} e^{ikx} e^{-(q-q_0)^2\sigma_z^2/4} e^{-i\omega_{qk}t}. \quad (2.19)$$

We now perform the same simplifications we used on (2.10), separating the integrals over q and k , and finding

$$\langle O(x, t) \rangle = e^{-iq_0 t} \psi(x, t), \quad (2.20)$$

where $\psi(x, t)$ is given by the same expression as (2.15) before, and we have ignored terms that are independent of x, t . Eq. 2.20 shows that $\langle O(x, t) \rangle$ remains localized within $|x| \lesssim \sigma_x$ for a time $t \lesssim q_0 \sigma_x^2$, as well as being highly oscillatory with time.¹ A wave packet and its

¹Since our interest is the CFT defined on the sphere of the global AdS boundary, our equations should only be used for small times, $t \ll 1$, where the approximation of the Poincare patch metric is valid. To follow the packet deeper into the bulk we would need to use the true global AdS evolution. Since it will take the packet a time of $\pi/2$ to reach the center of AdS, we expect its transverse spread even at the center to remain comparable to σ_x (with a sufficiently large choice of q_0). It is interesting to consider what would happen if we were actually interested in the CFT theory on the Minkowski boundary of the Poincare patch. Then as the packet moves towards the Poincare horizon located at infinite z , the CFT profile would spread over the entire x axis. This is not surprising; in this case the packet is traveling for an infinite amount of Poincare time.




Figure 2.4: A plot of the radial modes $f_{0l}(\rho)$ in (2.21) for fixed l ($l = 100$) and several different choices of Δ ($\Delta = 10, 100, 500$, in colors red, green, blue, respectively). The variable on the horizontal axis is the radial coordinate $r/L = \tan \rho$. The plot shows that increasing Δ leads to stronger confinement to the center of AdS. To violate UV/IR we choose a mode with $\Delta \gg l \gg 1/\delta$.

dual $\langle O \rangle$ at an instant of time are shown in Fig. 2.3. We see that the UV/IR prescription can be violated by arbitrarily large amounts with this example. To decrease the boundary size, σ_x , we simply need to increase q_0 in order to maintain $q_0 \gg \sigma_x^{-1}$.

Since we are taking q_0 to be large, one might worry that if there is other matter around then gravitational effects might invalidate the use of the free wave equation (2.7). However, any such effects would occur on scales set by the Planck scale. The Planck scale is related to the AdS scale by a power of $1/N$. Crucially, the lattice size on the boundary that UV/IR prescribes is independent of N , allowing us to take N as large as we want.

2.4.2 Non-relativistic Mode

In Sec. 7.2 and Sec. 2.4.1 we considered relativistic particles with large momentum in the radial direction. In this section we go to the other extreme of a non-relativistic particle with angular momentum. The radial location of the particle is determined by two competing forces. AdS is confining, and the more massive the particle, the more it is confined to the center of AdS. On the other hand, the centrifugal barrier from the particle's angular momentum pushes it out to larger radius. We will see that in order for the boundary image to have features on small scales, the particle needs to be given large angular momentum. By dialing the particle's mass to be sufficiently large, we make the particle well-confined to the center of AdS, and hence violate UV/IR.

As in the previous section, we consider a scalar field and for simplicity our equations are in AdS_3 . The mode solutions of the Klein-Gordon equation (2.7) in global coordinates (2.1) are given by,

$$\varphi_{nl}(\rho, \tau, \theta) = f_{nl}(\rho) e^{il\theta} e^{-i\omega_{nl}\tau}. \quad (2.21)$$

The energy $\omega_{nl} = \Delta + 2n + l$, where $\Delta \approx mL$ for large mass, and the radial part $f_{nl}(\rho)$ is some hypergeometric function whose explicit form we have suppressed. For small ρ AdS looks

like Minkowski space in spherical coordinates, and the radial modes reduce to the familiar spherical Bessel functions.

Consider a bulk solution that consists of just a single mode, (2.21). On the boundary, the expectation value of O for this state is simply proportional to (2.21) with the radial portion removed,

$$\langle O(\theta, \tau) \rangle \propto e^{il\theta} e^{-i\omega_{nl}\tau}. \quad (2.22)$$

Let us now put the CFT on a lattice of size δ . Since the CFT lives on a circle of radius 1, this amounts to restricting to modes with $l \lesssim \delta^{-1}$. Thus, to violate UV/IR we simply need to consider a bulk solution φ_{0l} with $l \gg \delta^{-1}$ and take the mass of the field sufficiently large to have this mode confined near the center. In other words, $\Delta \gg l \gg \delta^{-1}$. Since this example consists of an energy eigenstate, it illustrates that UV/IR is insufficient even for describing static bulk physics which remains confined within $\rho < \pi/2 - \delta$.

2.4.3 General Solution

We have seen examples of both relativistic and non-relativistic wave packets that cause difficulties for UV/IR. But perhaps these examples are special and for a generic bulk field configuration, the UV regulated CFT will be able to describe it fairly well. This is the question we address in this section.

Consider some solution $\Phi(\rho, \tau, \theta)$ of the wave equation (2.7). Expanding Φ in terms of modes (2.21),

$$\Phi(\rho, \tau, \theta) = \sum_{n,l} c_{nl} f_{nl}(\rho) e^{il\theta - i\omega_{nl}\tau}. \quad (2.23)$$

Assuming the modes are normalized with respect to the standard bulk Klein-Gordon norm, near the boundary $f_{nl}(\rho)$ takes the form $d_{nl}(\cos \rho)^\Delta$, where

$$d_{nl} = (-1)^n \sqrt{\frac{\Gamma(\Delta + n + |l|)\Gamma(\Delta + n)}{n!\Gamma(\Delta)^2\Gamma(n + |l| + 1)}}. \quad (2.24)$$

We can now find the boundary imprint of this solution,

$$\langle O(\tau, \theta) \rangle = \sum_{n,l} c_{nl} d_{nl} e^{il\theta - i\omega_{nl}\tau}. \quad (2.25)$$

The question now is: for a solution Φ relevant to the region $\rho < \pi/2 - \delta$, are the important features of the corresponding $\langle O \rangle$ on scales larger than δ ? On the boundary, a spatial lattice of size δ is equivalent to cutoff on l of δ^{-1} . In order to be able to truncate the sum in (2.25) to $l < \delta^{-1}$, it would need to be the case that c_{nl} is suppressed at these large l . However, this is clearly not generally true. At any fixed ρ in the bulk we need a complete basis of modes in the angular direction, which requires retaining modes of arbitrarily high l .

To be more precise, since we don't expect local field theory to be valid in the bulk below the Planck scale, the highest l we would actually need in order to describe $r < R$ would be R/l_{pl} . However, this would require a lattice spacing on the CFT that is far smaller than in the UV/IR prescription: in $\text{AdS}_{d+1}/\text{CFT}_d$, the lattice we would need to describe $\rho < \pi/2 - \delta$ would need to have a spacing of $\delta/N^{2/(d-1)}$.

2.5 Precursors

We have seen examples of states for which the CFT would lose some of the information contained in $\langle O \rangle$ if it were placed on a lattice. However, perhaps all the information about the state is still present in the CFT on a lattice, but instead of being contained in $\langle O \rangle$ is contained in the expectation value of some other operators? Without a better understanding of the AdS/CFT dictionary this remains a possibility, but one that is difficult to test. In this section we will just make two comments. The first is that we see no compelling reason why upon placing the CFT on a lattice, $\langle O \rangle$ would lose information while these other operators retain it. The second is that if this possibility is realized, then the UV/IR regulated version of AdS/CFT may not be very useful since the aspects of the bulk-boundary dictionary that we knew in the full AdS/CFT would not be applicable to the regulated version.

To discuss how the CFT encodes the bulk it is useful to ask the following question: how would a collection of CFT observers reconstruct some bulk field configuration $\Phi(\rho, \tau, \theta)$? Expanding Φ in terms of modes φ_{nl} as in (2.23), the goal of the CFT observers would be to determine all the coefficients c_{nl} . Measuring the expectation value of O over the whole sphere at one instant of time, $\langle O(\theta, \tau_0) \rangle$, would give them some of the information, but not all of it. In particular, since the boundary imprint in $\langle O \rangle$ of a single mode is proportional to $\exp(i l \theta - i \omega_{nl} \tau_0)$, the CFT observers would be able to distinguish among the different l quantum numbers. However, all the modes φ_{nl} of a fixed l but differing n would give the same imprint in $\langle O(\theta, \tau_0) \rangle$. By measuring $\langle O(\theta, \tau) \rangle$ over a range of times, the CFT observers could start distinguishing the modes with different n . Having $\langle O(\theta, \tau) \rangle$ for all θ and a time of $\Delta\tau = \pi$ would allow for a full reconstruction of Φ .

If AdS/CFT is complete then there should be a faster way of reconstructing the bulk; there should exist some CFT operators, named “precursors” in [148], which can be measured at one instant of time and immediately fully determine Φ . The precursor operators are expected to be highly nonlocal and their actual form has remained a mystery. In certain dynamical contexts precursors are essential. For instance, if two wave packets collide in the center of AdS, it takes a time of $\pi/2$ before causality allows the result of the collision to propagate to the boundary and become encoded in $\langle O \rangle$. During that $\pi/2$ interval of time, the result of the collision is encoded exclusively in the precursors.

These considerations lead us to believe that if we were to consider the CFT on some subset of the boundary, $\mathcal{R} \times S^{d-1}$, that does not contain the full time direction, then the only way to reconstruct some portions of the bulk would be by measuring precursors. However, this is not the situation UV/IR presents: the UV regulated CFT in UV/IR still lives on the full time direction.² Thus it would be unclear why precursors are essential in the UV/IR regulated version of AdS/CFT in contexts where they were not needed in the full AdS/CFT.

Nevertheless, it may be that precursors save UV/IR. In this case, we should ask what the bulk-boundary dictionary would be for the UV/IR regulated AdS/CFT. In the

²In UV/IR an energy density cutoff is prescribed for the CFT in addition to the spatial lattice. The energy cutoff can be regarded as making the time direction into a lattice. Thus, if in our examples UV/IR had failed due to the energy density cutoff, then the need for precursors might have been a good explanation. However, UV/IR failed already at the stage of imposing the spatial lattice.

full AdS/CFT one can construct a mapping, perturbatively in $1/N$, between local bulk operators $\hat{\phi}(\rho, \tau, \theta)$ and the CFT operator O smeared over space and time with some kernel. To leading order in $1/N$, this takes the form

$$\hat{\phi}(\rho, \tau, \theta) = \int d\tau' d\theta' K(\rho, \tau, \theta | \tau', \theta') O(\tau', \theta'). \quad (2.26)$$

The actual form of K was worked out in [90, 89], where the authors also pointed out that a UV cutoff on the CFT would destroy the mapping (2.26). Namely, [90] showed that the divergence in the two-point function of $\hat{\phi}$ at coincident points can only occur from the UV divergence of the two-point function of O . Thus, if the theory dual to the inner portion of AdS is the CFT with a UV cutoff, then the only way to express the bulk operators $\hat{\phi}$ in terms of CFT operators is directly in terms of the precursors.

2.6 A Proposal for Finding the States

We are interested in finding a theory which is capable of describing all the physics that can occur in the interior of a sphere of area A in AdS, while having a Hilbert space of the appropriate dimension. The proposal the UV/IR prescription gave is that the theory describing the interior $\rho < \pi/2 - \delta$ is the CFT on a lattice of size δ . We have seen in previous sections that this proposal faces difficulties. In this section we initiate a new search for the desired theory.

Finding such a theory is a difficult problem, and we will only try to address a more elementary question: what are the CFT states that this theory would contain? In other words, what is the set of CFT states that are sufficient to cover everything that could possibly happen in the region $0 < \rho < \rho_0$?

We need to know what CFT states are dual to a chair in the center of AdS, or a star, or any object contained within $\rho < \rho_0$. But we also need much more than that. The sphere at $\rho = \rho_0$ that we are discussing is imaginary; we aren't literally putting a shell there. So a particle can fly in and out of the region $\rho < \rho_0$, or two particles can come in, collide, and leave. We must include the CFT states dual to these processes. This is an essential point: if this "small world" consisting of $\rho < \rho_0$ is truly a holographic image, every physical process which occurs inside should be encoded in the hologram. The difficult part is that we should choose carefully which CFT states we are keeping, since the holographic principle only allows us to include $\exp(A/4l_{pl}^{d-1})$ independent states. We should note that there is no guarantee the question being posed has an answer. It may just be that one really needs more states, even an infinite number, and there is no hologram for a small world.

We begin this section by reviewing the Bekenstein bound. We then make our proposal for what the CFT states are that we want, and discuss how it passes the tests of the previous sections that UV/IR struggled with. We then discuss a puzzling aspect of our proposal, and indeed any proposal which keeps a finite number of states: ultraboosted states which pass through $\rho < \rho_0$ must be excluded from the description.

2.6.1 Proposal

Consider some weakly gravitating object that has energy E and can be enclosed in a sphere of radius R .³ For instance, the object could be a planet, or a gas inside of a spherical cavity with reflecting walls. The Bekenstein bound gives a bound on the entropy of the object,

$$S \leq 2\pi ER . \quad (2.27)$$

The requirement that the object not be within its Schwarzschild radius requires $E < R/2l_{pl}^2$, and transforms (2.27) into

$$S \leq \frac{A}{4l_{pl}^2} . \quad (2.28)$$

Eq. (2.28) can be interpreted as saying that if we have a spherical box with enclosing walls, then specifying $A/4l_{pl}^2$ numbers is sufficient to completely specify what is inside the box.

We now return to AdS, and consider placing one of these bound systems inside of AdS. Since AdS is confining, placing the system anywhere except the center will incur an energy cost. If we do place it in the center then in order for it to not undergo gravitational collapse, its energy must be bound by its size. These considerations motivate our proposal for the answer to the question of what CFT states are in the holographic theory. In coordinates (2.1), with the AdS radius L , and with the CFT living on a sphere of radius 1,

Proposal: The CFT states sufficient to fully describe the bulk for all $\rho < \rho_0$ are those with CFT energy less than ML , where M is the mass of a black hole of radius ρ_0 .

Let us discuss a few aspects of our proposal. First, we would like to make sure it covers all the states we considered in previous sections. The non-relativistic particle with angular momentum (Sec. 7.4) stays confined within $\rho < \rho_0$ and so is the kind of bound state the Bekenstein bound includes, and is easily covered. In Sec. 7.2 and Sec. 2.4.1 we considered a relativistic particle which enters and leaves our sphere. UV/IR had difficulty with it even in the case the particle's energy is measured in AdS units, regardless of how small l_{pl} is. In our proposal, the energy cutoff is in Planck units, and so these relativistic particles are included.

We also need to verify that our proposal has the correct number of states prescribed by the holographic principle. This follows from consistency of AdS/CFT. On the CFT side we want to count the number of states with energy less than ML . This will be dominated by states with energy close to ML , and up to order 1 factors, can be calculated from the standard thermodynamics of a free gas of N^2 species. For CFT_4 for large M , the log of the number of such states is of order $N^{1/2}(ML)^{3/4}$. We need to compare this with the area in Planck units of a large AdS_5 black hole of mass M . This is given by $(Ml_{pl}^3 L^2)^{3/4}/l_{pl}^3$. Using the relation $N^2 = (L/l_{pl})^3$, we see they are equal.⁴

³In this paragraph we are in 4-dimensional Minkowski space.

⁴Although our proposal should be applicable even for small spheres with sub-AdS radius, it is not particularly useful in this case since the number of CFT states is larger than what we want. In particular, there will be states which are dual to graviton energy eigenstates that are delocalized over the entire central AdS region.




Figure 2.5: A highly boosted state for which the center of mass energy, and hence backreaction, is small. In order for the theory on the sphere to only have a finite number of states, we must exclude this state. A plausible reason is that any attempt to probe it would form a black hole that is much larger than the sphere.

We note that our proposal is more concrete than the heuristic statement that CFT states of low energy are associated with bulk regions of small ρ and states of high energy are associated with the near boundary region, $\rho \approx \pi/2$. In particular, we are not proposing that CFT states of high energy are sufficient to describe the portion of the bulk that is near the boundary. This would obviously be wrong; a relativistic particle oscillating in AdS, like the one in Sec. 7.2, transverses nearly all of AdS while having constant energy. All we are saying is that the low energy states may be sufficient to completely describe the small ρ region.

2.6.2 Ultraboosted states

Having discussed the successes of our proposal, in this section we push it to the brink of failure. Namely, our proposal excludes the high energy states (those with energy greater than ML), and we need to know if this is valid. It is certainly justified to exclude a state if it has large backreaction, and this is the first avenue we pursue. Our discussion will focus on the two different frames of reference that are relevant: that of the sphere on which the theory lives, and the center of mass frame.

Consider the state inside the sphere at some instant of time as an object is passing through. The energy, E , with respect to the reference frame of the sphere is the energy that appears in our proposal as the cutoff energy determining which states are kept. There is also the center of mass reference frame, and we will let E_{CM} denote the energy in this frame. The backreaction, as characterized by curvature invariants, can be computed in any frame of reference and is easiest to analyze in the center of mass frame.

There are therefore two different possibilities: (a) E and E_{CM} are comparable. (b) $E \gg E_{CM}$. In Case (a) our proposal does well, since the states of large energy also have

large center of mass energy E_{CM} and consequently collapse into a black hole larger than our sphere. On the other hand, Case (b) (see Fig. 2.5) is more interesting. By taking some system of small energy and boosting it by an enormous amount, its energy E can be made arbitrarily large. By going to sufficiently large boost, we can make states which the theory living on the sphere can no longer describe.

Although it appears odd that sufficiently energetic states are invisible to the theory on the sphere, it is perhaps reasonable. The observables of the theory should correspond to some kind of physical probe of the interior. For instance, the sphere theory could send in particles to probe the interior. These probes will naturally be of low energy, as measured in the sphere reference frame. Now consider the state of an ultraboosted object passing through the interior. Any attempt to probe it in this way will form a large black hole, of radius much larger than where the sphere theory was living. So, in a physical sense, there is no way for the hologram to be able to describe these states.

2.7 Conclusion

Our goal has been to see the extent to which UV/IR gives us a theory, with $A/4l_{pl}^{d-1}$ degrees of freedom, that can fully describe the central AdS region out to a sphere of area A . We presented some examples which are in tension with UV/IR. We therefore took a step back, turning to the more basic question of which CFT states are sufficient to describe the interior of the sphere. The states of lowest energy seem like a promising candidate, which passes the tests UV/IR struggled with. Interestingly, an ultraboosted particle passing through would be invisible to the theory on the sphere.

One may wonder: is there any reason a region of spacetime should be described by a holographic theory with so few degrees of freedom? The most straightforward interpretation of entropy bounds would be that they do nothing more than quantify the simple observation that information requires energy, and too much energy confined to too small of a region leads to gravitational collapse. From this perspective, the holographic principle is an unbelievable extrapolation. And yet, AdS/CFT has partially realized the holographic principle; so perhaps it fully realizes it. Or, perhaps it doesn't and there only exist holograms with infinite information content.

2.8 Equations for Energy Shell

In this appendix we collect several equations from [99] relevant for finding $\langle T_{\mu\nu} \rangle$ for the oscillating particle in Sec. 7.2. The result was already sketched in Fig. 2.2 b.

To find $\langle T_{\mu\nu} \rangle$ we need to first find the gravitational field caused by the oscillating particle. The tail of this field at large r will then be proportional to $\langle T_{\mu\nu} \rangle$. Horowitz and Itzhaki [99] did this calculation in a more elegant way. They first found $\langle T_{\mu\nu} \rangle$ for a stationary particle at $r = 0$ and then applied a boost. Their answer is presented in Poincare patch coordinates, since in these coordinates a boost is just a dilatation and easy to implement. The Poincare patch metric is

$$ds^2 = \frac{L^2}{z^2} (-dt^2 + dx^2 + dz^2). \quad (2.29)$$

Defining $u = t - \sqrt{x^2}$ and $v = t + \sqrt{x^2}$, the result of [99] for AdS_5/CFT_4 is

$$\langle T_{uu} \rangle = \frac{8m\alpha^4}{3\pi^2} \frac{1}{(\alpha^2 + u^2)^3(\alpha^2 + v^2)} \quad (2.30)$$

$$\langle T_{vv} \rangle = \frac{8m\alpha^4}{3\pi^2} \frac{1}{(\alpha^2 + v^2)^3(\alpha^2 + u^2)} \quad (2.31)$$

$$\langle T_{uv} \rangle = \frac{4m\alpha^4}{3\pi^2} \frac{1}{(\alpha^2 + u^2)^2(\alpha^2 + v^2)^2}. \quad (2.32)$$

In Poincare patch coordinates, the particle starts at $z = \alpha$ at $t = 0$ and then falls towards larger z . For fixed v , $\langle T_{\mu\nu} \rangle$ can be seen to be peaked on the lightcone ($u = 0$) with width of order α .

Since our real interest is in the boundary theory on $\mathcal{R} \times S^3$, we need to do a conformal transformation and coordinate change to get from $\mathcal{R}^{3,1}$ to $\mathcal{R} \times S^3$. The Minkowski metric is

$$ds^2 = -dudv + \frac{(v-u)^2}{4} d\Omega_2^2. \quad (2.33)$$

Conformally rescaling by a factor of $\frac{1}{4}(1+v^2)(1+u^2)$ and changing coordinates: $u = \tan U$, $v = \tan V$, where $U = \frac{1}{2}(\tau - \theta)$ and $V = \frac{1}{2}(\tau + \theta)$ gives the desired

$$ds^2 = -d\tau^2 + (d\theta^2 + \sin^2 \theta d\Omega_2^2). \quad (2.34)$$

Applying these transformations to the $\mathcal{R}^{3,1}$ stress tensor (2.30) gives the $\mathcal{R} \times S^3$ stress tensor:

$$\langle T_{UU} \rangle = \frac{2m\alpha^4}{3\pi^2} \frac{(1+u^2)^3(1+v^2)}{(\alpha^2 + u^2)^3(\alpha^2 + v^2)} \quad (2.35)$$

$$\langle T_{VV} \rangle = \frac{2m\alpha^4}{3\pi^2} \frac{(1+v^2)^3(1+u^2)}{(\alpha^2 + v^2)^3(\alpha^2 + u^2)} \quad (2.36)$$

$$\langle T_{UV} \rangle = \frac{m\alpha^4}{3\pi^2} \frac{(1+u^2)^2(1+v^2)^2}{(\alpha^2 + u^2)^2(\alpha^2 + v^2)^2}, \quad (2.37)$$

where u and v were written above in terms of τ and θ . If we are interested in the energy component, then this is given by

$$\langle T_{\tau\tau} \rangle = \frac{1}{4}(\langle T_{UU} \rangle + 2\langle T_{UV} \rangle + \langle T_{VV} \rangle). \quad (2.38)$$

Chapter 3

Light-sheets and AdS/CFT

3.1 Introduction

In the study of the Lorentzian AdS/CFT correspondence [127, 185, 6, 89, 91], one may consider the boundary theory defined on a proper subset of the global conformal boundary. In this case, it is natural to expect that the bulk dual spacetime manifold may be extendible. That is, the theory on the subset should describe less than the maximally extended global bulk solution.

The most common example is the CFT on Minkowski space, $\mathbf{R}^{d-1,1}$, which is dual to the Poincaré patch of the bulk, with metric

$$ds^2 = R^2 \frac{-dt^2 + d\vec{x}_{d-1}^2 + dz^2}{z^2} . \quad (3.1)$$

These bulk and boundary regions are shown as subsets of global AdS in Fig. 3.1a. The Poincaré patch seems a “natural” choice for the bulk, at least in the above warped-product coordinate system, where the boundary corresponds to $z = 0$ and the Poincaré patch to $z > 0$. However, there is nothing special about this choice of coordinates. One could easily write down coordinates that cover a larger or smaller portion of the bulk that is bounded by the same portion of the conformal boundary as $z \rightarrow 0$ (see Fig. 3.1b). So what selects the Poincaré patch as the bulk dual of the CFT on Minkowski space?

The bulk dual region should be well-defined not only in the vacuum, but for arbitrary states in the semiclassical regime, perturbatively in $1/N$. Deep in the bulk, the metric will not be that of empty AdS space. In general all Killing symmetries will be broken, so they cannot be used to pick out a preferred bulk region. For example, consider a global bulk solution corresponding to a pair of neutron stars orbiting around the origin of the standard global coordinates. At what time (say, along its worldline) does each star enter and exit the bulk region dual to the CFT on $\mathbf{R}^{d-1,1}$?

Another well-known example is the maximally extended Schwarzschild-AdS black hole. The global conformal boundary consists of two disconnected copies of $\mathbf{R} \times \mathbf{S}^{d-1}$. In the Hartle-Hawking state (the Euclidean vacuum), the two components can be thought of as slices in a single complex manifold and are related by analyticity. However, one can consider more general states, for example by adding neutron stars near the left boundary




Figure 3.1: (a) The Poincaré patch of AdS, with the usual time slicing in the coordinates of Eq. (1.1). (b) Time slices of an arbitrary bulk coordinate system that covers the same near-boundary region as the Poincaré patch but a different region far from the boundary. This illustrates that there is no preferred coordinate system that would uniquely pick out a region described by the boundary, particularly if the bulk is not in the vacuum state.

(far from the black hole), and white dwarfs near the right boundary. Restricting attention to the CFT living on the left boundary, one would expect it to encode the nearby neutron stars, but not the white dwarfs on the far side of the black hole. But what is the basis of this expectation? And where does the reach of the left theory end: at the black hole horizon, or somewhere inside the black hole/white hole regions? Again, there should be an answer to this question that does not rely on special bulk symmetries or coordinate choices.

There are many other possible choices of subsets b of the global boundary, some of which are shown in Fig. 3.2.¹ Let us suppose that the CFT on b describes some portion $H(b)$ of the bulk. We will call $H(b)$ the holographic domain of b . It should be possible to construct this bulk region geometrically from b . The goal of this paper is to provide such a construction.²

¹We consider only subsets b which, viewed as manifolds in their own right, are globally hyperbolic and have the same number of spacetime dimensions as the global boundary. Otherwise the initial value problem of the CFT would be ill-defined. If b is a proper subset of the global boundary, then it is not obvious that the CFT on b must have a bulk dual, and we do not set out to prove this or establish under which conditions a bulk dual exists. The question we seek to address is: if b did have a bulk dual region, what would it be? In all cases, a semiclassical bulk dual can be assigned only to some subset of CFT states (excluding, for example, the equal superposition of two global CFT states corresponding to different bulk metrics), and only perturbatively in $1/N$. In the case of a small diamond, the bulk dual is the AdS-Rindler patch, so one would expect the geometric states to have a thermal character.

²In the special case that b is a diamond, motivated by the Ryu-Takayanagi proposal [154, 173, 68] for computing holographic entanglement entropy, one could try constructing $H(b)$ out of minimal surface which start and end on b . We will not explore this approach here, however it would be interesting to study how it relates to our construction. We will note however that Ryu-Takayanagi only applies for static situations;




Figure 3.2: The boundary of AdS; the dashed lines should be identified. Examples of globally hyperbolic subsets b are shown shaded. A causal diamond is a set of the form $I^-(q) \cup I^+(p)$, where q is boundary event in the future of the boundary event p . Let τ be the time along a geodesic from p to q in the Einstein static universe of unit radius ($ds^2 = -dt^2 + d\Omega_{d-1}^2$). With $\tau = 2\pi$, the causal diamond is the boundary of the Poincaré patch. A causal diamond with $\tau < 2\pi$ ($\tau > 2\pi$) is called “small” (“large”). An open interval (t_1, t_2) with $t_2 - t_1 < \pi$ ($t_2 - t_1 > \pi$) is called “short strip” (“tall strip”).

Outline Our strategy will be to bound the bulk dual region $H(b)$ by a subset $C(b)$ and a superset $L(b)$:

$$C(b) \subset H(b) \subset L(b) . \quad (3.2)$$

We choose $C(b)$ to be the minimum region the CFT needs to describe to be consistent with bulk causality and locality properties near the boundary. The set $L(b)$ will be constructed from light-sheets in a way that ensures that the bulk dual does not have more degrees of freedom (higher maximum entropy) than the CFT. We will then show that under some additional assumptions on $L(b)$, all three sets are equal. This fully determines $H(b)$ subject to the stated assumptions.

In Sec. 3.2, we propose a lower bound on the holographic domain. We argue that $H(b) \supset C(b)$, where $C(b) = J^+(b) \cap J^-(b)$ is the set of bulk points which lie both in the causal past and in the causal future of the conformal boundary region b . Otherwise, the boundary would fail to describe regions that can be explored by a bulk observer localized near the boundary.

In Sec. 3.3, we propose an upper bound on the holographic domain: $H(b) \subset L(b)$,

for time dependent situations one must use the proposal [104] of (minimal) extremal surfaces, which has been less well established. Additionally, as we will find in Section 6, allowing sufficient modifications of the boundary theory makes the bulk ambiguous in large regions. In this sense, the bulk dual naturally constructed from extremal surfaces would generally be too large.

determined by the requirement that the boundary theory should contain enough degrees of freedom to describe the bulk. The number of CFT degrees of freedom is controlled by the area A of the regulated boundary [165]. In order to ensure that the CFT can describe the holographic domain H , H must be contained in a bulk region that satisfies an entropy bound $S \lesssim A$. In Sec. 3.3.1, we review the covariant entropy bound, which states that the relevant entropy lies on *light-sheets* emanating from the boundary area. There are no generally valid bounds on the entropy at equal time, so light-sheets should play a preferred role in any holographic duality. A light-sheet is a null hypersurface generated by nonexpanding light rays orthogonal to a codimension two surface.

The boundary is of codimension one, so the construction of a bulk region out of light-sheets off the boundary first requires a foliation of the boundary b into time slices (which are codimension two surfaces from the bulk viewpoint). In Sec. 3.3.2, we examine whether it is nevertheless possible to define a covariant, slicing-independent upper bound $L(b)$ from light-sheets.³

Given a boundary time slicing, one may consider the union $L^+(b)$ of all future-directed light-sheets (one from each slice), or the union $L^-(b)$ of all past-directed light-sheets. However, both of these sets depend on the slicing of b , and the same is true for the union $L^+(b) \cup L^-(b)$. Thus, neither defines an upper bound on $H(b)$ covariantly.

We then consider the intersection $L(b) \equiv L^+(b) \cap L^-(b)$ as a candidate for an upper bound on $H(b)$. There are two apparent problems with this choice. The first is that $L(b)$, too, would appear to depend on the time slicing of the boundary set b . Secondly, it is easy to see that $L(b) \subset C(b)$. This would conflict with our expectation that $C(b) \subset H(b) \subset L(b)$, unless it can be shown that $L(b) = C(b)$. That is, the consistency of our arguments requires that any point that lies on a causal curve that begins and ends on the boundary region b must also lie on both a past and a future-directed light-sheet emanating from b , independently of the time slicing of b .

We prove this nontrivial result in Sec. 3.4. As a corollary, the slicing independence of $C(b)$ implies the slicing independence of $L(b)$. Assuming this is the correct choice of $L(b)$, it follows that the holographic domain $H(b)$ is completely determined:

$$H(b) = C(b) = L(b) . \tag{3.3}$$

The description of $H(b)$ in terms of light-sheets allows us to define a holographic bulk RG flow, and the equivalence with $C(b)$ makes the RG flow manifestly covariant. Combining the focussing theorem with the covariant entropy bound guarantees that the number of degrees of freedom is strictly nonincreasing along the flow. We discuss this construction in Sec. 3.5.

In Sec. 3.6 we consider natural definitions of the bulk dual region from the perspective of the field theory, which contains nonlocal operators that can probe deeply into the bulk. This approach is less general than our geometric approach, because the construction

³Light-sheets were used in Ref. [42] to determine the holographic domain of (effectively) half of the global boundary. This region is not globally hyperbolic and the theorem of Sec. 3.4 does not apply; but the division of the boundary selects for a preferred slicing. Recently, Hubeny has examined which bulk regions are explored by extremal surfaces of various dimensions in static situations; she finds, as we do for the holographic domain H , that black hole interiors are never included [101].

of such operators is not known for arbitrary b . In cases where it is known we find that a region \bar{H} , very similar to $H(b)$, is picked out as the region with an unambiguous bulk interpretation.

Sec. 3.7 is largely independent of the rest of this paper. We consider the extent to which AdS/CFT adds to our understanding of quantum gravity in regions with dominant self-gravity, such as black hole interiors and cosmological regions. We describe a thought-experiment involving the formation and evaporation of (smaller) black holes in such regions, and we argue that its description requires a nonperturbative bulk theory.

3.2 The Causally Connected Region C

In this section, we argue that the CFT on the boundary portion b must at least describe the set $C(b)$ of bulk points that are both in the past and in the future of b : $H(b) \supset C(b)$.

Consider a bulk excitation very close to the boundary, e.g. at $\frac{\pi}{2} - \rho = \epsilon \ll 1$ in global coordinates,

$$ds^2 = \frac{R^2}{\cos^2 \rho} (-d\tau^2 + d\rho^2 + \sin^2 \rho d\Omega_{d-1}^2) . \quad (3.4)$$

Such an excitation is represented on the boundary by excitations with support on a region of size ϵ , at the same transverse position [165]. Here we take the boundary theory to live on a unit sphere, but this property is essentially local. It holds in Poincaré as well as global coordinates. One expects, therefore, that it should hold for any boundary region b , as long as ϵ is much smaller than the characteristic temporal and spatial size of b . This implies that the boundary theory on b must describe at least the state and the dynamics of a bulk region sufficiently close to b . We will exploit the fact that this near-boundary region, in turn, is causally connected to the larger bulk region $C(b)$ to show that the CFT must at least describe $C(b)$.

Consider now a family of bulk observers localized at $\frac{\pi}{2} - \rho = \epsilon \ll 1$. They will require a finite proper acceleration of order R^{-1} , where R is the AdS curvature radius. The proper time for which this acceleration must be maintained is

$$\Delta t_{\text{proper}} = \frac{R}{\epsilon} \Delta \tau . \quad (3.5)$$

This diverges in the limit as $\epsilon \rightarrow 0$, but we will only need to consider the case of small but finite ϵ . Thus there appears to be no obstruction, in principle, to the existence of such bulk observers. What region can they explore? We will treat the observers as a collective, imagining that they densely fill Cauchy surfaces of the boundary b , moved into the bulk by ϵ . Such observers can receive signals from the causal past of b in the bulk, and they can send signals to the causal future of b . However, they cannot determine the state in the entirety of either of these regions without making additional assumptions.

The only region that can be actively explored and manipulated by observers near the boundary is the intersection of the past and the future of b ,

$$C(b) \equiv J^+(b) \cap J^-(b) . \quad (3.6)$$




Figure 3.3: The four null hypersurfaces orthogonal to a spherical surface B in Minkowski space. The two cones F_1, F_3 have negative expansion and hence correspond to light-sheets. The covariant entropy bound states that the entropy of the matter on each light-sheet will not exceed the area of B . The other two families of light rays, F_2 and F_4 , generate the skirts drawn in thin outline. Their cross-sectional area is increasing, so they are not light-sheets, and the entropy of matter on them is unrelated to the area of B .

This region can be probed by preparing local probes at an early time, allowing them to travel deeper into the bulk and then back to the observer at a late time. The outcome of such an experiment is completely determined by the state in $C(b)$, by causality. And conversely, local fields at any point in $C(b)$ can be manipulated by such an experiment. Since the experiments of these near-boundary observers are described by the CFT, then for a consistent duality to hold the CFT must describe at least the region $C(b)$:

$$C(b) \subset H(b) . \quad (3.7)$$

If the near boundary region can both probe and be affected by $C(b)$, then the same must be true for the boundary theory itself.

By contrast, knowledge of the state in $J^-(b) - J^+(b)$, say, is not necessary in order to compute the dynamics near the boundary or in any region that can be explored from the boundary. It is sufficient to specify initial conditions on the past boundary of $C(b)$ in the bulk. Conversely, since its past boundary need not be a Cauchy surface for $J^-(b) - J^+(b)$, the state in $J^-(b) - J^+(b)$ is not determined by the state in the region $C(b)$ which can be explored from the boundary.

3.3 The Light-sheet Region L

In this section we construct a region $L(b)$ which is guaranteed to be encodable in the CFT on the boundary portion b , in the sense that holographic entropy bounds guarantee that the maximum entropy of matter in $L(b)$ does not exceed the number of degrees of freedom of the CFT on b . For an exact duality to hold, $H(b)$ must be contained in some region with this property, and we will conjecture that $H(b) \subset L(b)$.

The holographic principle [172, 166, 21, 22] is a universal relation between area and quantum information. This relation manifests itself empirically as the Covariant Entropy




Figure 3.4: Penrose diagram of a collapsing star (shaded). At late times, the area of the star’s surface becomes very small (B). The enclosed entropy in the *spatial* region V stays finite, so that the spacelike entropy bound is violated. The covariant entropy bound avoids this difficulty because only future directed light-sheets are allowed by the nonexpansion condition. L is truncated by the future singularity; it does not contain the entire star.

Bound [21]: Let A be the area of an arbitrary (open or closed) codimension-two surface B , and let S be the entropy of the matter on any light-sheet ℓ of B :

$$S[\ell(B)] \leq \frac{A}{4} . \quad (3.8)$$

At the core of the covariant entropy bound is the notion of a light-sheet (Fig. 3.3). Light-sheets are null hypersurfaces generated by nonexpanding light rays orthogonal to the surface B . There are four orthogonal directions at every point, since the surface has two sides and we can consider past and future directed light rays. If the null energy condition is satisfied, at least two of these directions will have nonpositive expansion and thus give rise to light-sheets. The covariant entropy bound holds separately on each light-sheet. (For a review, see Ref. [23].)

3.3.1 Spacelike Holography vs Light-Sheets

The holographic principle is *not* the statement that the entropy in any spatial region V is limited by the area of the surface bounding that region. This “spacelike” entropy bound follows from the covariant bound in certain special cases [21], but in general it is false. Counterexamples are easily found in cosmology, inside black holes, and even in weakly gravitating systems [23]. As a general statement valid in all spacetimes, one must not think of holography as effectively projecting out a spacelike direction. Holography projects along a null direction, just as it does in a conventional hologram.

To illustrate the falsehood of any spacelike notion of holography, consider the surface area of a collapsing star (Fig. 3.4). The area approaches zero near the singularity, but the entropy of the star starts out finite and cannot decrease. Thus, $S(V) \gg A$ at late times: a violation of the spacelike entropy bound. However, a light-sheet off of a late-time surface will not penetrate the whole star, so the covariant bound, Eq. (3.8), is upheld.

Another example is shown in Fig. 3.5: It is possible to surround any matter system with a surface of arbitrarily small area. (For weakly gravitating systems the surface will




Figure 3.5: (a) A square system in 2+1 dimensions, surrounded by a surface B of almost vanishing length A . The entropy in the enclosed *spatial* volume can exceed A . (b) [Here the time dimension is projected out.] The light-sheet of B intersects only with a negligible (shaded) fraction of the system, so the covariant entropy bound is satisfied.

consist of elements that are highly boosted with respect to the rest frame.) Thus, the spacelike entropy bound is violated. By contrast, the light-sheets off of such surfaces do not contain most of the system, evading violation of the covariant entropy bound. Details and further examples are given in Ref. [23].

Notice that the above counterexamples to a spacelike entropy bound can easily be embedded into an asymptotically AdS spacetime. In particular, consider a timelike hypersurface that is a direct product, $b = \mathbf{R} \times \mathbf{S}^2$. (For definiteness, we consider AdS_4 but our arguments hold in any number of dimensions.) In the limit where the spatial two-spheres are large, b encloses a very large spacetime region. We can think of b as a regulated version of the conformal boundary. The bulk spacetime it encloses is described by a conformal field theory living on b , with a UV cutoff on a length scale comparable to the AdS curvature radius [165].

Yet, the hypersurface b can also be foliated into two-dimensional slices which have arbitrarily small area; and each such slice bounds a global bulk slice. With this slicing, a naive “spacelike” interpretation of holography would seem to imply that the bulk can be described by a theory with an arbitrarily small Hilbert space. This interpretation is clearly incorrect, as the bulk can have arbitrarily large entropy in the limit where b approaches the boundary.

Another example is furnished by the eternal Schwarzschild-AdS black hole solution. Let us again consider the regulated global boundary given by the direct product of time with a sphere of fixed radius larger than the black hole, $b = \mathbf{R} \times \mathbf{S}^2$. As shown in Fig. 3.6, this sphere encloses an infinitely large spatial region that extends all the way to the second, disconnected conformal boundary on the far side of the black hole. This region contains arbitrarily large entropy: for example, a dilute gas of $n \rightarrow \infty$ photons can be added near the second boundary with negligible backreaction. Of course, our intuition tells us that the CFT that lives on the second boundary should be “responsible” for those photons. The only problem is that spacelike holography does *not* tell us this. In fact, it cannot possibly tell us any such thing. Black hole event horizons are global objects; for example, if we decide much later to add mass to the black hole, this will affect the location of the horizon




Figure 3.6: An AdS-Schwarzschild black hole. A sphere on the regulated boundary encloses an infinitely large spatial region that extends all the way to the second, disconnected conformal boundary on the far side of the black hole.

at all earlier times. Thus, there is no local criterion that can prevent us from extending a spacelike holographic domain to the far side of the black hole.

As the above examples illustrate, the light-like nature of the holographic relation between entropy and area is crucial for making sense of AdS/CFT as a manifestation of the holographic principle. This relation is captured by the geometric construction of light-sheets.

3.3.2 Bounds on the Holographic Domain

As shown by Susskind and Witten [165], the CFT with a UV cutoff δ has a maximum entropy equal to the proper area of the spatial boundary of AdS in standard global coordinates. The latter is considered to be located in the bulk, a coordinate distance δ away from the true conformal boundary. This observation tells us that the CFT manifestly has the correct number of degrees of freedom demanded by the holographic principle, given by the area of the boundary, not by the size of the enclosed volume.

In this argument one assumes, of course, that the CFT describes no more and no less than the spacetime region “enclosed” in the regulated boundary. This is plausible in the case of global AdS with the standard slicing of the boundary into round spheres. The analysis of the previous subsection has shown, however, that this assumption is coordinate-dependent at best, and that it is ill-defined for cases such as the Poincaré patch, where the boundary slices do not “enclose” any particular region.

In order to claim more generally that the CFT uses an area’s worth of degrees of

freedom to describe the bulk, we must characterize some relevant bulk region, given the boundary. Note that the UV/IR connection is not the issue; by construction, the Susskind-Witten argument suffices to ensure that the maximum entropy on the CFT agrees with the area of the regulated boundary. This remains true for arbitrary foliations of the boundary, as long as the short-distance regulator in the CFT is imposed with respect to the chosen time-slices. We can always remove the UV cutoff at the end and think of the CFT as living on the true conformal boundary.

The nontrivial question is how far from the boundary (how deep into the bulk) the CFT description is valid. If this region is taken to extend too far from the boundary, then the bulk entropy might be larger than the boundary area, and thus larger than the maximum entropy of the CFT. In that case, the CFT Hilbert space would be too small to capture the bulk physics. The only way to ensure that the bulk entropy is sufficiently small is to appeal to the covariant entropy bound. This is why the relevant bulk region must be related to the boundary by light-sheets. We will now explore possible concrete proposals for this relation.

Consider a timelike hypersurface b embedded in an asymptotically AdS and a foliation into spacelike hypersurfaces. We may view b as a spacetime in its own right, with one less spatial dimension than the AdS spacetime it is embedded in. In order for the theory living on b to be well-defined, we require that b be globally hyperbolic and that each time slice be a Cauchy surface. Most relevant for the present discussion is the limiting case where b lies on the conformal boundary of global AdS: $b \subset \mathcal{I}$. (b can be a proper subset of \mathcal{I} , for example if b is the boundary of the Poincaré patch of AdS.) In this case, the metric of b is defined only up to conformal transformations. We will be interested only in situations where each time slice of b is “normal”, i.e., each slice t admits both a past-directed light-sheet ℓ_t^- and a future-directed light-sheet ℓ_t^+ . This is automatically the case for $b \subset \mathcal{I}$, as we show in appendix 3.8.

Let us consider the regions $L^+(b)$ and $L^-(b)$. $L^+(b)$ is defined as the union over t of the future-directed light-sheets ℓ_t^+ of each slice. Similarly, $L^-(b)$ is the union of all past-directed light-sheets. Two natural looking possible bounds on the holographic domain are $H(b) \subset L^+(b)$ and $H(b) \subset L^-(b)$. At the fundamental level, there is no distinction between the past and the future, so $L^+(b)$ and $L^-(b)$ should play a symmetric role as bounds on the holographic domain. Let us therefore consider the candidate bounds $H(b) \subset L^+(b) \cup L^-(b)$, and $H(b) \subset L^+(b) \cap L^-(b)$. The former bound guarantees that every point in $L^+(b) \cup L^-(b)$ lies either on a future-directed or a past-directed light-sheet from some time slice of b . The latter bound is stronger: if $H(b) \subset L^+(b) \cap L^-(b)$, then every point in $H(b)$ lies on at least one past- and one future-directed light-sheet.

The choice of time slicing on b is a coordinate choice and so cannot have fundamental significance. Remarkably, the set $L(b) \equiv L^+(b) \cap L^-(b)$ is indeed independent of how b is foliated, even though $L^+(b)$ and $L^-(b)$ individually do depend on the time slicing of b . In the following section we will prove a stronger theorem: $L(b) = C(b)$. (Recall that $C(b) \equiv J^+(b) \cap J^-(b)$ is the region causally connected to b , which we discussed in the previous section.) The fact that $L(b)$ is slicing-independent follows as a corollary, since $C(b)$ is slicing-independent by construction. The simplicity and slicing-independence of $L(b)$ make it especially attractive, and we will see in Sec. 3.5 that it also leads to a useful formulation




Figure 3.7: The union of all future directed light-sheets, $L^+(b)$, coming off the usual slicing of the boundary Minkowski space (left) covers precisely the Poincaré patch (the wedge-shaped region that lies both in the future of the boundary point A and in the past of D). On the right, we show a different time slicing of the same boundary region. One of these slices is shown in blue in the bulk Penrose diagram (center); it curves up at B and down at C . The future-directed light-sheet coming off the portion of the slice near C is nearly the same as the future lightcone of C (shown in red/short-dashed), which reaches far beyond the Poincaré patch to the far side of AdS. The bulk region covered by $L^+(b)$ will thus be nearly two Poincaré patches, consisting of the points that lie in the future of A but not in the future of D (long-dashed).

of holographic RG flow.

The other set that we were led to consider, the union $L^+(b) \cup L^-(b)$, is slicing-dependent (see Fig. 3.7). This disqualifies it from further consideration, if we insist, as we do in this paper, that a unique upper bound on the holographic domain should be constructed from light-sheets off a single, arbitrary slicing of the boundary region b .

It is important to note that there are other, more complicated ways of constructing a bulk region covariantly from light-sheets. We will not analyze such sets here, but we mention some of them for completeness and future consideration. For example, to treat the possible time-slicings in a democratic way, one could form the union, or the intersection, over all possible time-slicings \mathcal{T} , of the sets $L^+(b) \cup L^-(b)$: $L'(b) = \cap_{\mathcal{T}}(L^+(b) \cup L^-(b))$ and $L''(b) = \cup_{\mathcal{T}}(L^+(b) \cup L^-(b))$. Clearly, $L(b) \subset L'(b) \subset L''(b)$. However, the restriction $H(b) \subset L''(b)$ seems too weak to ensure that the boundary theory has enough degrees of freedom to describe the bulk. For a given time slicing, it may not be the case that a given point in $L''(b)$ lies on any light-sheet emanating from a slice. Worse, there may not exist any choice of time slicing for which all of $L''(b)$ is covered by the light-sheets from the slices. The stronger bound $H(b) \subset L'(b)$ (which however is weaker than the bound we examine here) does ensure these properties, and we intend to investigate it further in future work.

In summary, the requirement that the bulk have no more degrees of freedom than

the boundary, combined with arguments of symmetry and simplicity, has led us to propose the upper bound

$$L(b) \supset H(b) \tag{3.9}$$

on the holographic domain, where $L(b) \equiv L^+(b) \cup L^-(b)$ and $L^+(b)$ ($L^-(b)$) is the union of all future-directed (past-directed) light-sheets of the time-slices that foliate the boundary region b .

3.4 Proof that $L = C$

In the previous two sections we have argued that the holographic bulk domain H dual to the boundary region b must satisfy

$$C(b) \subset H(b) \subset L(b) . \tag{3.10}$$

In this section, we will prove that

$$C(b) = L(b) . \tag{3.11}$$

This implies that

$$C(b) = H(b) = L(b) , \tag{3.12}$$

so the holographic dual H is completely determined by our assumptions.⁴

It is obvious that $L(b) \subset C(b)$. But the converse inclusion is nontrivial, particularly since $L(b)$ is constructed from two sets $L^\pm(b)$ that depend on the slicing of b , whereas $C(b)$ is slicing-independent. By Eq. (3.12), $C(b) \subset L(b)$ is required for the consistency of the arguments we have put forward in the previous section. Thus, our proof also serves as a nontrivial consistency check. It is the main technical result of this paper.

We begin by stating our assumptions and definitions. Let B be a manifold with boundary, and let $b \subset \partial B$ be a timelike embedded submanifold in B . We will require that b is globally hyperbolic when considered as a spacetime on its own. We also assume that B has the property that $J^\pm(P)$ is closed for every $P \in B$. Note that we are not assuming inextendibility of either b or B , nor global hyperbolicity of B (from which the latter assumption would follow).

In the application of our theorem to AdS/CFT, we take b to be a portion of the conformal boundary of AdS. In this case, the spacetime B should be taken as the union of the unphysical conformally rescaled AdS spacetime and its boundary. Since our theorem only relies on properties of the spacetime which are preserved by conformal transformations, we are free to construct the proof in this unphysical spacetime.

An additional assumption is that the causal relation between any two points in b computed according to the causal structure of the lower-dimensional spacetime is the same as that according to the causal structure of B itself. This is essential for a physical duality to hold: the only way to guarantee that causality is preserved on both sides of the duality is to make the causal structures compatible in this way. (In asymptotically AdS spacetimes,

⁴For certain choices of b , H thus excludes bulk points that can be represented on b [148, 96]; see, however, Sec. 3.6. Our result is consistent with the fact that only C is needed to compute CFT correlation functions in b [134].




Figure 3.8: Consider an arbitrary boundary region b (enclosed by the solid line), and a point P in the bulk region $C(b)$ (orthogonal to the page). The causal future of P , $J^+(P)$, intersected with b is shown hatched. Roughly, the strategy of the proof is to demonstrate that there exists a time slice $s_{t_*(P)}$ on the boundary that is tangent to the lower boundary of future of P in b . We show that $s_{t_*(P)}$ is the earliest time slice that has any intersection with the future of P , and that P lies on the past light-sheet of $s_{t_*(P)}$.

this assumption follows from a theorem of Gao and Wald [69]. However, this theorem relies on additional assumptions that we have no reason to make here.)

Let $\tau : b \rightarrow \mathbf{R}$ be an arbitrary time function⁵ on b such that the equal time slices are Cauchy surfaces of b . The existence of τ is guaranteed by theorem 8.3.14 of Ref. [180]. Let $K \subset \mathbf{R}$ be the image of τ . Note that K is an open interval. Let $s_t \subset b$ be the Cauchy surface consisting of points with time t , $s_t = \{p \in b | \tau(p) = t\}$. Note that each s_t is spacelike. We will not demand that τ extends to a time function on B .

For the purpose of this theorem, we shall define $\ell_t^+ \subset B$ ($\ell_t^- \subset B$) as the set of future-directed (past-directed) null geodesics which are orthogonal to s_t with no conjugate points between s_t and the endpoint. (ℓ_t^\pm is a light-sheet associated to s_t if it is initially nonexpanding away from s_t . For the case of interest, where b is a portion of the conformal boundary of AdS, this always holds in the *physical* spacetime; see Appendix 3.8. In the proof we will not make use of the nonexpansion property and so will not demand that ℓ_t^\pm be a true light-sheet in the unphysical spacetime.) Let $L^\pm = \bigcup_t \ell_t^\pm$, and $L = L^+ \cap L^-$.

Theorem $L = C$ for any choice of time function on b . where $C \subset B$ is the set of points P that lie on a causal curve that begins and ends on b .

Corollary L does not depend on the choice of time function, even though L^+ and L^- do.

Proof Trivially, $L \subset C$. It remains to be shown that $C \subset L$, for all time functions τ (i.e., foliations of b). We will show that $C \subset L^-$, and similar arguments show that $C \subset L^+$. Let

⁵That is, τ is differentiable on b , and $\nabla^a \tau$ is a past-directed timelike vector field.

$P \in C$ (and P is not contained in b).

Definition. Let $K^+(P)$ be the subset of the real numbers consisting of all t such that $s_t \cap I^+(P)$ is not empty. That is, $\{s_t | t \in K^+(P)\}$ is the set of all time slices on b which contain endpoints of future-directed timelike curves from P .

Lemma A. $K^+(P)$ is nonempty. *Proof.* Since $P \in C$, there is a point $p \in b \cap J^+(P)$. All points in $b \cap I^+(p)$ are necessarily in $I^+(P)$ (corollary following theorem 8.1.2 of Ref. [180]). Since K is an open interval, there exists $t \in K$ with $t > \tau(p)$. Any inextendible timelike curve passing through p intersects s_t because s_t is a Cauchy surface, which shows that $s_t \cap I^+(p)$, and consequently $b \cap I^+(P)$, is nonempty. Therefore $K^+(P)$ is nonempty.

Lemma B. For any $t \in K^+(P)$, the set $A_t = b \cap J^+(P) \cap J^-(s_t)$ is compact. *Proof.* Since $P \in C$, $b \cap J^-(P)$ is not empty. Let $q \in b \cap J^-(P)$. Let $j^+(q)$ be the causal future of q within the spacetime b , and define $j^-(s_t)$ similarly. Then $j^+(q) \cap j^-(s_t)$ is compact (theorem 8.3.12 of Ref. [180]). But because the causal relation between points in b is the same whether we treat them as events in the spacetime b or the spacetime B , we also have that $j^+(q) \cap j^-(s_t) = b \cap J^+(q) \cap J^-(s_t)$. Since this set is compact, it is closed as a subset of B . $J^+(P)$ is a closed subset of $J^+(q)$, and it follows that $b \cap J^+(P) \cap J^-(s_t) = J^+(P) \cap (b \cap J^+(q) \cap J^-(s_t))$ is a closed subset of $b \cap J^+(q) \cap J^-(s_t)$. Therefore $b \cap J^+(P) \cap J^-(s_t)$ is compact.

Definition. Let $t_*(P) = \inf_{t \in K^+(P)} t$.

Lemma C. For any $t \in K^+(P)$, $t_*(P) = \min_{p \in A_t} \tau(p)$. In particular, this means that there is a surface $s_{t_*(P)}$ and that $s_{t_*(P)} \cap J^+(P)$ is not empty. *Proof.* By Lemma B, A_t is compact and hence $\tau(p)$ attains a minimum value τ_{\min} . Consider a point $p \in s_{\tau_{\min}} \cap J^+(P)$ and an inextendible future-directed timelike curve γ in b starting at p . All points on γ other than p are necessarily in $b \cap I^+(P)$ (corollary following theorem 8.1.2 of Ref. [180]), and γ necessarily intersects $s_{t'}$ for every $t' \in K$ with $t' > \tau_{\min}$ because $s_{t'}$ is a Cauchy surface. This implies that all such times t' are in $K^+(P)$, and hence $\tau_{\min} \geq t_*(P)$. If $\tau_{\min} > t_*(P)$, then there is some $t' \in K^+(P)$ with $\tau_{\min} > t'$. Then $t > t'$, and so $s_{t'} \cap A_t$ is not empty by the definition of A_t . But then $\tau_{\min} \leq t'$ by the definition of τ_{\min} , which contradicts $\tau_{\min} > t'$. Therefore we conclude that $\tau_{\min} = t_*(P)$.

Lemma D. $K^+(P)$ is an open subset of the real numbers. This implies that $s_{t_*(P)} \cap I^+(P)$ is empty. *Proof.* Let γ_p be an inextendible timelike curve in b which passes through a point $p \in b \cap I^+(P)$. We can choose to parametrize γ_p by the time function τ , which means that $\gamma_p : K \rightarrow b$ is a continuous function satisfying $\tau(\gamma_p(t)) = t$. Since $I^+(P)$ is open, it follows that the inverse image $\gamma_p^{-1}[I^+(P)]$ is open in K , and therefore open in \mathbf{R} because K is an open interval. By applying $\tau \circ \gamma_p$, we see that $\gamma_p^{-1}[I^+(P)] \subset K^+(P)$. Then $\bigcup_{p \in b \cap I^+(P)} \gamma_p^{-1}[I^+(P)]$ is an open subset of \mathbf{R} equal to $K^+(P)$.

Proof of theorem. Lemma C and Lemma D together demonstrate that there is a surface $s_{t_*(P)}$ such that $s_{t_*(P)} \cap I^+(P)$ is empty while $s_{t_*(P)} \cap J^+(P)$ is nonempty. Let $p \in s_{t_*(P)} \cap J^+(P)$. By the corollary following theorem 8.1.2 of Ref. [180], there is a null geodesic connecting p to P . Furthermore, since $s_{t_*(P)} \cap I^+(P)$ is empty, this null geodesic cannot be deformed to a timelike curve connecting $s_{t_*(P)}$ to P . Then by theorem 9.3.10 of Ref. [180], this null geodesic must be orthogonal to $s_{t_*(P)}$ and have no conjugate points between P and p . This shows that $P \in \ell_{t_*(P)}^- \subset L^-$.

3.5 Covariant Renormalization Group Flow

We repeatedly made use of the UV/IR connection [165] in motivating our constructions. Entropy bounds play a very important role in UV/IR: the bulk region within the IR cutoff must have entropy limited by the area in Planck units of a time slice on the cutoff surface. However, as we have stressed, this only holds true for very special time slices and the covariant entropy bound must be used in general to bound the bulk entropy. Holographic renormalization group flows [5, 52, 61, 98] aim to refine UV/IR, but in all standard approaches manifest covariance is lost and the status of entropy bounds is unclear. Here we outline an approach which reproduces the standard results and remedies both of these problems. Our construction gives an improved bulk radial flow, however it does not address the open question of finding a precise field theory RG representation of the bulk flow.




Figure 3.9: The covariant bulk RG flow presented here reproduces the standard bulk RG flow in certain coordinate systems. Here we illustrate the construction in global coordinates of Anti-de Sitter space. For a given coordinate time cutoff δ , the union over t of the intersection surfaces $s_t^\delta = \ell_{t+\delta/2}^- \cap \ell_{t-\delta/2}^+$ form a timelike hypersurface b^δ in the bulk (left). The cross-sectional area of a given light-sheet will be greater on the surface b^δ than on $b^{\delta'}$ (right). The difference $(A - A')/4$ bounds the entropy on the red light-sheets going from b^δ to $b^{\delta'}$, meaning that the bound applies to the entire darkly shaded wedge between them. The lightly shaded region between the hypersurfaces $b^{\delta'}$ and b^δ is covered by such wedges.

Choose a time function τ on b . Then for any t_1, t_2 the set $b[t_1, t_2] = \{p \in b | t_1 < \tau(p) < t_2\}$ satisfies the conditions of our theorem and we can associate to it the region $H(b[t_1, t_2])$. Now introduce a cutoff timescale δ to the theory. In the bulk, we should remove the union (over t) of the sets $H(b[t - \delta/2, t + \delta/2])$ from $H(b)$. The remaining bulk region, H^δ , is the region described by the cutoff CFT. This prescription is similar to the construction of [26], where it was shown that the IR cutoff surface as normally defined can

be reproduced using only causality.

Denote the boundary of H^δ by b^δ . The time function on b automatically induces a time function on b^δ .⁶ By construction, b^δ is the union of sets of the form $s_t^\delta = \ell_{t+\delta/2}^- \cap \ell_{t-\delta/2}^+$, which we can take as time slices on b^δ . The maximum proper time in b^δ between s_t^δ and $s_{t+\delta}^\delta$ will be of order the AdS time. More generally, we can let δ depend on t .

To change the cutoff from δ to δ' , we have two equivalent options: First, we can return to the true boundary b and repeat the construction with δ' in place of δ . Second, we can use the surface b^δ , together with its induced time function, as the starting point for the construction, with cutoff $\delta' - \delta$. The geometry is the same either way, because the light-sheets from b^δ are continuations of the light-sheets from b .

By the definition of the L , and since $H = L$, all time slices of b^δ are *normal*. That is, their orthogonal light-rays are everywhere nonexpanding in the direction of the bulk RG flow. A generic time slicing on b^δ will *not* admit ingoing past- and future-directed light-sheets at all points on all time-slices; only the slicing induced by the flow has this property. And generic hypersurfaces other than those induced by the flow may not admit any slicing with this property. Note that this property is inherited from the remarkable property of the conformal boundary described in the Appendix: any slicing of the conformal boundary is everywhere normal in the physical metric.

The fact that all time slices of b^δ are normal is highly nontrivial. It leads to two additional, attractive features that distinguish this geometric flow from, say, the flow along spacelike geodesics:

- Entropy bounds hold for both the UV and the IR regions.
- These bounds guarantee that the effective number of degrees of freedom is nonincreasing along the flow.

Consider first the IR region. The covariant entropy bound guarantees that any future- or past-directed light-sheet from any slice s_t^δ has entropy less than the area of that slice, in Planck units. Since the area of the light-sheet is nonincreasing as it moves away from the boundary, the number of degrees of freedom is nonincreasing along the RG flow.

Now consider the UV region, i.e., follow the light-sheet from $s_t^{\delta'}$ backward to some cutoff $\delta < \delta'$ closer to the true boundary. The cross-sectional area will be larger on b^δ and we can consider the area difference $A - A'$. A generalized version of the covariant entropy bound [65] implies that the entropy on the *partial light-sheet* between A and A' is bounded by $(A - A')/4$. Note that this bound applies to both light-sheets bounding the wedge-shaped region between the slices $s_{t+(\delta'-\delta)/2}^\delta$ and $s_{t-(\delta'-\delta)/2}^\delta$ on b^δ and $s_t^{\delta'}$ on $b^{\delta'}$, and hence to the entire information content of the regions that are integrated over in the RG flow.⁷

RG flows are normally defined in terms of proper distances, times, or energies, whereas the above construction is in terms of an arbitrary coordinate time. To define

⁶This will be true for sufficiently well-behaved time functions t ; the precise conditions will be investigated elsewhere.

⁷In general, the area $A(t)$ of different slices on b^δ will not be independent of t . It seems likely that requiring this independence will lead to preferred choices of $\delta(t)$ and thus for the flow. Also, while the areas will be automatically finite at finite δ in the case of the global boundary, they will be infinite in general and must be regulated.

proper distances on the boundary, one has to choose a conformal frame, and the choice of time function in our construction is analogous. Heuristically, one can imagine choosing the conformal factor so that τ becomes proper time, which means that it is conjugate to energy.

We have emphasized the geometric aspects of the RG flow, i.e., the bulk side. The question of what precisely the removal of near-boundary regions corresponds to in the CFT remains subtle, and it will be interesting to revisit it in light of the geometric flow we have described. In particular, it would be nice to understand whether a duality holds for the wedge-shaped regions associated with thin boundary strips, and for the Rindler portion of the bulk dual to small causal diamonds on the boundary.

As the example of the global boundary (Fig. 3.9) illustrates, the description of the bulk region $H(b[t - \delta/2, t + \delta/2])$, for small time intervals δ , cannot involve the full set of CFT degrees of freedom, since approximately $\exp(A/4)$ CFT states correspond to bulk states localized to the interior of the bulk hypersurface b^δ . If a duality exists, the CFT degrees of freedom relevant for $H(b^\delta)$ will not involve certain nonlocal operators that occupy scales larger than δ .

3.6 Nonlocal Operators and the Bulk Dual

We constructed a candidate for a bulk region $H(b)$ dual to b geometrically, using considerations of causality and holographic entropy bounds. Except for appealing to the UV/IR relation [165], we did not make use of detailed properties of the AdS/CFT duality or the boundary theory. In this section we will explore a different approach to this question, namely the use of nonlocal operators in the CFT to probe deeply into the bulk region. We will examine the relationships between this approach and the geometric one. Note that the operator approach is available only for choices of b where the construction of boundary duals to local bulk operators is known, so it is less general than our geometric construction.

The methods of the operator approach were recently discussed by Heemskerk, Marolf, and Polchinski [96], and we refer the reader to that work for further details. The problem is to identify a subset of operators $\{\phi(x)\}$ within the set of all operators in the CFT on b , indexed by a position x in a semiclassical spacetime of higher dimension, which can be identified as local bulk operators. An important point is that the definition of the local bulk operators depends on the background metric. Our analysis below pertains to the case where this metric is held fixed up to perturbative corrections in $1/N$, so this issue should not pose difficulties.

The CFT definition of $\phi(x)$ will involve nonlocal CFT operators known as precursors [148, 171, 66]. An (in principle) explicit construction is available in the case of global AdS, and the resulting operators make use of an entire Cauchy surface of the global boundary. For this reason, and also for simplicity, we will spend most of this section focused on the case of a short strip $b = \mathcal{S}$,

$$\mathcal{S} = (-\tau_0, \tau_0) \times S^d, \text{ with } \tau_0 \ll 1, \quad (3.13)$$

which we normally think of as being embedded in the global boundary (see Fig. 3.10), using the coordinates of Eq. 3.4. From the global point of view, the set of operators on the Cauchy

surface $\tau = 0$ of the boundary is complete, and so in particular contains the operator $\phi(x)$ for every point x in the entire global bulk.




Figure 3.10: A cross-section of Anti-de Sitter space, showing a short strip region \mathcal{S} centered around $\tau = 0$ on the boundary, and the bulk region $S(\mathcal{S})$ spacelike separated from \mathcal{S} . A local operator at the origin of the bulk can be written in terms of local operators on the boundary smeared over the boundary region spacelike-related to the origin, within the green wedges. This region is much larger than \mathcal{S} (red thick line), stretching from $\tau = -\pi/2$ to $\tau = +\pi/2$.

Now consider modifications of the CFT Hamiltonian \mathcal{H} outside b , and let us define \bar{H} as the largest bulk region such that operators $\phi(x)$ for $x \in \bar{H}$ can be represented in terms of CFT operators in b in an unambiguous way. There are some modifications of the CFT Hamiltonian which continue to yield a well-defined bulk Hamiltonian. For instance, we can insert a local source in the CFT whose effect in the bulk is to cause a particle to propagate causally inward from the boundary. As we discuss below, certain modifications of the CFT Hamiltonian which make use of nonlocal operators can lead to ambiguities in the bulk Hamiltonian. In the case where we allow only those modifications leading to well-defined bulk Hamiltonians, we identify a region S on which the operators $\phi(x)$ have unambiguous CFT representations in b . In the more general case of modifications which lead to an ambiguous bulk Hamiltonian, we will identify \bar{H} as the subset of S for which $\phi(x)$ still has an unambiguous CFT representation in b . We will note that \bar{H} is closely related to H .

The region S The only known construction of $\phi(x)$ in the boundary theory consists of two steps and applies either in global AdS or the Poincaré patch. In the first step, one

writes a local bulk operator $\phi(x)$ as a smeared local operator on the boundary:

$$\phi(x) = \int dy' K(x|y') \mathcal{O}(y') + O(1/N) . \quad (3.14)$$

Here y' denotes a boundary coordinate, while x is the bulk coordinate.⁸ The smearing function $K(x|y)$ is not unique, but a convenient choice is nonzero if and only if y spacelike-related to x [90, 89, 107]. One can think of $K(x|y)$ as providing the solution to the equation of motion for $\phi(x)$ given boundary data on the asymptotic boundary; this is a spacelike analogue of the standard initial value problem. Note that for any choice of x in global AdS, the boundary support of $K(x|y)$ is larger than the region \mathcal{S} . The second step of the construction uses unitary evolution in the CFT to write all of the local CFT operators $\mathcal{O}(y')$ appearing in Eq. 3.14 in terms of nonlocal operators defined at $\tau = 0$ (or some other timeslice within the short strip \mathcal{S}).

Now let us identify a subset $S(\mathcal{S})$ of the bulk such that the above construction of the CFT operator $\phi(x)$ is independent of modifications of the boundary Hamiltonian outside b ,⁹ provided that these modifications continue to lead to a well-defined bulk Hamiltonian. In Ref. [96], it was shown that S consists of the bulk points that are neither in the causal future of the future boundary of \mathcal{S} , nor in the causal past of its past boundary. We will refer to such points somewhat loosely as being “spacelike-related” to \mathcal{S} (see Fig. 3.10).

At first, this result may seem surprising, so it is worth reviewing the argument for it [96]. First, fix a fiducial CFT Hamiltonian \mathcal{H} over the entire global boundary, and follow the above procedure to construct an operator \mathcal{O} satisfying $\phi(x) = \mathcal{O}$, for some x spacelike-related to \mathcal{S} (the special case $x = 0$ is pictured in Fig. 3.10). But there is another way we can produce the operator \mathcal{O} : We can use causal bulk evolution to write $\phi(x)$ in terms of $\phi(x')$ for points x' in the future of the x :

$$\phi(x) = \int_{\Sigma} dx' [\phi(x') \nabla G_{\text{adv}}(x|x') - \nabla \phi(x') G_{\text{adv}}(x|x')] + O(1/N) , \quad (3.15)$$

where Σ consists of a bulk Cauchy surface and, possibly, a portion of the global boundary. We can use Eq. 3.14 to write the $\phi(x')$ appearing here in terms of operators on the boundary. By evolving $\phi(x)$ sufficiently far forward into the future, the y' -support of $K(x'|y')$ will lie in the region $\tau > 0$ for every x' appearing in Eq. 3.15. Now we evolve this new smeared operator back to $\tau = 0$ to obtain a second operator \mathcal{O}' . However, $\mathcal{O}' = \phi(x) = \mathcal{O}$, and so these two procedures actually give the same answer. Now suppose we modify the CFT Hamiltonian from our fiducial choice \mathcal{H} to $\tilde{\mathcal{H}}$, and we stipulate that $\tilde{\mathcal{H}}$ only differs from \mathcal{H} for $\tau < 0$. We can repeat the procedure to compute new operators $\tilde{\mathcal{O}}$ and $\tilde{\mathcal{O}}'$, which are equal to each other and to $\phi(x)$. The claim is that, since the calculations of \mathcal{O}' and $\tilde{\mathcal{O}}'$ refer only to the $\tau > 0$ region of the bulk and boundary, and $\mathcal{H} = \tilde{\mathcal{H}}$ in that region, the computations are identical and so manifestly we have $\mathcal{O}' = \tilde{\mathcal{O}}'$. Therefore $\tilde{\mathcal{O}} = \mathcal{O}$, and the change in fiducial Hamiltonian did not change the operator assignment. An analogous argument can be made for modifications to the CFT Hamiltonian for $\tau > 0$.

⁸The presence of a source will modify this equation in the appropriate way given by Green’s identity.

⁹We impose the boundary condition that the boundary state at $\tau = 0$ remain unchanged.




Figure 3.11: According to the Hamiltonian on the boundary strip \mathcal{S} , no source acts at the origin in the bulk, so the expectation value of ϕ vanishes everywhere. At the time τ_0 outside the strip, a source term for the nonlocal boundary operator dual to $\phi(0,0)$ can be added to the boundary Hamiltonian. This causes the expectation value of ϕ to be nonzero in the future of $(0,0)$, in contradiction with the earlier conclusion about the same bulk points. Thus, unless we possess information about the exterior of \mathcal{S} on the boundary guaranteeing that such operators do not act, the bulk interpretation of regions outside $\bar{H}(\mathcal{S}) = C(\mathcal{S})$ is potentially ambiguous.

The region \bar{H} The above argument relied crucially on the existence of a well-defined bulk Hamiltonian. However, there are reasonable modifications of the CFT Hamiltonian \mathcal{H} for which this will not be the case (see Fig. 3.11). At the time τ_0 on the boundary, let us add to \mathcal{H} a source for the nonlocal CFT operator \mathcal{O} dual to a local bulk operator at the origin, $\phi(\tau = 0, \rho = 0)$:

$$\mathcal{H} \rightarrow \mathcal{H} + J\delta(\tau - \tau_0)\mathcal{O} \quad (3.16)$$

We are working perturbatively, so the operator \mathcal{O} is the one constructed using the above method and the unmodified Hamiltonian \mathcal{H} . In the bulk, this source acts completely locally as a source for ϕ at the origin at $\tau = 0$. Note that this bulk point is spacelike related to the boundary slice τ_0 .

Now let us compute the expectation value of the bulk operator ϕ at the origin, at some infinitesimal time ϵ after the source acts. This operator can be constructed by the usual methods, but those methods require a knowledge of the bulk Hamiltonian in the region $S(\mathcal{S})$. This Hamiltonian is ambiguous: from the viewpoint of the strip \mathcal{S} , the source does not act, since it acts in the CFT only at τ_0 . Then the bulk evolution should be computed from the usual bulk Hamiltonian without source, and we obtain $\langle \phi(\epsilon, 0) \rangle = 0$. From the

viewpoint of the CFT at the time τ_0 , the bulk region near $(0,0)$ contains a source. Then the bulk evolution should take this source into account and we obtain $\langle\phi(\epsilon,0)\rangle \neq 0$. Thus, there is no unique assignment of a bulk field value at $(\epsilon,0)$. In this sense, $(\epsilon,0)$ should not be considered a bulk point dual to the strip.

Let us now consider a general boundary region b and construct a bulk region \bar{H} such that the interpretation of what happens in \bar{H} is unambiguous (Fig. 3.12). The ambiguities we discussed arise from inserting nonlocal CFT operators into the Hamiltonian on Cauchy surfaces of the boundary which do not intersect the region b . A modification of the CFT Hamiltonian on such a Cauchy surface σ can lead to an ambiguous bulk Hamiltonian in the region $S(\sigma)$ spacelike-related to σ . Thus we find that \bar{H} is the complement of the union of all $S(\sigma)$, where σ is any Cauchy surface for the global boundary which lies in the complement of b .




Figure 3.12: The shaded region shows bulk points spacelike related to a global boundary Cauchy surface σ . The union of all such sets over the collection of boundary Cauchy surfaces which do not intersect \mathcal{S} has an ambiguous bulk interpretation when the boundary Hamiltonian is allowed to vary outside of \mathcal{S} . The unambiguous region, $\bar{H}(\mathcal{S})$, is the complement of this union. In this example, we see that $\bar{H}(\mathcal{S}) = C(\mathcal{S})$.

The region \bar{H} is closely related to $C(b)$ and hence to $H(b)$. It is easy to see that $H = C \subset \bar{H}$. The study of a number of examples suggests that $H \neq \bar{H}$ if and only if an event horizon is present in the bulk. It would be nice to study \bar{H} and its relation to H further. The discussion in the following section may be relevant.

3.7 Quantum Gravity Behind Event Horizons

In this section we discuss an issue that is somewhat orthogonal to the main subject of this paper: the degree to which the reconstruction of bulk regions, perturbatively in $1/N$, allows us to claim that AdS/CFT provides a full quantum gravity theory for regions behind event horizons, such as the interior of a black hole, or cosmological regions. We will construct an experiment behind the horizon whose outcome is known but not captured by such methods.

The CFT provides a full quantum gravity theory for observers near the boundary.¹⁰ It completely settles the issue of whether the formation and evaporation of a black hole is a unitary process. It is crucial for this argument that the time evolution is carried out on the boundary, where it is manifestly unitary; the duality is used only at early and late times in order to recover the in and out states in the bulk [167, 146, 78, 64].

To what extent can we regard the CFT *also* as a full quantum gravity theory for an observer falling into the black hole? Perhaps, by reordering the degrees of freedom, one could interpret the CFT as providing a nonperturbative definition of quantum gravity for the infalling observer? This would require that the bulk dual region is ambiguous, at the nonperturbative level. This may be the case, but the dictionary that would provide this definition is not known at the required level of precision.

The black hole interior is clearly encoded in the CFT if one makes use of bulk equations of motion to evolve the infalling data back out of the black hole and to the boundary. But in the same approximation, we can also generate a xeroxing paradox [169]: at the semiclassical level, there is no manifest obstruction to evolving to global bulk slices that contain both the black hole interior and the Hawking cloud.

Perhaps we should restrict the bulk evolution by hand to the causal patch of an infalling observer, and settle for this approximate description of the black hole interior? The finiteness of entropy bounds inside black holes imply that there cannot be exact observables associated with the infalling observer at late times. Thus, the description of the infalling observer should be less precise than that of the observer near the boundary, who has access to exact observables. Perhaps the need to use bulk evolution is simply a reflection of this intrinsic limitation?

In fact, however, it is clear that the infalling observer requires a theory that goes beyond semiclassical bulk evolution. This can be seen from the following thought-experiment. Consider an infalling observer who performs scattering experiments inside a Schwarzschild-AdS black hole of radius R , after crossing the event horizon. The scattering occurs at high enough energy to produce a (smaller) black hole of mass m , which then evaporates. We are free to choose parameters so that the scattering effectively takes place in flat space:

$$R_{\text{AdS}} \gg R \gg f(m) \gg 1, \quad (3.17)$$

where $f(m)$ is at least the evaporation timescale of the black hole (m^d in d spatial bulk dimensions). For sufficiently large but finite $R/f(m)$, the infalling observer can confirm

¹⁰Note that all observers that remain outside the black hole have the same causal diamond, consisting of the exterior of the black hole. The covariant RG flow we described in Sec. 3.5 can be thought of as moving the observers deeper into the bulk, but note that the flow never enters the black hole.

the unitarity of black hole formation and evaporation to any required precision. But this conflicts with the result that would be obtained from the semiclassical analysis (Hawking’s calculation).

We conclude that the proper description of an infalling observer requires a quantum gravity theory beyond the semiclassical approximation. The CFT on the boundary does not provide us with that theory, since its application to the infalling observer relies on semiclassical bulk evolution. No other way of relating local operators inside the black hole to boundary operators is known; therefore, we cannot replace this bulk evolution by boundary evolution, as we did in the case of scattering experiments performed by a near-boundary observer.

Of course, in the limit used in the thought-experiment, one could imagine “cutting out” the spacetime region of size $f(m)$ that contains the scattering experiment. One could embed this region in an asymptotically AdS spacetime with $R_{\text{AdS}} \gg f(m)$ and use the flat space S-matrix, which can be computed using AdS/CFT [167, 146, 78, 64]. In fact, this argument is what gives us confidence that the process is indeed unitary. But how would this prescription generalize? A theory of the infalling observer that relied on this type of cutting and pasting would not be applicable to the highly dynamical regions deeper inside the black hole, nor to the spacelike singularity, which cannot be so transplanted.

An exact version of this cut-and-paste process is available if the black hole is formed by sending in a spherical null shell from the boundary of AdS [100]. By causality, the bulk region in the past of the shell is the same no matter whether we decide to send in the shell or not. If we do send in the shell, then this region includes a portion w of the black hole interior. If we do not, then the same bulk region w is dual to operators on the boundary, which can be evolved to nonlocal operators W on a single global boundary time slice preceding the insertion of the shell. This illustrates that cut-and-paste is well-defined precisely in the limit where it yields no information about regions that could *only* exist in the interior of black holes, such as singularities and a highly dynamical geometry.

The absence (so far) of an intrinsic bulk theory at the nonperturbative level appears to impose crucial limitations on our ability to describe black hole interiors and cosmological regions¹¹ via AdS/CFT, beyond what follows from the approximate methods that were already available to generate bulk evolution.

3.8 Light-sheets from the conformal boundary of AdS

In this appendix we show that the light-rays orthogonal to any spacelike slice of the conformal boundary of AdS have nonpositive expansion into the bulk in the physical metric, and thus generate light-sheets. This is an important property of light-sheets in asymptotically AdS spacetimes. It guarantees that the maximum entropy of the bulk holographic domain H is nonincreasing under the covariant RG flow described in Sec. 3.5.

¹¹A similar conclusion [75] applies to the interior of other event horizons, such as an FRW universe [128]. In this case there are two natural choices of conformal frame on the boundary [7]: one in which the CFT is well-behaved, and another in which the coefficient of a relevant operator diverges in finite time. It is tempting to interpret this violent behavior in terms of an infalling observer hitting the big crunch singularity behind the horizon, but it can be understood more simply as the arrival of the bulk domain wall on the boundary.

In Poincaré coordinates, the AdS metric is:

$$ds^2 = \frac{1}{z^2}(dz^2 + \eta_{\mu\nu}dx^\mu dx^\nu) , \quad (3.18)$$

where $\eta_{\mu\nu}$ is the metric for d dimensional Minkowski space. Consider the conformally rescaled metric

$$d\tilde{s}^2 = dz^2 + \eta_{\mu\nu}dx^\mu dx^\nu , \quad (3.19)$$

which is $d + 1$ dimensional Minkowski space. Take some small region of the boundary and let $\tilde{\theta}$ be the expansion of some congruence of infinitesimally neighboring light rays in this space. That is

$$\tilde{\theta} = \frac{d \log \tilde{\mathcal{A}}}{d\tilde{\lambda}} , \quad (3.20)$$

where $\tilde{\mathcal{A}}$ is the infinitesimal area spanned by the light rays and $\tilde{\lambda}$ is the affine parameter. In order for the null geodesics to remain affinely parameterized after the conformal transformation, the affine parameter must transform as (Appendix D of [180])

$$\frac{d\tilde{\lambda}}{d\lambda} = cz^2 , \quad (3.21)$$

where c is a constant and z^2 is the conformal factor. The area will transform as

$$\mathcal{A} = \frac{\tilde{\mathcal{A}}}{z^{d-1}} . \quad (3.22)$$

Using (3.20), (3.21), and (3.22) we find that the expansion in the AdS spacetime (3.18) is

$$\theta = cz^2\tilde{\theta} - c(d-1)z\tilde{k}^z , \quad (3.23)$$

where

$$\tilde{k}^z = \frac{dz}{d\tilde{\lambda}} .$$

Since both $\tilde{\theta}$ and \tilde{k}^z are defined in $d + 1$ dimensional Minkowski space using a congruence of null rays orthogonal to a spacelike (and hence nowhere null) foliation, they are finite. Thus, at the boundary ($z = 0$), we have $\theta = 0$.

Chapter 4

Null Geodesics, Local CFT Operators, and AdS/CFT for Subregions

4.1 Introduction

The AdS/CFT correspondence [127, 185] provides an important tool for obtaining insight into quantum gravity. Yet even today, the seemingly basic question of how bulk locality is encoded in the boundary theory—in other words, which CFT degrees of freedom describe a given geometrical region in the bulk—has resisted a simple, precise answer.

In this chapter, we investigate the related question of AdS/CFT subregion dualities. That is, we consider the possibility that a CFT restricted to a subset of the full AdS boundary is dual to a geometric subset of the AdS bulk. There is no obvious reason that a geometric region on the boundary has to correspond to a geometric region in the bulk, but there are strong arguments for such a subregion duality in certain simple cases [48], and intriguing hints [165, 154, 104] that it may be true more generally.

The problem of precisely what bulk region should be associated with a given boundary region is complicated and has been explored recently by [40, 51, 103]. We will not propose or adopt a rule for constructing such an association. Instead, we focus on one nice feature of the global AdS/CFT duality which does not generalize to arbitrary subregions, namely the ability to reconstruct the bulk using local CFT operators in the classical limit [90, 89, 107, 96]. Specifically, we will emphasize the role of continuity of the bulk reconstruction, and propose a simple geometric diagnostic testing whether continuous reconstruction holds for a given subregion.

To motivate our investigation, first consider the full global AdS/CFT duality. We will work in Lorentzian signature and fix the Hamiltonian of the CFT, which corresponds to fixing all the non-normalizable modes in the bulk. Now take the $G_N \rightarrow 0$ limit in the bulk; the bulk theory reduces to solving classical field equations in a fixed background. The non-normalizable modes are fully determined and non-dynamical, but there are still many allowed solutions because of the normalizable modes. CFT data on the boundary should be sufficient to specify a particular bulk solution. Normalizable modes in the bulk




Figure 4.1: Here we show the AdS-Rindler wedge inside of global AdS, which can be defined as the intersection of the past of point A with the future of point B . The asymptotic boundary is the small causal diamond defined by points A and B . The past lightcone of A and the future lightcone of B intersect along the dashed line, which is a codimension-2 hyperboloid in the bulk. There is a second AdS-Rindler wedge, defined by the points antipodal to A and B , that is bounded by the same hyperboloid in the bulk. We refer to such a pair as the “right” and “left” AdS-Rindler wedges.

approach zero at the boundary, but a nonzero boundary value can be defined by stripping off a decaying factor,

$$\phi(b) \equiv \lim_{z \rightarrow 0} z^{-\Delta} \Phi(b, z) , \quad (4.1)$$

where z is the usual coordinate that approaches zero at the boundary, b stands for the boundary coordinates, and Φ is a bulk field. We will also use the notation $B = (b, z)$ where convenient. By the “extrapolate” version of the AdS/CFT dictionary [6], these boundary values are dual to expectation values of local operators,

$$\phi(b) = \langle \mathcal{O}(b) \rangle . \quad (4.2)$$

We can now ask a classical bulk question: do the boundary values ϕ determine the bulk solution everywhere? This is a nonstandard type of Cauchy problem, because we are specifying data on a surface that includes time.

In a simple toy model where the bulk contains only a single free field with arbitrary mass, Hamilton *et al.* [90, 89] showed explicitly that this boundary data does specify the bulk solution completely in global AdS. The fact that the boundary data specifies the bulk solution can be considered the classical, non-gravitational limit of AdS/CFT. It is a nontrivial fact that expectation values of local CFT operators are sufficient to reconstruct the bulk field in this case.

A proposed subregion duality must pass the same test. Is the CFT data in a boundary subregion sufficient to reconstruct the bulk solution within the corresponding

bulk subregion? In principle, the CFT data is quite complicated. The simplification that occurred in global AdS/CFT, that expectation values of local boundary operators were sufficient, may or may not carry over to other cases, and our task is to properly account for when it does. This is a problem in the theory of classical differential equations which we can hope to solve. Simple examples show that the problem is subtle, however, and to properly capture the physics of the problem we need to differentiate between bulk reconstruction and *continuous* bulk reconstruction.

The simplest illustration comes from AdS-Rindler space, which can be described as follows. In the global duality, the CFT is formulated on a sphere cross time and the associated bulk is global AdS. Let us divide the boundary sphere at some time across the equator. In the bulk, the extremal surface ending on the boundary equator is a hyperboloid, and we can use Rindler-type coordinates in AdS so that this extremal surface is a Rindler horizon. The northern hemisphere on the boundary extends naturally into a small causal diamond on the boundary, namely the region determined by time evolution of the data in the northern hemisphere. The corresponding bulk region is a Rindler wedge, shown in Fig. 4.1, which we will call AdS-Rindler space.

Does the global boundary data, restricted to the small boundary diamond, determine the bulk solution in the corresponding AdS-Rindler wedge? Hamilton *et al.* [89] also addressed this question. They determined that a particular analytic continuation of the boundary data was necessary to reconstruct the bulk. Here we provide a different answer that does not rely on analytic continuation of the boundary data. We claim that there *is* a direct map from the boundary data to the bulk field, but that the map is *not continuous*. This leaves the physical interpretation open to doubt.

There are two reasons to focus on the question of continuity. First, if the subregion duality is correct, we would expect that measuring boundary data to finite precision should determine the bulk data to a corresponding precision. This is only true if the bulk solution depends continuously on the boundary data. Second, the question of continuous reconstruction seems to be mathematically robust; we will be able to make heuristic contact with nice mathematical theorems about when continuous reconstruction is possible.

Continuous reconstruction fails because there are finite excitations in the bulk Rindler wedge with an arbitrarily small imprint on the boundary data. The physics of these excitations is simple: there exist null geodesics that pass through the bulk Rindler wedge, but avoid the boundary diamond. One can construct solutions where geometric optics is an arbitrarily good approximation and the energy is concentrated along such a null geodesic. In this way, we can construct solutions that are finite in the bulk but have arbitrarily small boundary data in the Rindler wedge.

We can also ask a slightly different mathematical question, which is closely related to bulk reconstruction from the boundary data but simpler to analyze: the question of unique continuation. Suppose we are given the bulk solution in some region near the boundary, and we want to continue the solution further into the bulk. In the AdS context, evolution inward is roughly dual to RG flow in the CFT. This question is closely related to the previous one, and again can be diagnosed with null geodesics [174]. In the case of the bulk Rindler wedge we find that unique continuation fails as well. We cannot evolve the solution radially inward in this case.

Given the connection to continuity and local reconstruction, as well as geometrical simplicity, we are motivated to propose a diagnostic for continuous bulk reconstruction from local CFT operators:

Does every null geodesic in the bulk subregion have an endpoint on the corresponding boundary subregion?

Despite the failure of this diagnostic for AdS-Rindler, there are good reasons to think that this particular subregion duality actually holds. The Rindler wedge can be thought of as an eternal black hole with a hyperbolic horizon. This suggests that a duality holds, by analogy with the ordinary eternal black hole: the CFT in the Hartle-Hawking state may be restricted to one boundary component, and the resulting thermal state is dual to one of the two exterior regions of an eternal AdS-Schwarzschild black hole [48].

Since continuous reconstruction from CFT one-point functions fails for this subregion, we learn that nonlocal boundary operators must play an important role in the duality even in the classical limit. Generalizing this result, we learn that nonlocal CFT operators [148, 171] are important when subregions are small enough that the boundary region no longer captures all null rays passing through the bulk.

The remainder of the chapter is organized as follows. In Section 4.2 we review the general procedure for reconstructing the bulk solution from boundary data which was employed by Hamilton *et al.* in their work. We also show how to determine continuity of the reconstruction map. The general method is applied to global AdS, AdS-Rindler space, the Poincare patch, and Poincare-Milne space. In Section 4.3 we formulate the geometric diagnostic of capturing null geodesics and relate it to continuity of the reconstruction map, making contact with results in the mathematics literature. We apply the diagnostic to the black hole geometries, as well, without finding an explicit reconstruction map. In Section 4.4, we exhibit arguments that a subregion duality does exist for AdS-Rindler space. In Section 4.5, we note that this can be reconciled with the failure of continuous reconstruction from local fields if the duality involves nonlocal boundary operators in an essential way.

4.2 The Reconstruction Map

4.2.1 General Formulas

We begin this section by reviewing the procedure for obtaining a bulk solution from boundary data using eigenmodes of the wave equation, generalizing the approach of Ref. [89]. A classical, free bulk field Φ can be expanded in terms of orthonormal modes F_k which depend on a collection of conserved quantities k :

$$\Phi(B) = \int dk a_k F_k(B) + \text{c.c.} \quad (4.3)$$

Near the boundary, the modes F_k have the asymptotic form $F_k(B) \sim r^{-\Delta} f_k(b)$. Thus we find the boundary field $\phi = \lim_{r \rightarrow \infty} r^\Delta \Phi$ has the expansion

$$\phi(b) = \int dk a_k f_k(b) + \text{c.c.} \quad (4.4)$$

Given $\phi(b)$, we can ask whether it is possible to reconstruct $\Phi(B)$. Recall that $\phi(b)$ is dual to a one-point function in the CFT, hence this is equivalent to asking whether the bulk field is determined by CFT one-point functions. This is possible when the a_k can be extracted from ϕ through an inner product of the form

$$a_k = W_k \int db f_k^*(b) \phi(b) , \quad (4.5)$$

where W_k is a weighting factor. Equivalently, the boundary mode functions should satisfy the orthogonality relation

$$\int db f_k^*(b) f_{k'}(b) = W_k^{-1} \delta_{k,k'} . \quad (4.6)$$

There is no guarantee that a relation such as (4.6) will hold in general. We will see both possibilities in the examples below.

Given (4.5), it is a simple matter to solve for $\Phi(B)$:

$$\Phi(B) = \int dk \left[W_k \int db f_k^*(b) \phi(b) \right] F_k(B) + \text{c.c.} \quad (4.7)$$

We emphasize that at this stage (4.7) is, in principle, a recipe for computing the bulk field in terms of the boundary field.

However, there is an important simplification when the order of integration over k and b can be exchanged. Then we have

$$\Phi(B) = \int db K(B|b) \phi(b) , \quad (4.8)$$

where

$$K(B|b) = \int dk W_k f_k^*(b) F_k(B) + \text{c.c.} \quad (4.9)$$

This is a nontrivial simplification which does not occur in all cases. We will see below that when the order of integration is illegitimately exchanged, as in the example of the AdS-Rindler wedge, the integral over k in (4.9) does not converge [89].¹

Non-convergence of the integral in (4.9) is due to growth of the eigenmodes at large k . The large k behavior of the modes is closely related to the question of continuity of the reconstruction map, $\phi(b) \mapsto \Phi(B)$. To examine continuity, we need to adopt definitions for the bulk and boundary norms. On the boundary, we will follow Ref. [9] and use the norm

$$\|\phi\|_b^2 = \int db |\nabla_b \phi|^2 + |\phi|^2 . \quad (4.10)$$

Here $|\nabla_b \phi|^2$ is positive-definite, not Lorentzian, even though we are in a Lorentzian space-time. In other words, the norm looks like an integral of an energy density (over both space

¹With certain extra assumptions on the fields, however, [89] is able to construct a complexified smearing function.

and time), not an action. We will leave its exact form unspecified here, but will be explicit in the examples below. The correct norm to choose is an open question, and a different choice may affect the answer. Our choice is motivated by related results in the mathematics literature, but it may not be a natural choice for this problem. For now, this norm will serve to illustrate the possible answers to the continuity question. Because of (4.6), we will find that $\|\phi\|_b^2 \propto \int p(k)W_k^{-1}|a_k|^2$, where $p(k)$ is a quadratic polynomial in the conserved momenta.

In the bulk, a convenient and natural norm is given by the energy of the solution. Adopting the standard Klein-Gordon normalization for the modes $F_k(B)$, the energy is given by

$$\|\Phi\|_B^2 = E[\Phi] = \int dk |\omega(k)||a_k|^2, \quad (4.11)$$

where $\omega(k)$ is the frequency written as a function of the conserved quantities (one of which may be the frequency itself). The reconstruction map is continuous if and only if there is a constant $C > 0$ such that

$$\|\Phi\|_B^2 \leq C\|\phi\|_b^2. \quad (4.12)$$

That is, a bulk solution of fixed energy cannot have arbitrarily small imprint on the boundary. Equivalently, by going to momentum space, the product $\omega(k)W_k/p(k)$ must be bounded from above. In the remainder of this section we apply these general formulas to several specific cases to find smearing functions and check continuity. We restrict ourselves to a 2+1-dimensional bulk for simplicity.

4.2.2 Global AdS

The AdS_{2+1} metric in global coordinates is

$$ds^2 = -\frac{1}{\cos^2 \rho} dt^2 + \frac{1}{\cos^2 \rho} d\rho^2 + \tan^2 \rho d\theta^2. \quad (4.13)$$

The Klein-Gordon equation in these coordinates reads

$$-\cos^2 \rho \partial_t^2 \Phi + \frac{\cos^2 \rho}{\tan \rho} \partial_\rho (\tan \rho \partial_\rho \Phi) + \frac{1}{\tan^2 \rho} \partial_\theta^2 \Phi = m^2 \Phi. \quad (4.14)$$

The normalizable solutions are

$$F_{nl} = N_{nl} e^{-i\omega t} e^{il\theta} \cos^\Delta \rho \sin^{|l|} \rho \mathcal{F}_{nl}(\rho) \quad (4.15)$$

where

$$N_{nl} = \sqrt{\frac{\Gamma(n+|l|+1)\Gamma(\Delta+n+|l|)}{n!\Gamma^2(|l|+1)\Gamma(\Delta+n)}}, \quad (4.16)$$

$$\mathcal{F}_{nl}(\rho) = {}_2F_1(-n, \Delta+n+|l|, |l|+1, \sin^2 \rho), \quad (4.17)$$

and the frequency is $\omega = \Delta + 2n + |l|$. The boundary modes are

$$f_{nl} = \lim_{\rho \rightarrow \pi/2} \cos(\rho)^{-\Delta} F_{nl} = (-1)^n e^{il\theta - i\omega t} \sqrt{\frac{\Gamma(\Delta+n+|l|)\Gamma(\Delta+n)}{n!\Gamma^2(\Delta)\Gamma(n+|l|+1)}}. \quad (4.18)$$

Following the general procedure outlined above, we can compute the smearing function

$$K(\theta, t, \rho | \theta', t') = \sum_{n,l} \frac{1}{4\pi^2} \frac{\Gamma(\Delta)\Gamma(n+|l|+1)}{\Gamma(\Delta+n)} (-1)^n e^{-i\omega(t-t')} e^{il(\theta-\theta')} \cos^\Delta \rho \sin^{|l|} \rho \mathcal{F}_{nl}(\rho) + c.c. \quad (4.19)$$

This can be summed to obtain the result of Ref. [89].

The boundary norm in this case is given by²

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_{-T/2}^{T/2} dt d\theta \left((\partial_t \phi)^2 + (\partial_\theta \phi)^2 + \phi^2 \right) = 4\pi^2 \sum_{nl} (\omega^2 + l^2 + 1) \frac{\Gamma(\Delta+n+|l|)\Gamma(\Delta+n)}{n!\Gamma^2(\Delta)\Gamma(n+|l|+1)} |a_{nl}|^2. \quad (4.20)$$

The reconstruction map is continuous if and only if the following quantity is bounded:

$$\frac{\omega W_{nl}}{1 + \omega^2 + l^2} = \frac{\omega}{4\pi^2(1 + \omega^2 + l^2)} \frac{n!\Gamma^2(\Delta)\Gamma(n+|l|+1)}{\Gamma(\Delta+n+|l|)\Gamma(\Delta+n)}. \quad (4.21)$$

This ratio clearly remains finite for all values of n and l , thus proving continuity.

4.2.3 AdS-Rindler

We now turn to the AdS-Rindler wedge, which in 2+1 dimensions has the metric

$$ds^2 = \frac{1}{z^2} \left[- \left(1 - \frac{z^2}{z_0^2} \right) dt^2 + \frac{dz^2}{1 - \frac{z^2}{z_0^2}} + dx^2 \right]. \quad (4.22)$$

The Rindler horizon is located at $z = z_0$, while the AdS boundary is at $z = 0$. The Klein-Gordon equation is

$$-\frac{z^2}{1 - z^2/z_0^2} \partial_t^2 \Phi + z^3 \partial_z \left(\frac{1}{z} \left(1 - \frac{z^2}{z_0^2} \right) \partial_z \Phi \right) + z^2 \partial_x^2 \Phi = m^2 \Phi. \quad (4.23)$$

The normalizable solutions are

$$F_{\omega k} = N_{\omega k} e^{-i\omega t} e^{ikx} z^\Delta \left(1 - \frac{z^2}{z_0^2} \right)^{-i\hat{\omega}/2} {}_2F_1 \left(\frac{\Delta - i\hat{\omega} - i\hat{k}}{2}, \frac{\Delta - i\hat{\omega} + i\hat{k}}{2}, \Delta, \frac{z^2}{z_0^2} \right) \quad (4.24)$$

where $\hat{\omega} = \omega z_0$, $\hat{k} = k z_0$, and

$$N_{\omega k} = \frac{1}{\sqrt{8\pi^2 |\omega|}} \left| \frac{\Gamma(\frac{\Delta+i\hat{\omega}+i\hat{k}}{2})\Gamma(\frac{\Delta+i\hat{\omega}-i\hat{k}}{2})}{\Gamma(\Delta)\Gamma(i\hat{\omega})} \right|. \quad (4.25)$$

The boundary modes are then

$$f_{\omega k} = \lim_{z \rightarrow 0} z^{-\Delta} F_{\omega k} = N_{\omega k} e^{ikx - i\omega t}. \quad (4.26)$$

²In global coordinates, the norm in position space is properly defined as an average over time. This is related to the fact that the frequencies are discrete.

We can attempt to construct the smearing function following Eq. 4.9, but, as discussed below that equation, we will find that the integral over k does not converge:

$$K(x, t, z|x', t') = \tag{4.27}$$

$$\frac{1}{4\pi^2} \int dk d\omega e^{ik(x-x')} e^{-i\omega(t-t')} z^\Delta \left(1 - \frac{z^2}{z_0^2}\right)^{-i\hat{\omega}/2} {}_2F_1\left(\frac{\Delta - i\hat{\omega} - i\hat{k}}{2}, \frac{\Delta - i\hat{\omega} + i\hat{k}}{2}, \Delta, \frac{z^2}{z_0^2}\right) \tag{4.28}$$

$$= \infty . \tag{4.29}$$

This divergence is due to the exponential growth in k of the hypergeometric function when $k \gg \omega$ [89],

$${}_2F_1\left(\frac{\Delta - i\hat{\omega} - i\hat{k}}{2}, \frac{\Delta - i\hat{\omega} + i\hat{k}}{2}, \Delta, \frac{z^2}{z_0^2}\right) \sim \exp[\hat{k} \sin^{-1}(z/z_0)] . \tag{4.30}$$

The boundary norm is given by

$$\int dt dx ((\partial_t \phi)^2 + (\partial_x \phi)^2 + \phi^2) = \int d\omega dk 4\pi^2 N_{\omega k}^2 (1 + \omega^2 + k^2) |a_{\omega k}|^2 . \tag{4.31}$$

We see that the ratio which must be bounded in order that continuity hold is

$$\frac{\omega W_{nl}}{1 + \omega^2 + k^2} = \frac{2\omega^2}{1 + \omega^2 + k^2} \left| \frac{\Gamma(\Delta)\Gamma(i\hat{\omega})}{\Gamma(\frac{\Delta+i\hat{\omega}+i\hat{k}}{2})\Gamma(\frac{\Delta+i\hat{\omega}-i\hat{k}}{2})} \right|^2 . \tag{4.32}$$

This ratio remains bounded for fixed k , but when $k \gg \omega$ it grows like $\exp(\pi\hat{k})$. So we find both that the smearing function does not exist and that continuity fails.

Physical Interpretation In this case, the problem with reconstructing the bulk solution occurs regardless of the bulk point we are interested in. The discontinuity can be understood physically. At first, it is surprising that modes with $\omega < k$ are even allowed; in the Poincare patch, obtained as the $z_0 \rightarrow \infty$ limit of AdS-Rindler, they are not.³ Near the Rindler horizon frequency is redshifted relative to its value at infinity, while momentum is unaffected. So a local excitation with proper frequency comparable to its proper momentum appears at infinity as a mode with $\omega < k$. The modes with $\omega < k$ are confined by a potential barrier that keeps them away from the boundary; for large k the height of the barrier is proportional to k^2 . This causes the boundary data to be suppressed relative to the bulk by a WKB factor $\exp(-\int \sqrt{V}) \sim \exp(-\pi k)$.

We have seen that there is no smearing function in this case because a divergence at large momentum prevents us from exchanging the order of integration. To understand the physical meaning of this divergence, we can ask about computing a more physical quantity, which will regulate the divergence. Instead of trying to find an expression for the bulk field

³In Ref. [162], in the context of the BTZ black hole, it is suggested that these modes are connected with finite temperature effects, and the associated exponential factors are interpreted as Boltzmann weights. We consider this to be very suggestive, but have not found a concrete connection to this work.

at a specified bulk point, consider instead a bulk field smeared with a Gaussian function of some width σ in the transverse direction,

$$\Phi_\sigma(t, x, z) \equiv \int dx' \exp\left(-\frac{(x' - x)^2}{\sigma^2}\right) \Phi(t, x', z) . \quad (4.33)$$

We only smear in the x direction because the only divergence is in k , and we drop various numerical factors and polynomial prefactors that will be unimportant for our conclusion. We will also set $z_0 = 1$ (which is always possible by an appropriate scaling of coordinates) for the remainder of this section.

The smeared field has a perfectly fine expression in terms of local boundary fields. We can use symmetries to place the bulk point at $t = x = 0$; then

$$\Phi_\sigma(0, 0, z) = \int dt' dx' K_\sigma(0, 0, z|x', t') \phi(x', t') \quad (4.34)$$

with

$$K_\sigma(0, 0, z|x', t') = \int d\omega dk e^{i\omega t' - kx' - k^2 \sigma^2} (1 - z^2)^{-i\omega/2} {}_2F_1\left(\frac{\Delta - i\omega - ik}{2}, \frac{\Delta - i\omega + ik}{2}, \Delta, z^2\right) . \quad (4.35)$$

The important question is the large k behavior of this function. To get a feeling for it, replace the hypergeometric function by its large k limit,

$${}_2F_1 \approx g(\omega, \Delta, z) k^{\Delta-1} \cosh(2k\theta) \quad (4.36)$$

where θ depends on the distance from the boundary, $\sin \theta = z$, and g a function that does not depend on k . We ignore the polynomial prefactor and focus on the exponential dependence. Performing the integral, we get

$$K_\sigma(0, 0, z|x', t') \text{ " = " } \tilde{g}(t, \Delta, z) \exp\left(\frac{\theta^2}{\sigma^2} - \frac{x'^2}{\sigma^2} - 2i\frac{\theta}{\sigma^2}x'\right) \quad (4.37)$$

where the quotation marks indicate that this is only a cartoon of the correct answer that captures the large momentum behavior of the smearing function. Now we can write the smeared bulk field in terms of the boundary values,

$$\Phi_\sigma(0, 0, z) = \int dx' dt' K_\sigma(0, 0, z|x', t') \phi(x', t') . \quad (4.38)$$

What is the behavior of this function as we localize the bulk field by taking the width small, $\sigma \rightarrow 0$? K_σ is strongly dependent on σ : the maximum value of K_σ is exponentially large at small σ , $K_\sigma^{\max} = \exp(\theta^2/\sigma^2)$, where again θ is related to the distance from the horizon, ranging from $\theta = \pi/2$ at the horizon to $\theta = 0$ at the boundary. It varies rapidly, with characteristic wavenumber θ/σ^2 , and has a width set by σ .

The physical length over which the bulk point is smeared is $\sigma_{\text{phys}} = \sigma/z$, and up to an order-one factor we can approximate $\theta \approx z$. Restoring factors of the AdS radius L , we

find that the smeared smearing function K_σ is a rapidly oscillating function with maximum value

$$K_\sigma^{\max} \sim \exp\left(\frac{L^2}{\sigma_{\text{phys}}^2}\right). \quad (4.39)$$

Note that the dependence on the radial location has disappeared upon writing things in terms of the physical size. Attempting to measure the bulk field at scales smaller than the AdS radius requires exponential precision in the boundary measurement, because we are trying to compute an order-one answer (the bulk field value) by integrating an exponentially large, rapidly oscillating function multiplied by the boundary field value.

We note here an interesting technical feature of this construction. We chose to compute a bulk operator smeared with a Gaussian profile in the transverse direction. Normally, the exact form of a smeared operator is not physically relevant. In particular, we can ask whether it is possible to construct an analogous function K for smeared bulk operators which have smooth but compact support in the transverse direction. Unfortunately this is impossible. In order to overcome the exponential divergence at large k in the mode functions, we had to smear against a bulk profile which dies off at least exponentially fast at large k . Such a function is necessarily analytic in x , and hence will not have compact support. Therefore we cannot truly localize our smeared bulk operators in the above construction; some residual leaking to infinity is required.

4.2.4 Poincare Patch

The Poincare patch is the canonical example of a subregion duality that works. With our chosen norms, we will find that continuity actually *fails* in the Poincare patch, even though a smearing function exists. This suggests that the Poincare patch may already reveal subtleties that we claim exist in the AdS-Rindler case. However, we will see that the nature of the discontinuity is very different from that of the AdS-Rindler wedge. Later, in Section 4.3, we will argue that this discontinuity may be a harmless relic of our choice of norm, and that a more reliable answer is given by the geometric criterion presented there.

The metric of the Poincare patch is

$$ds^2 = \frac{dz^2 - dt^2 + dx^2}{z^2}, \quad (4.40)$$

and the Klein-Gordon equation in these coordinates reads

$$-z^2 \partial_t^2 \Phi + z^3 \partial_z \left(\frac{1}{z} \partial_z \Phi \right) + z^2 \partial_x^2 \Phi = m^2 \Phi. \quad (4.41)$$

In this case we label the eigenmodes by k and q , with $q > 0$. The frequency is given by $\omega = \sqrt{q^2 + k^2}$. Properly normalized, the modes are $F_{qk} = (4\pi\omega)^{-1/2} e^{ikx} z \sqrt{q} J_\nu(qz)$. We have introduced the notation $\nu = \Delta - 1 = \sqrt{1 + m^2}$. Then the boundary modes are

$$f_{qk} = \lim_{z \rightarrow 0} z^{-\Delta} F_{qk} = \frac{q^{\nu + \frac{1}{2}}}{2^\nu \Gamma(\Delta)} \frac{e^{i(kx - \omega t)}}{\sqrt{4\pi\omega}}. \quad (4.42)$$

The smearing function can easily be computed,

$$K(x, t, z|x', t') = \int dqdk \frac{2^\nu \Gamma(\Delta)}{4\pi^2 q^{\nu-1} \omega} e^{ik(x-x')} e^{-i\omega(t-t')} z J_\nu(qz) + \text{c.c.} , \quad (4.43)$$

and this matches with the result of Ref. [89].

The boundary norm is

$$\int dt dx ((\partial_t \phi)^2 + (\partial_x \phi)^2 + \phi^2) = \int dqdk \frac{\pi q^{2\nu}}{4^\nu \Gamma^2(\Delta)} (1 + \omega^2 + k^2) |a_{qk}|^2 . \quad (4.44)$$

The ratio which must remain bounded for continuity to hold is

$$\frac{\omega W_{qk}}{1 + \omega^2 + k^2} = \frac{4^\nu \Gamma^2(\Delta) \omega}{\pi q^{2\nu} (1 + \omega^2 + k^2)} . \quad (4.45)$$

For large q, k this remains bounded, but as $q \rightarrow 0$ it does not. The physics of the problem is the following. Starting with any solution, we can perform a conformal transformation that takes

$$z \rightarrow \lambda z, \quad x \rightarrow \lambda x, \quad t \rightarrow \lambda t . \quad (4.46)$$

For large λ , this moves the bulk solution towards the Poincare horizon and away from the Poincare boundary, resulting in a small boundary imprint. Under this scaling, $q \rightarrow \lambda^{-1} q$, so it is exactly the small q behavior above that allows for such an “invisible” solution.

As stated above, we believe that this discontinuity may merely be a problem of the choice of norm. In particular, this is an “infrared” discontinuity, and the difficulties of the AdS-Rindler wedge were “ultraviolet” in character. The smearing function seems to be sensitive only to the ultraviolet discontinuities, which suggests that those are more troublesome. Furthermore, in Section 4.3 we will see that the Poincare patch (marginally) passes the geometric test of continuity while the AdS-Rindler wedge clearly fails. A remaining problem for future work to provide a more concrete connection between “ultraviolet” and “infrared” discontinuities and the existence or non-existence of a smearing function.

4.2.5 Poincare-Milne

Poincare-Milne space is the union of the collection of Milne spaces at each value of z in the Poincare patch. It is useful to contrast the Poincare-Milne case with the AdS-Rindler case considered above. The reason is that the conformal boundary of Poincare-Milne space is identical to that of the AdS-Rindler wedge, but the Poincare-Milne bulk is larger, as shown in Fig. 4.2.⁴ We expect that the boundary theory of the AdS-Rindler boundary is dual to the AdS-Rindler bulk space and not more [40, 51, 103], and so it is an important check on our methods that they do not provide false evidence for a Poincare-Milne subregion duality. While we have no proof that the free theory constructions we have considered so far cannot be extended to Poincare-Milne, we can show that the most obvious construction breaks down in a very curious way.

⁴For definiteness we discard the future light cone of the point E in the figure, so that the boundary is exactly AdS-Rindler, with no extra null cone.




Figure 4.2: Here we depict Poincaré-Milne space, together with an AdS-Rindler space that it contains. The bulk of Poincaré-Milne can be defined as the intersection of the past of point A with the future of line BE . Clearly this region contains the AdS-Rindler space which is the intersection of the past of A and the future of B . Furthermore, the asymptotic boundary of the Poincaré-Milne space and the AdS-Rindler space is identical, being the causal diamond defined by A and B on the boundary.

The metric of Poincaré-Milne space is

$$ds^2 = \frac{dz^2 - dt^2 + t^2 dx^2}{z^2}, \quad (4.47)$$

where we restrict to $t > 0$. The Klein-Gordon equation in these coordinates reads

$$-z^2 t^{-1} \partial_t (t \partial_t \Phi) + z^3 \partial_z \left(\frac{1}{z} \partial_z \Phi \right) + z^2 t^{-2} \partial_x^2 \Phi = m^2 \Phi. \quad (4.48)$$

The z -dependence and x -dependence of the normalizable eigenmodes are identical to the Poincaré patch case, and the t -dependence comes from solving the equation

$$-t^{-1} \partial_t (t \partial_t \psi) - t^{-2} k^2 \psi = q^2 \psi. \quad (4.49)$$

The general solution to this equation is a linear combination of Hankel functions, $\psi = AH_{ik}^{(1)}(qt) + Be^{\pi k} H_{ik}^{(2)}(qt) = AH_{ik}^{(1)}(qt) + B[H_{ik}^{(1)}(qt)]^*$.

As we will demonstrate, no equation like 4.6 can hold for solutions to this equation. To see this, it is convenient to define $\tilde{\psi} = (qt)^{1/2} \psi$. Then we have

$$-\partial_t^2 \tilde{\psi} - \frac{k^2 + 1/4}{t^2} \tilde{\psi} = q^2 \tilde{\psi}. \quad (4.50)$$

This is a Schrödinger equation for a scattering state in an attractive $1/t^2$ potential. To simplify the calculation, we will normalize the solutions so that $A = 1$ always. The standard

expectation from quantum mechanics is that B is then completely determined as a function of q , and in particular we will only have a single linearly independent solution for a given value of q . However, from the bulk point of view there should always be two solutions for any q , corresponding to the positive and negative frequency modes. Indeed, the coefficient B is usually determined by the boundary condition $\tilde{\psi}(0) = 0$, but here that is trivially satisfied for all B . Hence B is a free parameter. We will now demonstrate another strange fact about this potential, that eigenmodes with different values of q are not orthogonal, which shows that Eq. 4.6 does not hold.

To see this, consider two solutions $\tilde{\psi}_1$ and $\tilde{\psi}_2$ corresponding to q_1 and q_2 . We have

$$(q_1^2 - q_2^2) \int_0^\infty dt \tilde{\psi}_1^* \tilde{\psi}_2 = \tilde{\psi}_1^* \partial_t \tilde{\psi}_2 - \partial_t \tilde{\psi}_1^* \tilde{\psi}_2 \Big|_0^\infty . \quad (4.51)$$

We can compute the inner product once we know the asymptotic behavior of the solutions near $t = \infty$ and $t = 0$.

First, we use the large argument asymptotic form of the Hankel function,

$$H_{ik}^{(1)}(qt) \approx \sqrt{\frac{2}{\pi qt}} e^{i(qt - \pi/4)} e^{k\pi/2} , \quad (4.52)$$

so that

$$\tilde{\psi}_i \approx \sqrt{\frac{2}{\pi}} e^{k\pi/2} \left(e^{i(q_i t - \pi/4)} + B_i e^{-i(q_i t - \pi/4)} \right) . \quad (4.53)$$

Then we find

$$\lim_{t \rightarrow \infty} \frac{1}{q_1^2 - q_2^2} (\tilde{\psi}_1^* \partial_t \tilde{\psi}_2 - \partial_t \tilde{\psi}_1^* \tilde{\psi}_2) = 2e^{\pi k} (1 + B_1^* B_2) \delta(q_1 - q_2) , \quad (4.54)$$

where we have used the fact that $\lim_{x \rightarrow \infty} e^{-iqx}/q = \pi\delta(q)$ and $\delta(q_1 + q_2) = 0$ when q_1 and q_2 are both positive. The result is proportional to a δ -function, as it had to be. For large t the solution approaches a plane wave, and plane waves of different frequencies are orthogonal.

Near $t = 0$ we use the small argument expansion

$$H_{ik}^{(1)}(qt) \approx \frac{1 + \coth \pi k}{\Gamma(1 + ik)} \left(\frac{qt}{2} \right)^{ik} - \frac{\Gamma(1 + ik)}{\pi k} \left(\frac{qt}{2} \right)^{-ik} , \quad (4.55)$$

so that

$$\tilde{\psi}_i \approx C_i t^{ik+1/2} + D_i t^{-ik+1/2} , \quad (4.56)$$

where C_i and D_i are determined in terms of B_i and q_i . Then we have

$$\lim_{t \rightarrow 0} \tilde{\psi}_1^* \partial_t \tilde{\psi}_2 - \partial_t \tilde{\psi}_1^* \tilde{\psi}_2 = 2ik (C_1^* C_2 - D_1^* D_2) . \quad (4.57)$$

In order to ensure orthogonality, this combination has to vanish for arbitrary choices of the parameters. This is clearly not the case. We note in passing that imposing an extra constraint of the form $D = e^{i\delta} C$, with δ a new independent parameter, will make the wavefunctions orthogonal. Tracing through the definitions, one can see that this also fully determines B in terms of q and δ , and that $|B| = 1$ as expected by unitarity. The choice of δ

corresponds to a choice of self-adjoint extension, necessary to make the quantum mechanics well-defined (for additional discussion of this point see Ref. [59], and see references therein for more on the $1/t^2$ potential in quantum mechanics). As we pointed out above, however, such a prescription is not relevant for our current task, as it would eliminate a bulk degree of freedom.

4.2.6 AdS-Rindler Revisited

We would like to emphasize that the above analysis of Poincare-Milne space is not a no-go theorem. As an example, we now show that AdS-Rindler space, analyzed in a certain coordinate system, suffers from the same pathologies. By a change of coordinates, one can show that the AdS-Rindler wedge can be written in a way that is precisely analogous to Poincare-Milne:

$$ds^2 = \frac{dz^2 - x^2 dt^2 + dx^2}{z^2}, \quad (4.58)$$

where we restrict to the region $x > 0$. Using this coordinate system, and following the usual procedure, we encounter problems very similar to those of Poincare-Milne space discussed above. The x -dependence of the eigenmodes is found by solving

$$-x^{-1} \partial_x (x \partial_x \psi) - x^{-2} \omega^2 \psi^2 = -q^2 \psi. \quad (4.59)$$

This is equivalent to a Schrödinger equation in the same potential as before, except now we are finding bound states instead of scattering states. The analysis is completely analogous to the Poincare-Milne case. There is a continuous spectrum of bound states (unusual for quantum mechanics!), and they are not generically orthogonal. Thus we cannot carry out the program of mapping boundary data to bulk solutions. In this scenario, the choice of a self-adjoint extension would involve quantizing the allowed values of q , and by restricting q correctly we can find a set of orthogonal states. While that is appropriate for quantum mechanics, here the bulk physics is well-defined without such a restriction.

4.3 A Simple, General Criterion for Continuous Classical Reconstruction: Capturing Null Geodesics

In this section, we propose a general, geometric criterion for classical reconstruction of the bulk from the boundary. To our knowledge, the case of AdS has not been analyzed explicitly. However, mathematicians such as Bardos *et al.* [9] have analyzed the analogous situation in flat spacetime: Consider a field that solves the classical wave equation in some region Ω of Minkowski space with a timelike boundary $\partial\Omega$, with Neumann boundary conditions everywhere on the boundary. Now suppose the boundary value of the field is given in some region $R \subset \partial\Omega$ of the boundary. When is this sufficient to determine the bulk field everywhere in Ω ?

The central result is that every null geodesic in Ω should intersect R in order for continuous reconstruction to be possible. The basic intuition is that if there is some null geodesic that does not hit R , then by going to the geometric optics limit we can construct solutions that are arbitrarily well localized along that geodesic. These solutions




Figure 4.3: This is one of many null geodesics which passes through AdS-Rindler space without reaching the AdS-Rindler boundary. The four highlighted points on the trajectory are (bottom to top) its starting point on the near side of the global boundary, its intersection with the past Rindler horizon, its intersection with the future Rindler horizon, and its endpoint on the far side of the global boundary.

are “invisible” to the boundary observer who only can observe ϕ in the region R , in the sense that the boundary imprint can be made arbitrarily small while keeping the energy fixed.

It is not surprising that capturing every null geodesic is a necessary condition for continuous reconstruction, and this will be the important point for us. In many situations, however, the null geodesic criterion is actually sufficient. As long as every null geodesic hits R , the entire bulk solution can be reconstructed. (The theorems are quite a bit more general than we have described here, applying to general second-order hyperbolic partial differential equations, and generalizing to nonlinear problems.)

We propose to extrapolate this condition to AdS and its asymptotic boundary, and subregions thereof. The statement is that continuous reconstruction of a bulk subregion is only possible if every null geodesic in that subregion reaches the asymptotic boundary of that subregion. Applying this to a small diamond on the boundary, we conclude that there is no bulk region for which boundary data on the small diamond can be continuously mapped to a bulk field. As shown in Fig. 4.3, it is possible to find a null geodesic through any bulk point that does not intersect a small diamond on the AdS boundary.

A rigorous generalization of the null geodesic criterion to the case of AdS is desirable. In global AdS, at least for the special case of the conformally coupled scalar field, the theorems of Ref. [9] are already strong enough in their current form to ensure continuity. That is because the problem is equivalent to a particular wave equation in a (spatially) compact region with boundary, i.e., the Penrose diagram. And indeed, there we found that the reconstruction was continuous in the way predicted by the theorems.

As stated above, a subtlety arises for the Poincare patch. From the point of view of null geodesics, the Poincare patch is a marginal case. In the Penrose diagram, the boundary of the Poincare patch seems to be just barely large enough to capture all null geodesics passing through the bulk. Why, then, did we find that the reconstruction map is discontinuous, in apparent violation of the theorems of Ref. [9]? In fact, the Poincare patch

just barely fails the criterion because the boundary region is not an open set, as required by the theorem that guarantees continuous reconstruction. We believe this may explain the “infrared” discontinuity we found, and we also believe that a different choice of norm could cure the problem. The existence of an explicit smearing function shows that the problems of the Poincare patch are not fatal.

AdS-Rindler space is of an entirely different character. As we mentioned above, it is clear that there are null geodesics which pass through the bulk and do not intersect even the closure of the boundary. We believe that this is why the discontinuity is in the “ultraviolet,” and also why the smearing function does not exist.

4.3.1 Unique Continuation, Null Geodesics, and RG Flow

There is another important physical question which brings null geodesics to the fore, and it is less subtle than continuity. The trouble with continuity, as we have seen, is that precise statements depend on a choice of boundary norm, and we have been unable to specify a natural choice for this problem. However, even without a boundary norm, we can ask the bulk question of unique continuation of a solution in the radial direction. In AdS/CFT, the radial evolution of the fields is related to a renormalization group flow of the CFT [5, 52, 61, 98]. Let r be a radial coordinate such that $r = \infty$ is the boundary, which represents the UV of the CFT. In the CFT, the IR physics is determined by the UV physics, which suggests that a bulk field configuration near $r = \infty$ can be radially evolved inward and determine the field configuration for all r . This intuition can be checked for any given proposed subregion duality.

It is a well-studied problem in mathematics to take a classical field, which solves some wave equation, specified in the region $r > r_*$ and ask if it can be uniquely continued to the region $r < r_*$. If we ask the question locally, meaning that we only ask to continue in a neighborhood of $r = r_*$, then the answer is simple and apparently very robust: the continuation is unique if and only if all null geodesics that intersect the surface $r = r_*$ enter the known region $r > r_*$. (This is usually stated by saying that the extrinsic curvature tensor of the surface, when contracted with any null vector, should have a certain sign.) The intuition here is the same as with continuous reconstruction: if a null geodesic grazes the surface but does not enter the region where we are given the solution, then we can construct geometric optics type solutions that are zero in the known region, but nonzero inside [174].

By this same reasoning, one might conclude that reconstruction from the boundary is not unique when there are null geodesics which avoid the boundary, as opposed to the reconstruction being merely discontinuous as stated previously. The resolution has to do with the technical definitions behind the phrasing, which differ slightly between the two questions. In the present context, the non-uniqueness of the solution comes from going all the way to the geometric optics limit along some geodesic which does not enter $r > r_*$. But this is a singular limit, and one might wish to exclude such configurations from being solutions to the equation. That is the choice we implicitly made in previous sections when we talked about continuity. Continuity is broken because of the same type of geometric optics solutions with a singular limit, but we do not have to include the limiting case itself; continuity only depends on the approach to the limit. So the null geodesic criterion, and

the reasoning behind it, is the same even though certain technical aspects of the description change based on convenience for the particular question being asked. The point of discussing unique continuation at all is that the boundary is not involved in the question, and so a boundary norm need not be chosen.

For the case of the AdS-Rindler wedge, the same analysis of null geodesics as above indicates that unique continuation fails as well. Knowing the solution for $r > r_*$ does not determine the solution for smaller r . Furthermore, the Poincare patch is again a marginal case for this question. Using the standard z coordinate, then for any z_* there are null geodesics which do not deviate from $z = z_*$.

4.3.2 The Diagnostic in Other Situations

To get a sense for how seriously to take our diagnostic, we can apply it to a variety of familiar situations to test its implications.

AdS black hole formed in a collapse Suppose we begin at early times with matter near the AdS boundary, and then at some later time it collapses to make a large black hole. In this case, every null geodesic reaches the boundary. For a given geodesic, just follow it back in time: at early times there is no black hole and no singularity, and we know that all null geodesics in AdS hit the boundary. So for a black hole formed in a collapse, every null geodesic is captured by the boundary, and it is likely that continuous reconstruction of the bulk is possible, both inside and outside the horizon.

Eternal AdS black holes and black branes In the case of an eternal black hole, there are some null geodesics that never reach the boundary; they go from the past singularity to the future singularity. The bulk can be continuously reconstructed from the boundary data only outside $r = 3G_N M$. (3 is the correct numerical factor in 3+1 dimensions. More generally, the bulk can be reconstructed down to the location of the unstable circular orbit.)

We will show this explicitly, focusing initially on a spherical black hole in 3+1 dimensions. The metric is

$$ds^2 = -f(r)dt^2 + \frac{dr^2}{f(r)} + r^2 d\Omega_2^2 \quad (4.60)$$

with $f(r) = 1 + r^2/L^2 - 2G_N M/r$. The null geodesics are extrema of the action

$$S = \int d\lambda g_{\mu\nu} \dot{x}^\mu \dot{x}^\nu = \int d\lambda \left(f\dot{t}^2 - \frac{\dot{r}^2}{f} - r^2 \dot{\Omega}^2 \right). \quad (4.61)$$

Identifying the conserved quantities $E = f\dot{t}$ and $l = r^2\dot{\Omega}$, the equation of motion can be read off from the condition that the worldline is null:

$$0 = g_{\mu\nu} \dot{x}^\mu \dot{x}^\nu = \frac{E^2}{f} - \frac{\dot{r}^2}{f} - \frac{l^2}{r^2}. \quad (4.62)$$

This derivation leads to a simple equation for null geodesics,

$$\dot{r}^2 + V_{\text{eff}}(r) = E^2 \quad \text{with } V_{\text{eff}} = \frac{fl^2}{r^2}. \quad (4.63)$$




Figure 4.4: On the left we show the effective potential for a null geodesic in a spherical black hole, and on the right the same for a planar black brane. In the case of a spherical black hole, there is a potential barrier which traps some null geodesics in the $r < 3G_N M$ region. Therefore continuous reconstruction from the boundary is not possible for the region $r < 3G_N M$. In the planar case, there are null geodesics reach arbitrarily large finite r without making it to the boundary. Hence there is no bulk region which can be continuously reconstructed from the boundary data.

The effective potential has a maximum at $r = 3G_N M$, independent of l (see Fig. 4.4). So null geodesics that begin outside this radius will inevitably reach the boundary, either in the past or the future. But there are null geodesics that exit the past horizon, bounce off the potential barrier, and enter the future horizon. Because of these, it will be impossible to reconstruct the bulk region near the horizon.

In this case, rather than conclude that there is anything wrong with the correspondence, the natural interpretation is that our classical analysis is breaking down. The “lost” null geodesics are being lost because they fall into the singularity. To recover this information, we will need to go beyond the classical approximation and resolve the singularity.

We can also ask about unique continuation. Starting with the data at large r , we can try to integrate in to find the solution at smaller r . This process will work fine down to $r = 3G_N M$. However, trying to continue the solution across $3G_N M$ will be impossible.

In the case of a black brane with a planar horizon in AdS_D , the effective potential for the null geodesics becomes

$$V_{\text{eff}} = a - \frac{b}{r^{D-1}} \quad (4.64)$$

where a and b are positive constants. Unlike the spherical black hole, there is no local maximum in the effective potential. For every value of r , there are null geodesics which exit the past horizon, travel to that value of r , then exit the future horizon. So there is no bulk region that can be continuously reconstructed from the boundary data.

General conclusion about black hole reconstruction In the cases of eternal black holes and black branes, the presence of singularities led to the existence of null geodesics which did not reach the boundary, and consequently regions of the bulk which could not

be reconstructed from the boundary data. This is not a sign that AdS/CFT is breaking down, but rather an indication that our classical reconstruction procedure is not valid. We know that classical physics breaks down in the neighborhood of the singularity, but the null geodesic criterion suggests that there is a problem even in low-curvature regions. Since the problematic null geodesics begin and end on singularities, it is possible that the physics of singularities needs to be resolved before this question can be answered. A second possibility is that nonlocal boundary operators in the CFT encode the physics of the missing bulk regions. As we emphasize in Section 4.4, this latter possibility is the expected outcome for AdS-Rindler space, where we believe there is an exact duality between particular bulk and boundary subregions.

4.4 Arguments for an AdS-Rindler Subregion Duality

In this section we exhibit several arguments in favor of a subregion duality for AdS-Rindler space, despite the failure of continuous reconstruction from local boundary fields.

4.4.1 Probing the Bulk

In the previous section, we asked whether we could classically reconstruct the bulk field $\Phi(B)$ from data on the boundary $\phi(b)$. In essence, we restricted ourselves to considering only one point functions $\langle \mathcal{O}(b) \rangle$ on the boundary, and sought to reconstruct bulk fields from integrals of these local boundary operators. However, from an operational standpoint, there is no reason to expect this to be the most efficient way of reconstructing the bulk in general. The boundary theory is equipped with many inherently nonlocal operators. For instance, higher point correlation functions such as $\langle \mathcal{O}(b_1)\mathcal{O}(b_2) \rangle$ could provide a much better probe of the bulk than one point functions.⁵

From a physical standpoint, basic properties of AdS/CFT and causality [165, 42] are enough to argue that the theory on the boundary diamond should be capable of reconstructing, at the very least, the AdS-Rindler bulk [40, 51, 103]. Consider a bulk observer Bob who lives near the boundary. The boundary theory should be able to describe Bob, and thus it would be inconsistent for Bob to have information about the bulk which the boundary theory does not. Since Bob can send and receive probes into regions of the bulk which are in the intersection of the causal future and causal past of his worldline, he can probe the entire bulk diamond. Thus, it should be the case that the entire bulk diamond can be reconstructed from data on the boundary.

The question of classical reconstruction—restricting to one-point functions on the boundary—amounts to only allowing Bob to make measurements of the field value at his location. If the value of the field decays rapidly near the boundary, Bob would need extremely high resolution to resolve the field. Allowing higher point functions on the boundary amounts to allowing Bob to send and receive probes into the bulk which directly measure

⁵It may be the case that higher point functions, which can be obtained by solving classical bulk equations of motion with quantum sources, encounter similar obstructions in the classical limit. However the boundary theory also contains many additional nonlocal operators, such as Wilson loops, which we expect to behave differently in this regime.




Figure 4.5: The geometry defining the Hartle-Hawking state for AdS-Rindler. Half of the Lorentzian geometry, containing the $t > 0$ portion of both the left and right AdS-Rindler spaces, is glued to half of the Euclidean geometry. The left and right sides are linked by the Euclidean geometry, and the result is that the state at $t = 0$ is entangled between the two halves.

the field away from the boundary. This could potentially be a far more efficient way of reconstructing the bulk.

4.4.2 Hyperbolic Black Holes

A CFT dual for AdS-Rindler arises as a special case of the AdS/CFT duality for hyperbolic black holes. The conformal boundary of AdS-Rindler can be viewed as the Rindler patch of Minkowski space, by Eq. 4.58. The CFT vacuum, when restricted to the Rindler patch, appears as a thermal Unruh state, indicating the presence of a thermal object in the bulk. Indeed, the AdS-Rindler metric in Eq. 4.22 with the replacement $z = 1/r$ is exactly the $\mu = 0$ case of the metric of the hyperbolic black hole studied in Ref. [57]:

$$ds^2 = - \left(\frac{r^2}{L^2} - 1 - \frac{\mu}{r^{d-2}} \right) dt^2 + \left(\frac{r^2}{L^2} - 1 - \frac{\mu}{r^{d-2}} \right)^{-1} dr^2 + r^2 dH_{d-1}^2, \quad (4.65)$$

where the spatial hyperbolic plane has the metric

$$dH_{d-1}^2 = \frac{d\xi^2 + dx_i^2}{\xi^2}. \quad (4.66)$$

A CFT dual for hyperbolic black holes follows from an adaptation of Maldacena's analysis of the eternal AdS black hole [130], which generalizes easily to the hyperbolic case. In particular, hyperbolic black holes have a bifurcate Killing horizon, allowing for a definition of a Hartle-Hawking state from a Euclidean path integral (Fig. 4.5). The boundary consists of two disconnected copies of $\mathbf{R} \times H^{d-1}$ (the boundary diamonds). The boundary Hartle-Hawking state is defined through a Euclidean path integral performed on $I_{\beta/2} \times H^{d-1}$, where $I_{\beta/2}$ is an interval of length $\beta/2$ and β is the inverse Rindler temperature.




Figure 4.6: The Penrose diagram for a hyperbolic black hole. In the $\mu = 0$ case, regions I and IV become the right and left AdS-Rindler wedges. In this case, the singularity is only a coordinate singularity, so the spacetime can be extended to global AdS.

(Of course, the Hartle-Hawking state for AdS-Rindler is equivalent to the global vacuum. This follows since $I_{\beta/2} \times H^{d-1}$ is conformal to a hemisphere, which is half of the boundary of Euclidean AdS.) The right and left wedges, regions I and IV in Fig. 4.6, are entangled:

$$|\psi\rangle = \sum_n e^{-\beta E_n/2} |E_n\rangle_R |E_n\rangle_L . \quad (4.67)$$

Restricting to only region I or IV therefore yields a thermal density matrix.

Excitations above the Hartle-Hawking vacuum can be constructed through operator insertions in the Euclidean geometry. In these states, all particles that enter and leave region I through the Rindler/hyperbolic black hole horizon will be entangled with particles in region IV. One may therefore question to what extent region I can be reconstructed without access to region IV. Small excitations above the Hartle-Hawking vacuum, with energy below the temperature $1/2\pi L$, will appear as an indiscernible fluctuation in the thermal noise when restricted to region I. More energetic states, however, are Boltzmann-suppressed. The density matrix in I will, to a good approximation, accurately register the presence of particles with energy above $1/2\pi L$. Hence, the boundary theory of region I, i.e., its density matrix, should encode at least the high-energy states in the bulk region I [50].

4.5 Discussion

If an AdS/CFT duality is to make sense physically, it should be the case that a physicist with a large but finite computer can simulate the CFT and learn something about the bulk. Knowing particular boundary observables to some accuracy should determine the bulk to a corresponding accuracy. In the case of global AdS/CFT, Hamilton *et al.* [89] found simple boundary observables—local, gauge invariant operators—which are sufficient to reconstruct the bulk. We have shown that this reconstruction is continuous, meaning that it is a physical duality in the above sense.

In the case of the proposed AdS-Rindler subregion duality, we have seen that these operators are not sufficient to perform the same task. We have shown that, given Eq. 4.10 as our choice of boundary norm, the classical reconstruction map in AdS-Rindler is not continuous. This indicates that we must specify the boundary theory to arbitrary precision to learn anything about the bulk, signaling a breakdown in the physicality of the correspondence.

It is true that our argument for the breakdown depends on the specific boundary norm we choose. We are always free to pick a different norm, for instance one which better respects the symmetries of the boundary theory, and it may be useful to investigate this possibility further. However, the null geodesic criterion gives a simple and intuitive picture of the failure of classical reconstruction, and we would find it surprising if a natural choice of norm could cure the difficulties.

The failure of our diagnostic does not necessarily signal the death of a AdS-Rindler subregion duality. The crucial point is that besides taking the classical limit, we additionally assumed that bulk operators could only be expressed as integrals of local boundary quantities. By removing this extra assumption, a full duality may be recovered—and it would seem surprising if, in general, local boundary quantities were always sufficient for classical reconstruction in all situations. The CFT contains many nonlocal operators, such as complicated superpositions of Wilson loops [148, 171], in addition to local ones. Our results suggest that these additional operators are necessary to see locality in the bulk, even in the classical limit.

Chapter 5

AdS Black Holes, the Bulk-Boundary Dictionary, and Smearing Functions

5.1 Introduction

The Lorentzian AdS/CFT [127, 185, 79] dictionary in extrapolate form gives a simple relation between a bulk operator close to the boundary, and a boundary operator. If $\Phi(B)$ is a bulk operator, where B denotes a bulk coordinate $B = (r, t, \Omega)$, and b is a boundary coordinate $b = (t, \Omega)$, then [6, 78, 91]

$$\lim_{r \rightarrow \infty} r^\Delta \Phi(B) = O(b). \quad (5.1)$$

This relates a local bulk operator at large r to a local boundary operator. But what is the CFT dual of $\Phi(B)$ at finite r ? A natural proposal is

$$\Phi(B) = \int db' K(B|b') O(b') + \mathcal{O}(1/N), \quad (5.2)$$

where $K(B|b')$ is some smearing function.

There is no reason a dictionary as simple as (5.2) has to be true. Our goal in this chapter will be to make progress on establishing when a mapping like (5.2) is and isn't possible. In pure global AdS, the smearing function $K(B|b')$ was found in [89]. We will show a smearing function as in (5.2) does not exist in AdS-black hole backgrounds, for *any* bulk point B .

Eq. 5.2 is in some ways an extraordinary statement. It allows us to express an n -point function of bulk operators Φ in terms of smeared n -point functions of boundary operators O . Thus it says that the entire bulk state is encoded in terms of operators O . Yet, the operators O are special: they are related via (5.1) to the large r limit of local bulk operators. Most field theory operators, for instance Wilson loops, are not of this form.¹

¹One should keep in mind that in (5.2) we are smearing on the boundary over both space and time; if one were to use the CFT Hamiltonian to evolve the right side of (5.2) to a single time then one would generate an operator with Wilson loops. Nevertheless, the ability to avoid Wilson loops if one is allowed to compute correlation functions of the O for different times is in itself nontrivial.

Having the precise form of the smearing function $K(B|b')$ is an important component of the AdS/CFT dictionary. For any bulk point B , $K(B|b')$ will presumably have most of its support on some subregion of the boundary. Thus, $K(B|b')$ would tell us, independent of the state, which subregion [40, 51, 103] of the global AdS boundary is “responsible” for a bulk point B .²

As a result of (5.1), Eq. 5.2 has a purely bulk interpretation. It states that a bulk operator at point B can be expressed in terms of smeared bulk operators at large radius. Alternatively, in the Schrödinger picture it states that the bulk wavefunction restricted to a large but fixed $r = R$ and provided for some sufficient time extent, completely encodes the bulk state for all $r < R$. It is not obvious if this is a true or false statement. If a particle is sitting at the center of AdS, its wavefunction at large r will be small but nonvanishing; perhaps that is enough to determine the wavefunction everywhere? Or perhaps there are some states for which the wavefunction has vanishingly small support at $r = R$, making (5.2) impossible?

While (5.2) is an operator statement, determining $K(B|b')$ is a classical field theory problem: Given $\varphi(b') \equiv \Phi(r = R, t', \Omega')$, how does one reconstruct $\Phi(B)$? This is a nonstandard boundary value problem, with data being specified on a timelike surface. However, having a smearing function is a more stringent requirement than simply having an algorithm for determining $\Phi(B)$ for any given $\varphi(b')$. For instance, it may be the case that for any particular bulk solution, even if $\varphi(b')$ is extremely small, one can pick an appropriate resolution so as to see it and reconstruct $\Phi(B)$. However, it could be that no matter how good a resolution one picks, there always exist field configurations having a near boundary imprint $\varphi(b')$ that is below the resolution scale. In such a case there wouldn't be a smearing function; for a smearing function implies a state-independent way of reconstructing. In a sense one has to pick the resolution beforehand without knowledge of which field configurations will be under consideration. As a CFT statement, the absence of a smearing function means that certain aspects of the bulk are not well encoded in the smeared CFT operators $O(b')$, but rather in the more general Wilson loops.

Constructing a smearing function is straightforward in static, spherically symmetric spacetimes. One solves the bulk equations of motion through a mode decomposition: $\Phi_{\omega lm}(r, t, \Omega) = \phi_{\omega l}(r) Y_{lm}(\Omega) e^{-i\omega t}$. The bulk is reconstructed mode by mode, using the boundary imprint to extract the coefficient of each mode. In some cases this can be used to construct a smearing function. However in other cases, for reasons discussed above and which we will make precise in Sec. 5.2, the candidate smearing function is a divergent sum. Pure AdS falls into the first category, while AdS-black holes are in the second.

In AdS, just like in flat space, at small r there is a centrifugal barrier which reflects the modes. However, black holes have the property that at a finite distance from the horizon the centrifugal barrier peaks and the potential dies off as the horizon is approached. Unlike in pure AdS, modes with $\omega \ll l$ become admissible, and are trapped behind the centrifugal barrier. As l is increased with ω kept constant, the barrier grows, and the imprint of the modes at large r decays exponentially in l .

In Sec. 5.2 we review the mode sum approach to obtaining a smearing function. In Sec. 5.3 we rewrite the Klein-Gordon equation for a scalar field in a static, spherically

²We should note that in general $K(B|b')$ will not be unique.

symmetric background as a Schrödinger equation. For large l , the potential has roughly two competing terms: the centrifugal barrier and the AdS barrier $\sim r^2$. For any radius r , no matter how large, there is an l sufficiently high so that the centrifugal barrier dominates. In Sec. 5.4 we use WKB to show that in AdS-black hole backgrounds this effect gives rise to the exponential behavior in l for the modes.

It may seem surprising that our ability to describe the bulk at large r , where the metric is well approximated by pure AdS, could be affected by the presence of a small black hole deep in the bulk. In Sec. 5.4.1 we show that while it is true the behavior of the modes near the boundary is always well approximated by Bessel functions $r^{-d/2} J_\nu(\sqrt{\omega^2 - l(l+d-2)}/r)$, the relation between ω and l depends on the entire bulk geometry. In pure AdS, ω is quantized as $\omega_n = 2n + l + \Delta$, while in the AdS-black hole ω is continuous and independent of l . As a result, $\omega \ll l$ is allowed for AdS-black holes, leading the Bessel function to have imaginary argument and correspondingly exhibit exponential growth.

In Sec. 5.5 we take the first steps towards generally establishing for which asymptotically AdS spacetimes a smearing function exists. In Sec. 5.5.1 we consider a static, spherically symmetric spacetime and argue that only the behavior of the high l modes is relevant for this question. We find in this limit the potential in the Schrödinger-like equation describing the modes simplifies significantly. We find that any barrier in the large l potential leads to exponential behavior in l of the modes and prevents a smearing function. Here the appropriate limit leading to an exponentially suppressed tail involves sending ω to infinity as well sending l to infinity, while keeping the ratio ω/l constant. Thus, we will find that even a small, dense star in AdS can prevent a smearing function from existing for some bulk points. However, unlike the black hole, we are not necessarily prevented from constructing a smearing function at large r in general. In Sec. 5.5.2 we consider general spacetimes and examine the possibility of the existence of trapped null geodesics (geodesics with neither endpoint on the boundary) as a proxy for the smearing function not existing. We find that in static spherically symmetric spaces the null geodesic equation is that of a classical point particle moving in a potential identical to the one found in Sec. 5.5.1 as being relevant for the smearing function question. Therefore a smearing function does not exist if there are trapped null geodesics.

5.2 Smearing functions

We work in Lorentzian $\text{AdS}_{d+1}/\text{CFT}_d$ with fixed boundary Hamiltonian, and correspondingly all nonnormalizable bulk modes turned off. We let B denote a bulk coordinate, $B = (r, t, \Omega)$, and b a boundary coordinate $b = (t, \Omega)$. If we consider a scalar field $\Phi(B)$ in the bulk, then excited states are obtained by acting with $\Phi(B)$ on the vacuum. As the boundary is approached, $\Phi(r \rightarrow \infty)$ will decay to 0. However, we can extract the leading term φ in the decay, $\Phi(B) \rightarrow \varphi(b)/r^\Delta$, where the conformal dimension $\Delta = d/2 + \sqrt{d^2/4 + m^2}$. The extrapolate version of the AdS/CFT dictionary instructs us to identify φ with a local boundary operator: $\varphi(b) \leftrightarrow O(b)$.

As a result we can construct a relation between $\Phi(B)$ and the CFT operators $O(b)$ by relating the tail φ of Φ at the boundary to $\Phi(B)$ through use of bulk equations

of motion. This is a nonstandard boundary value problem where the data is specified on a timelike surface. Unlike usual time evolution where the field at point is determined by the field in the causal past of that point, here we have little intuition about which portion of the boundary is needed to determine $\Phi(B)$.

Smearing function as a mode sum

In the limit of infinite N , the bulk field operator $\Phi(B)$ obeys the free wave equation and its reconstruction from boundary data can be implemented through Fourier expansion. Letting $\Phi_k(B)$ be the orthogonal solutions to the Klein-Gordon equation (where k is a collective index), we do a mode expansion of $\Phi(B)$ in terms of creation and annihilation operators a_k ,

$$\Phi(B) = \int dk a_k \Phi_k(B) + h.c. \quad (5.3)$$

Taking B to the boundary and letting $\varphi_k = \Phi_k r^\Delta$ gives

$$O(b) = \int dk a_k \varphi_k(b) + h.c. \quad (5.4)$$

In some cases the boundary mode functions $\varphi_k(b)$ are orthogonal. If they are we can invert (5.4)

$$a_k = \int db O(b) \varphi_k^*(b), \quad (5.5)$$

where we have with hindsight chosen to normalize the modes Φ_k so that φ_k are orthonormal.³ Inserting (9.12) into (5.4) gives

$$\Phi(B) = \int dk \left[\int db' \varphi_k^*(b') O(b') \right] \Phi_k(B) + h.c. \quad (5.6)$$

Exchanging the integrals over k and b gives

$$\Phi(B) = \int db' K(B|b') O(b'), \quad (5.7)$$

where

$$K(B|b') = \int dk \Phi_k(B) \varphi_k^*(b') + h.c. \quad (5.8)$$

Potential divergences of the smearing function

Eq. 6.6 is the equation for a smearing function and will be the focus of the rest of the chapter. In all the cases we will consider, (9.12) will exist, but the integral in (6.6) may or may not converge. In the limit of infinite N the question of the existence of a smearing function in some background can therefore be equivalently stated as the question of convergence of the integral in (6.6).⁴ In cases when a smearing function exists in the

³One is of course free to choose any normalization for the Φ_k ; however if the φ_k are not orthonormal, (6.6) will get modified by the appropriate factor.

⁴There is a potential loophole. The smearing function could exist without the integral (6.6) converging. If this were the case, the smearing function would have to be a function whose Fourier transform is not well-defined.

$N = \infty$ limit, one can then include corrections to (5.7) perturbatively in $1/N$ [107, 96]. We will only be concerned with the smearing function at infinite N .

The bulk modes Φ_k that appear in (6.6) need to be normalized so that the boundary modes φ_k they asymptote to are orthonormal. This is at the heart of the problem of constructing smearing functions. When we decompose some bulk solution $\Phi(B)$ in terms of modes, we generally don't expect each mode to be weighted equally. Rather, there are some modes which may have small coefficients. However, the smearing function is state independent and has no way of knowing which modes will get small weight.

When working in a general background, not all modes are equal. Some modes may need to pass through enormous barriers in the potential on their way to the boundary and consequently suffer a huge damping. All the modes Φ_k are normalized so that the $\varphi_k \equiv \Phi_k r^\Delta$ they asymptote to on the boundary are orthonormal. As a result, modes that had to pass through a large barrier will be extremely large at small r . For any particular solution this wouldn't bother us, as these modes would have a small expectation value of a_k . As a result, (5.6) would converge. However, without having the small a_k to dampen the modes at small r , the integral (6.6) appearing in (5.7) might diverge. As we will see later, this is precisely what happens in AdS-Rindler and in AdS-Schwarzschild.

Pure AdS smearing function

The metric for global AdS $_{d+1}$ can be written as

$$ds^2 = -(1+r^2)dt^2 + \frac{dr^2}{1+r^2} + r^2 d\Omega_{d-1}^2. \quad (5.9)$$

The smearing function $K(B|b')$ was constructed in [89] (see also [16, 90]).⁵ Notably, it has support on boundary points b' that are spacelike (or null) separated from B (shown in Fig. 5.1). It takes a different form in even and odd dimensions, and is simpler when $d+1$ is even:

$$K(B|b') = \left[\sqrt{1+r^2} \cos(t-t') - r \cos(\Omega - \Omega') \right]^{\Delta-d}. \quad (5.10)$$

The spacelike support of $K(B|b')$ gives it some peculiar features. If one uses $K(B|b')$ to construct $\Phi(B)$ through (5.7) and considers the limit of taking B to the boundary, it is not manifest that $\Phi(B) \rightarrow r^{-\Delta} O(b)$. In fact, UV/IR [165] seems to suggest one should only need some compactly supported portion of the boundary to construct $\Phi(B)$ if B is close to the boundary. However, the smearing function (5.10) does not reflect this intuition. Indeed, the limit of (5.10) when B is close to the boundary,

$$K(r \rightarrow \infty, t, \Omega | b') \rightarrow r^{\Delta-d} [\cos(t-t') - \cos(\Omega - \Omega')]^{\Delta-d} \quad (5.11)$$

is not at all peaked at small $t-t'$ and $\Omega - \Omega'$.

⁵To avoid any potential confusion, we note that in Lorentzian AdS/CFT the smearing function problem is distinctly different from the one Witten's bulk-boundary propagator [185] addressed in Euclidean space. The Lorentzian version of Witten's bulk-boundary propagator is a bulk Green's function with one point taken to the boundary (as a result, unlike (5.10) it does manifestly approach a delta function for coincident points). However, it is a smearing function for the nonnormalizable modes, which are dual to sources for the CFT; whereas we are interested in a smearing function for normalizable modes.




Figure 5.1: To construct the bulk operator $\Phi(B)$, the CFT operator $O(b')$ is smeared with the smearing function $K(B|b')$ as indicated in (5.7). (a) The support of the pure AdS smearing function $K(B|b')$ is all boundary points b' spacelike separated from B (hatched region). (b) Had the AdS-Rindler smearing function existed, it would have only made use of the boundary region that overlaps with $J^+(q) \cap J^-(p)$ (the intersection of the causal future of q and causal past of p), where q and p are chosen so that $J^+(q) \cap J^-(p)$ just barely contains B . Any changes outside this bulk region $J^+(q) \cap J^-(p)$ would have been manifestly irrelevant for computing $\Phi(B)$.

AdS-Rindler smearing function

Perhaps the smearing function (5.10) is not optimal and uses more boundary data than actually necessary? The minimal possible boundary region (shown in Fig. 5.1b) can be found by picking boundary points p, q such that B is just barely contained in the intersection of the causal future of q with the causal past of p , $J^+(q) \cap J^-(p)$. The intersection of $J^+(q) \cap J^-(p)$ with the boundary yields the smallest boundary region allowed by causality [26]. A convenient coordinate system to use which covers only this region is AdS-Rindler, which in AdS_3 takes the form

$$ds^2 = -(r^2 - 1)dt^2 + \frac{dr^2}{r^2 - 1} + r^2 dx^2. \quad (5.12)$$

In Ref. [89] construction of a smearing function of this form was attempted, but the procedure fails. The solution for the modes in terms of a hypergeometric function is [112]

$$\Phi_{\omega k}(r, t, x) = r^{-\Delta} \left(\frac{r^2 - 1}{r^2} \right)^{-i\omega/2} F \left(\frac{\Delta - i\omega - ik}{2}, \frac{\Delta - i\omega + ik}{2}, \Delta, \frac{1}{r^2} \right) e^{i(kx - \omega t)}. \quad (5.13)$$

For $k \gg \omega$ the modes have an exponential growth in k . As a result, the integral (6.6) doesn't converge. Note that although AdS-Rindler asymptotes to the Poincare Patch at large r , modes with $k \gg \omega$ are forbidden in Poincare Patch but allowed in AdS-Rindler.

Had an AdS-Rindler smearing existed, it would have guaranteed a smearing function for points B in the large r region of any asymptotically AdS geometry. The field and metric at any point outside of $J^+(q) \cap J^-(p)$ would have been manifestly irrelevant. In the absence of an AdS-Rindler smearing function, all we have is the global smearing function.

Since it makes use of the entire spacelike separated region from B , changes to the field anywhere in the bulk could potentially have an impact on reconstruction of $\Phi(B)$. While we wouldn't expect a small perturbation of the metric at the center of AdS to have a significant impact on the form of the smearing function, a black hole in the center is a major change to metric and the existence of a smearing function is no longer guaranteed.

Our goal will be to understand in which circumstances the smearing function does and doesn't exist; when (6.6) does and doesn't converge. In the following section we analyze the bulk modes in an AdS-Schwarzschild background.

5.3 Solving the wave equation

In this section we rewrite the wave equation for a scalar field in the form of a Schrödinger equation, allowing us to easily analyze the solutions.

We consider a scalar field $\Phi(B)$ in a static, spherically symmetric background,

$$ds^2 = -f(r)dt^2 + \frac{dr^2}{h(r)} + r^2 d\Omega_{d-1}^2. \quad (5.14)$$

To leading order in $1/N$ the field Φ satisfies the free wave equation

$$\frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu \Phi) - m^2 \Phi = 0. \quad (5.15)$$

Separating Φ as

$$\Phi(r, t, \Omega) = \phi(r) Y(\Omega) e^{-i\omega t} \quad (5.16)$$

gives for the radial field $\phi(r)$,

$$\frac{\omega^2}{f} \phi + \frac{1}{r^{d-1}} \sqrt{\frac{h}{f}} \partial_r (\sqrt{f h} r^{d-1} \partial_r \phi) - \frac{l(l+d-2)}{r^2} - m^2 \phi = 0. \quad (5.17)$$

Letting $\phi(r) = u(r)/r^{\frac{d-1}{2}}$ and changing variables to a tortoise-like coordinate $dr_* = dr/\sqrt{fh}$ turns (5.17) into a Schrödinger-like equation

$$\frac{d^2 u}{dr_*^2} + (\omega^2 - V(r)) u = 0, \quad (5.18)$$

with a potential

$$V(r) = f \left[\frac{(d-1)(fh)'}{4r} + \frac{(d-1)(d-3)}{4} \frac{h}{r^2} + \frac{l(l+d-2)}{r^2} + m^2 \right]. \quad (5.19)$$

We examine the form of the potential in global AdS and AdS-Schwarzschild:




Figure 5.2: The wave equation can be recast as a Schrödinger equation (5.18). We plot the global AdS_4 potential (5.21) for $l = 3$ for a massless field. The plot on the left is in terms of the radial coordinate r appearing in the AdS metric (5.20). The plot on the right is in terms of the tortoise coordinate r_* , and is the one relevant for solving (5.18). The two are related through $r = \tan r_*$. The tortoise coordinate has the effect of compressing the potential at large r , while leaving small r unaffected. The AdS barrier occurs at r_* very close to $\pi/2$; its narrowness allows the modes to decay only as a power law: $\phi \sim r^{-\Delta}$.

Global AdS

Global AdS_{d+1} has the metric

$$ds^2 = -(1+r^2)dt^2 + \frac{dr^2}{1+r^2} + r^2 d\Omega_{d-1}^2, \quad (5.20)$$

and correspondingly a potential

$$V_{Global}(r) = (1+r^2) \left[\frac{d^2-1}{4} + m^2 + \frac{(d-1)(d-3) + 4l(l+d-2)}{4r^2} \right]. \quad (5.21)$$

The potential for global AdS is plotted in Fig. 5.2. The potential is dominated at small r by the angular momentum barrier $l(l+d-2)/r^2$, and at large r by the AdS barrier proportional to r^2 . At intermediate radius, these terms balance and the potential attains a minimum set by the angular momentum l . The minimum of the potential, which at large l is approximately $l(l+d-2)$, sets the lower bound on ω .

AdS-Schwarzschild

AdS-Schwarzschild is of the form

$$ds^2 = - \left(1 + r^2 - \left(\frac{r_0}{r} \right)^{d-2} \right) dt^2 + \frac{dr^2}{1 + r^2 - \left(\frac{r_0}{r} \right)^{d-2}} + r^2 d\Omega_{d-1}^2. \quad (5.22)$$

giving a potential

$$V_{BH}(r) = \left[1 + r^2 - \left(\frac{r_0}{r} \right)^{d-2} \right] \left[\frac{d^2-1}{4} + m^2 + \frac{(d-1)^2 r_0^{d-2}}{4 r^d} + \frac{(d-1)(d-3) + 4l(l+d-2)}{4r^2} \right]. \quad (5.23)$$




Figure 5.3: A plot of the AdS_4 -Schwarzschild ($r_0 = 1$) potential (5.23) as a function of the radial coordinate r . The plot on the left is for $l = 4$, and the one on the right for $l = 10$. Unlike for pure AdS, ω is not bounded from below by l ; for a fixed ω , l can be arbitrarily high. The barrier an ω mode must pass through grows as l increases. This results in the $\omega \ll l$ modes having exponential behavior in l . Intuitively, these modes become ever more confined near the horizon with increasing l .

The potential for AdS-Schwarzschild is shown in Fig. 5.3. For large r the behavior is the same as for pure AdS. However, the behavior is different for r close to the horizon: the factor of f vanishes at the horizon, forcing the potential to vanish as well. The vanishing of the AdS-Schwarzschild potential at the horizon allows ω to be arbitrarily small, regardless of the value of l . This is in contrast with pure AdS.

BTZ

A nonrotating BTZ black hole has a metric similar to AdS-Rindler (5.12). In the form of (5.14), $f(r) = h(r) = r^2 - M$, giving a potential

$$V_{BTZ}(r) = (r^2 - M) \left(\frac{3}{4} + m^2 + \frac{l^2 + M/4}{r^2} \right). \quad (5.24)$$

The BTZ potential has similar properties to that of AdS-Schwarzschild.

5.4 Black hole smearing functions and large angular momentum modes

In this section we explain why global AdS admits a smearing function while AdS-black hole backgrounds do not. The reason is simple: if a black hole is present, modes with $l \gg \omega$ (and l arbitrarily large) become allowed. These modes are highly suppressed at large r by the centrifugal barrier. An attempt to calculate the smearing function via a mode sum immediately gives a divergence when performing the sum over l at a fixed ω . We will show there is no smearing function in two ways. First in Sec. 5.4.1 we show the existence of these $l \gg \omega$ modes in itself, independent of the details of the metric, prevents

a smearing function for bulk points at large r . In Sec. 5.4.2 we use WKB to directly solve for the modes, showing there is no smearing function for any bulk point B . We should note that there are other cases where a smearing function fails to exist, even without a horizon and the associated $l \gg \omega$ modes, as we will show in Sec. 5.5.

In Sec. 5.4.1 we review how in pure AdS modes oscillate as e^{iqz} where $z = 0$ is the boundary and $q^2 = \omega^2 - l(l + d - 2)$. If modes existed with $q^2 < 0$, then they would grow exponentially as $e^{\kappa z}$ where $\kappa^2 = -q^2$. Since black hole backgrounds asymptotically approach pure AdS, their $q^2 < 0$ modes will display this exponential behavior in l . Connecting with the discussion in Sec. 5.2, this means the sum in (6.6) will not converge, and hence these modes forbid a smearing function for bulk points near the boundary.

In Sec. 5.4.2 we show that in the limit of high l , the potential (5.19) considerably simplifies, with only the centrifugal barrier remaining. Using WKB we solve to find the modes. The result shows an exponential behavior in l for these $l \gg \omega$ modes, for any bulk point. Thus we find there is no smearing function for any point in any static spherically symmetric spacetime with a horizon.

5.4.1 Asymptotic behavior of the wave equation

For an asymptotically AdS space, at large r the $f(r)$, $h(r)$ in (5.14) have the limit $f(r)$, $h(r) \rightarrow r^2$. Changing variables to $z = 1/r$, we write the metric as

$$ds^2 = \frac{-dt^2 + dz^2 + d\Omega^2}{z^2}. \quad (5.25)$$

For small angles (5.25) resembles the metric of the Poincare Patch. The wave equation for the radial modes is

$$z^2 \phi'' - z(d-1)\phi' + (z^2 q^2 - m^2)\phi = 0, \quad (5.26)$$

where $q^2 = \omega^2 - l(l + d - 2)$. Substituting $\phi(z) = z^{d/2}\psi(z)$ yields

$$z^2 \psi'' + z\psi' + (z^2 q^2 - \nu^2)\psi = 0, \quad (5.27)$$

where we defined $\nu^2 = m^2 + d^2/4$. For $\omega^2 > l(l + d - 2)$ this gives $\psi(z) = J_\nu(qz)$ and hence

$$\phi(z) = z^{d/2} J_\nu(qz), \quad (5.28)$$

which resembles the usual solution in the Poincare Patch.⁶ Modes with $\omega^2 < l(l + d - 2)$ have negative q^2 . Of course these modes don't exist in pure AdS, but they do in AdS-Schwarzschild. Defining $\kappa^2 = -q^2$, we get $\psi(z) = J_\nu(i\kappa z) \equiv e^{i\nu\pi/2} I_\nu(\kappa z)$.⁷ Since the Bessel function $J_\nu(x)$ oscillates, $I_\nu(x)$ grows exponentially. We can see the exponential growth directly. Letting $\phi(z) = z^{\frac{d-1}{2}} u(z)$, (5.26) becomes

⁶The other solution, $Y_\nu(qz)$, is discarded because it doesn't have the correct behavior $\phi \rightarrow z^\Delta$ near the boundary that is expected of a normalizable mode.

⁷In the context of *Euclidean* AdS/CFT in the Poincare Patch one has this scenario of $q^2 < 0$. There are two solutions: $I_\nu(\kappa z)$ and $K_\nu(\kappa z)$. The $K_\nu(\kappa z)$ solution is kept while the $I_\nu(\kappa z)$ is discarded precisely because of its exponential growth in the bulk. Of course, for us $K_\nu(\kappa z)$ cannot be kept since it grows exponentially as the boundary is approached and so is nonnormalizable.

$$u'' - \left[\kappa^2 + \frac{\nu^2 - 1/4}{z^2} \right] u = 0 . \quad (5.29)$$

In the limit of $z \gg \sqrt{\nu^2 - 1/4}/\kappa$, Eq. 5.29 is solved by $u = e^{\kappa z}$. Thus, we see that for large l and $l \gg \omega$ the modes behave as

$$\phi(z) = z^{\frac{d-1}{2}} e^{lz} . \quad (5.30)$$

Our use of the pure AdS metric is justified for sufficiently small z . However for any $z > 0$, there exists an l sufficiently large such that $z \gg \sqrt{\nu^2 - 1/4}/\kappa$ is satisfied. Since computing a smearing function involves summing over l arbitrarily large, we are guaranteed to reach regime (5.30) at sufficiently high l .

Smearing function for static spherically symmetric spacetimes

In the case of static spherically symmetric spacetimes, the solutions (5.16) can be inserted into the smearing function (6.6), giving

$$K(r, t, \Omega | t', \Omega') = \int d\omega e^{i\omega(t-t')} \sum_{l, m_i} \phi_{\omega, l}(r) Y_{lm_i}(\Omega) Y_{lm_i}^*(\Omega') , \quad (5.31)$$

where m_i denotes all the angular quantum number, m_1, \dots, m_{d-2} , and we have normalized the time dependent piece with respect to the boundary Klein-Gordon norm. The radial modes $\phi_{\omega, l}$ are solutions to the radial wave equation (5.18) and should be normalized so that $\phi_{\omega l} \rightarrow r^{-\Delta}$ as $r \rightarrow \infty$. For AdS-Schwarzschild, the energies ω are continuous and so we have written an integral over ω ; for global AdS this would be replaced by a discrete sum over n as $\omega_n = 2n + l + \Delta$.

If only modes with $\omega > \sqrt{l(l+d-2)}$ are allowed then, as we saw above, the near boundary solution (5.28) is, when properly normalized,

$$\phi_{\omega l}(r) = 2^\nu \Gamma(\nu + 1) \frac{J_\nu(q/r)}{r^{d/2} q^\nu} . \quad (5.32)$$

Inserting the $\phi_{\omega l}(r)$ into (5.31), we see the sum converges. On the other hand, for modes with $\omega < \sqrt{l(l+d-2)}$, and in particular the high l ones with solution (5.30), the sum over l in (5.31) is hopelessly divergent.

5.4.2 Large angular momentum and WKB

Our goal here is to directly show the exponential behavior in l of the modes $\phi_{\omega l}(r)$ for large l and $l \gg \omega$ for any bulk point. The smearing function doesn't exist due to the modes with arbitrarily large l , which is why this is a sufficient limit to consider. We will also see how the details of the metric become irrelevant in the large l limit, with the centrifugal barrier dominating the potential (5.19).

Modes with energy ω have a turning point at $r = r_t$ which satisfies $V(r_t) = \omega^2$. In the limit of $l \gg \omega$, the turning point approaches the horizon, $r_t \approx r_h$. For $r > r_t$ the modes always have $\omega^2 < V$ and thus decay. For $r > r_t$, we can use WKB:

$$u(r) = \frac{1}{\sqrt{p}} \exp\left(-\int_{r_i}^r p \, dr_*\right), \quad (5.33)$$

where $p^2 = V - \omega^2$.

We will only be interested in the exponential term, so we drop the $1/\sqrt{p}$ prefactor. As discussed in Sec. 5.2, in order to compute the smearing function we need to normalize all the bulk modes so that their boundary limit (upon stripping off r^Δ) is normalized with respect to the boundary norm. In terms of $u(r)$, we need its coefficient to approach 1 as $r \rightarrow \infty$. Thus,

$$u(r) = \exp\left(\int_r^\infty \frac{dr'}{\sqrt{f(r')h(r')}} \sqrt{V(r') - \omega^2}\right), \quad (5.34)$$

where we used the relation between the radial coordinate and the tortoise coordinate, $dr_* = dr/\sqrt{fh}$.

The key point is that for any point outside the horizon, $r > r_h$, there is an l sufficiently large such that the potential (5.19) can be approximated by

$$V(r) = f \frac{l^2}{r^2}, \quad (5.35)$$

where for simplicity we used $l(l+d-2) \approx l^2$. In (5.34) it is sufficient to only integrate for some finite distance δ away from r to see the exponential behavior in l ,

$$u(r) > \exp\left(\int_r^{r+\delta} \frac{dr'}{\sqrt{f(r')h(r')}} \sqrt{V(r') - \omega^2}\right). \quad (5.36)$$

For any δ we want, there is an l sufficiently large such that the potential (5.19) can be approximated by (5.35) for all radii between r and $r + \delta$. Thus, using the approximate potential (5.35) and neglecting ω^2 we get,

$$u(r) > \exp\left(l \int_r^{r+\delta} \frac{dr'}{r' \sqrt{h(r')}}\right). \quad (5.37)$$

This demonstrates the exponential growth in l of the modes. This is true for any bulk point r ; the only difference is the larger r , the greater the l before the exponential growth (5.37) sets in.

In the limit of large r we can approximate $f(r) \approx h(r) \approx r^2$. This yields $u(r) \rightarrow e^{l/r}$. Recalling $\phi = u/r^{(d-1)/2}$, this reproduces (5.30). Additionally, (5.37) matches the exponential growth in l of the exact hypergeometric function solution (5.13) found in [112] for the BTZ black hole.

5.5 Smearing functions for other spacetimes

We have seen if there is a horizon the potential (5.19) vanishes at the horizon and consequently any frequency $\omega > 0$ is allowed. The arguments of the previous section show





Figure 5.4: We are interested in finding for which static spherically spacetimes without horizons a smearing function exists. The smearing function involves the sum (5.31) over modes, which can be grouped into 3 different regimes. Only *A* poses a threat to the convergence of (5.31). At large r the metric, and consequently the potential (5.19), looks like that of pure AdS (5.38). At smaller r , in regime *A* the angular momentum l is so large that all terms in the potential except for the centrifugal barrier (5.40) are irrelevant.

there is no smearing function. In this section we examine more generally when a smearing function exists. In Sec. 5.5.1 we consider a general static spherically symmetric spacetime (5.14), and find a simple criteria on the metric (5.41) which guarantees there will be modes with exponential behavior in l for high l , and hence there will not be a smearing function for some bulk points. In Sec. 5.5.2 we examine the possibility of trapped null geodesics as a proxy for a smearing function not existing for an arbitrary spacetime. In the special case of static spherically symmetric spacetimes we demonstrate that the existence of trapped null geodesics prevents a smearing function.

5.5.1 Static spherically symmetric spacetimes

In this section we examine the existence of a smearing function for spacetimes of the form (5.14) which do not possess horizons.

The question of the existence of a smearing function is the question of the convergence of the sum (5.31) over ω and l at a given value of r . To answer this question, we will need to estimate the size of each mode with a given ω and l , for every ω that is allowed. It is convenient to divide the (ω, l) -plane into three regions, shown in Fig. 5.4, according to the sizes of ω and l relative to certain large values ω_0 and l_0 which depend only on f and h and will be defined carefully below. Region *A* consists of all modes with $l > l_0$, region *B* consists of all modes with $\omega > \omega_0$ and $l < l_0$, and region *C* consists of the remaining modes with $\omega < \omega_0$ and $l < l_0$.



(a)

Figure 5.5: The wave equation can be recast as a Schrödinger equation (5.18) with a potential $V(r)$ and an energy ω^2 . Here we sketch a possible potential (5.19) for which a smearing function doesn't exist. At large r , $r > R$, the potential looks like that of pure AdS (the figure has been compressed; the distance between r_2 and R is really much larger). At smaller r the potential, for large l , is approximated by (5.40). If $f(r)/r^2$ ever has positive slope, as shown above, some of the modes ω (dashed line) will have to tunnel through the barrier. Consequently, the sum (5.31) will diverge for $r < r_2$.

Approximating the potential

To aid our calculation, we will approximate the behavior of the potential for large, small, and intermediate values of r .

Large r

Since the metric approaches pure AdS at large r , we can approximate $f(r) \approx h(r) \approx r^2$ for $r > R$, where R is some sufficiently large radius that depends on f and h . The potential for $r > R$ thus takes the form

$$V(r) \approx l(l+d-2) + \left(\frac{d^2-1}{4} + m^2 \right) r^2, \quad r > R. \quad (5.38)$$

Small r

As $r \rightarrow 0$, $h(r) \rightarrow 1$ to avoid a conical singularity, and $f(r) \rightarrow f_0 > 0$. Thus we can find some small $\epsilon > 0$ such that

$$V(r) \approx f_0 \left(\frac{l(l+d-2) + (d-1)(d-3)/4}{r^2} \right), \quad r < \epsilon. \quad (5.39)$$

Note that the form of the potential implies that ω is quantized. Aside from the constant f_0 , the potential for $r < \epsilon$ does not depend on the details of the geometry.

Intermediate r

For $\epsilon < r < R$ it will be useful to do a separate analysis for modes lying in the three different (ω, l) -regions shown in Fig. 5.4.

A: Since f , h , and their derivatives are bounded functions for $r < R$, we can find an l_0 sufficiently large so that for all $l > l_0$ all terms in the potential (5.19) except the centrifugal barrier are irrelevant for all $r < R$,

$$V(r) \approx f(r) \frac{l(l+d-2)}{r^2}, \quad r < R \quad \text{and} \quad l > l_0. \quad (5.40)$$

This potential agrees with our small- r approximation above when $r < \epsilon$. Note that the potential in this region has an overall scaling with l .

B: We choose an ω_0 sufficiently large such that $\omega_0^2 \gg V(r)$ for all $\epsilon < r < R$ and $l < l_0$. For modes in region B the potential is negligible at intermediate values of r .

C: For modes in region C all of the details of the potential are important, and there is no useful approximation.

Convergence of sum for smearing function

Since ω and l are quantized, there are only a finite number of modes in region C , so that part of the sum (5.31) converges. In region B the modes experience the same potential as in pure AdS (aside from an inconsequential scaling of f_0 at small r), so that part of the sum will converge as well. This only leaves region A to analyze.

Let us suppose the potential (5.40) has positive slope for some range of r ,

$$\frac{d}{dr} \left(\frac{f(r)}{r^2} \right) > 0 \quad \text{for some } r. \quad (5.41)$$

We will now show if this occurs then there is no smearing function for some bulk points due to an exponential divergence in region A . In Fig. 5.5 we sketch an example of potential for which (5.41) occurs. Consider the limit of large ω and large l . This is the classical limit of the Schrödinger equation (5.18), as can be seen from the fact that the range of r_* is finite and fixed, while the potential $V(r)$ and energy ω are getting large. Thus, it is guaranteed that there exists a mode ω lying within any classically allowed energy interval. If (5.41) is satisfied in a neighborhood of $r = r_1$, then that neighborhood consists of classical turning points for an interval of possible values of ω . Let $r_2 > r_1$ be any point in the classically forbidden region for these values of ω .⁸ Then the field at any r in the range $r_1 < r < r_2$ (or any r in the classically allowed region $r < r_1$) is larger than that at r_2 by a WKB factor of

$$\exp \left(\int_r^{r_2} \frac{dr}{\sqrt{f\hbar}} \sqrt{V - \omega^2} \right) = \exp \left(l \int_r^{r_2} \frac{dr}{\sqrt{f\hbar}} \sqrt{\frac{f}{r^2} - \left(\frac{\omega}{l} \right)^2} \right). \quad (5.42)$$

There is a subtlety here: in addition to this decaying solution, there is also an exponentially growing solution and the eigenstate will in general be a linear combination of the two if there is a second classically allowed region when $r > r_2$ (as in the scenario of Fig. 5.5). If both solutions contribute with comparable coefficients, then the eigenstate will not be exponentially larger at r than it is at r_2 as we have claimed. That phenomenon occurs, for instance, in the symmetric double well potential familiar from quantum mechanics. However, that kind of behavior is special to the symmetric, degenerate case. As long as the energy differences between approximate eigenstates localized on either side of the barrier is larger than the exponentially small tunneling factor, the true eigenstates of the system will be exponentially well localized on one side of the barrier, and we can restrict attention to those localized in the $r < r_1$ region.

⁸We treat r_1 and r_2 as if they are less than R for the purpose of approximating the potential. However, if they are larger than R it makes little difference. In the large- l limit the extra terms in (5.38) become irrelevant for $r < r_2$.

It is clear that (5.42) can be made arbitrarily large by making l large. Specifically, let $\alpha = \omega/l$ where ω and l are the modes considered above which are suppressed and give behavior (5.42). Now consider the portion of the sum (5.31) concentrated on the line of fixed α . Thus, (5.31) will not converge and there will not be a smearing function for points $r < r_2$. We should note that unlike the black hole case considered earlier, which did not have a smearing function for any bulk point, for a potential like in Fig. 5.5 there is a smearing function for bulk points at large r .

A remaining question is the converse of our statement: if $f(r)/r^2$ has non-positive slope for all r , is the smearing function guaranteed to exist? In this case there are no turning points at intermediate values of r , and hence no opportunity for exponential WKB factors. However, it is possible that the magnitude of the slope of $f(r)/r^2$ is small for some range of r , and then the $1/\sqrt{p}$ factor in (5.33) can become large. It is conceivable that one could still make sense of (5.31), despite power law growth in the summand, through regulation and analytic continuation. This requires further analysis, and is something we intend to investigate in subsequent work.

5.5.2 Trapped null geodesics

We have established a smearing function does not exist in a black hole background. In a general static spherically symmetric spacetime, we have shown that it does not exist if the metric has the property (5.41). In a more general spacetime without a high degree of symmetry, the mode sum approach to constructing a smearing function is inapplicable. This motivates us to search for a simply proxy for the existence of a smearing function. In this section we explore the following proposal: there is a smearing function iff all null geodesics have at least one endpoint on the boundary.

A smearing function allows one to make the statement (5.2) about the mapping between bulk and boundary operators. However, finding a smearing function is a classical field theory problem. At the level of individual modes we saw a smearing function for a point B in the bulk fails if there are modes whose imprint on the regulated boundary is exponentially small compared to their value B . Throughout this chapter we held the boundary imprint fixed and saw the value at B grow arbitrarily large, causing (6.6) to diverge. Keeping the field value at B fixed, this would correspond to modes with boundary imprint becoming arbitrarily small.

Since field configurations are built out of modes, we can state this as: a smearing function fails to exist if there are bulk field solutions $\Phi(B)$ with arbitrarily small boundary imprint. When one takes the geometric optics limit, field solutions become arbitrarily well localized along null geodesics. At a heuristic level, this motivates the simple criterion of trapped null geodesics, which we will now explore quantitatively.⁹

⁹In [9, 174] it was shown for a large class of hyperbolic differential equations that the diagnostic if reconstruction of a bulk field depends on boundary data continuously (for a particular choice of boundary norm) is that there not be any trapped null geodesics. In [31] null geodesics were explored in the context of subregion dualities due to these results and with the motivation of establishing if a collection of boundary observers can physically reconstruct the bulk field in a subregion of AdS. This question is of secondary concern to us here; our interest is rather in the nature of the bulk-boundary dictionary (if (5.2) is possible).




Figure 5.6: The equation for a null geodesic is that of a particle traveling in a 1-d potential (5.44). The potential is plotted for (a) pure AdS and (b) AdS-Schwarzschild ($M = 1$ in AdS_4). In pure AdS all null geodesics have an endpoint on the boundary, as can be seen from the figure on the left. This is in contrast to spacetime with horizons (right figure) which have some null geodesics which are trapped as a result of the potential U vanishing at the horizon. More generally, whenever there are trapped null geodesics, then there is no smearing function for some points in the bulk.

Geodesic equation

To find the motion of null geodesics in the spacetime (5.14), we note that the timelike Killing vector gives the conserved quantity $E = f\dot{t}$, and the Killing vector in one of the angular directions, θ , gives $L = r^2\dot{\theta}$. Here we are using the notation $\dot{x}^\mu \equiv dx^\mu/d\lambda$ where λ is the affine parameter.

The geodesic equation can be written as

$$\frac{f}{h}\dot{r}^2 + L^2\frac{f}{r^2} = E^2. \quad (5.43)$$

This is just the Newtonian energy conservation equation for a particle with position-dependent mass moving in a potential

$$U = L^2\frac{f}{r^2}. \quad (5.44)$$

In black hole backgrounds the potential (5.44) vanishes at the horizon, leading to trapped null geodesics. More generally, (5.44) tells us there are no trapped geodesics iff $U' < 0$ for all r .

It is interesting to note that even a small dense star in AdS can have trapped null geodesics. All the star needs is to have a radius \mathcal{R} which lies in the range $2M < \mathcal{R} < 3M$. The metric for $r > \mathcal{R}$ is described by the Schwarzschild metric. Since for the Schwarzschild metric $U'(r) > 0$ for $r < 3M$, null geodesics will get trapped at small r .

Trapped null geodesics \Rightarrow no smearing function

In 5.1 we saw that for the question of the existence of a smearing function, only the behavior of the high l modes was relevant. In this regime (labeled A in Fig. 5.4)

the potential in the Schrödinger equation for the modes was well approximated by (5.40). Yet this is exactly the same as the classical particle potential (5.44) for a null geodesic. Our condition for a smearing function not existing (5.41) is the same as the condition for the existence of trapped null geodesics. Thus we conclude that if there are trapped null geodesics, then a smearing function does not exist.

We note that since null geodesics are only sensitive to the local metric, trapped null geodesics cannot tell us for which regions of the bulk there is no smearing function. If a null geodesic is confined to $r < r_t$, this indicates there is no smearing function for $r < r_t$, but it says nothing about a smearing function for $r > r_t$. The existence of a smearing function for $r > r_t$ depends, as we explained above, on the existence of additional classically allowed regions with $V(r) < \omega^2$ for $r > r_t$. A classical null geodesic confined to $r < r_t$ cannot probe these aspects of the potential.

5.6 Conclusions

In this chapter we have further explored one of the approaches to establishing the dictionary between bulk and boundary operators. In this approach, a bulk operator is expressed in terms of bulk operators at asymptotically large radius, which are then mapped to local boundary operators through (5.1). While this approach works in pure global AdS, we have argued it can fail if there are bulk modes which have an arbitrarily small tail at large radius. We have shown that AdS-Schwarzschild backgrounds are a case where this smearing function approach fails as a result of modes with arbitrarily high angular momentum l , but fixed boundary energy, ω .

Understanding in general when a spacetime has a region for which a bulk operator cannot be expressed in terms of smeared local boundary operators remains an important future problem. We have shown that for static, spherically symmetric spacetimes, this question can be answered by considering the behavior of modes with large ω and large l . These modes satisfy a Schrödinger-like equation with a potential that is the same as the potential experienced by classical null geodesics. These results suggest that a smearing function may not exist for some bulk points in any spacetime which has trapped null geodesics.

The extent to which the absence of a smearing function modifies the bulk-boundary dictionary remains to be seen. It is possible one can obtain an approximate smearing function by imposing a cutoff in the bulk and excluding high l and high ω modes from the sum (5.31) defining the smearing function. Additionally, as we discussed, the existence of a smearing function is a more stringent requirement than simply being able to reconstruct a bulk field solution given some particular boundary data. To this extent, even though (5.7) may not exist, (5.6) exists at $N = \infty$. However, unlike a smearing function, it is unclear that (5.6) can be generalized to situations with broken spherical and time-translation symmetry, and so may be of limited use. Another option would be to try to construct a smearing function which uses the complexified boundary, as done in [89] for AdS-Rindler, and perhaps this would give clues as to what the real spacetime representation of the bulk-boundary map is.

When a smearing function does exist, it means that bulk data provided at large radius for a sufficient time extent completely determines the bulk everywhere. In a way, it

makes holography seem less powerful; a spatial direction has just been replaced by a time direction. Of course, the power of AdS/CFT is due to the CFT Hamiltonian which one can use to evolve the right side of (6.6) to a single time. The resulting operator is highly nonlocal and known as a precursor [148]: an operator which encodes what happened deep in the bulk long before causality allows a local operator $O(b)$ to know about it.

The absence of a smearing function makes the holographic dictionary more elusive. Bulk evolution combined with boundary evolution are not even in principle sufficient to answer the question of what the precursors are. Other methods must be developed to find the dictionary between bulk operators and nonlocal boundary operators.

Chapter 6

Scanning Tunneling Macroscopy, Black Holes, and AdS/CFT Bulk Locality

6.1 Introduction

This chapter concerns reconstructing the state of the anti-de Sitter (AdS) bulk from the conformal field theory (CFT) boundary. Finding which CFT quantities encode the bulk and, if so, in what ways has been actively pursued recently, e.g. [106, 97, 152, 111, 116, 60, 58, 102, 4, 31, 118, 110]. One standard approach is to start with local CFT data and use bulk equations of motion to evolve radially inward. The questions we seek to answer are: under what circumstances is this evolution possible, how deep into the bulk can we evolve, on what scales can the bulk be reconstructed, and what assumptions about the bulk state must be made?

We first address these questions in a simpler context: the electromagnetic field in Minkowski spacetime. Given boundary field data on a *timelike* codimension-1 hypersurface ($z = 0$), can the electromagnetic field be determined everywhere ($z > 0$)? Suppose the bulk is filled with homogeneous air. The only solutions to the wave equation consistent with translation symmetries are traveling waves, and reconstruction is trivially achieved. On the other hand, suppose that the bulk is filled with air for $0 \leq z < z_g$ and with glass for $z > z_g$. The translation symmetry being broken, there is now a new class of solutions: waves which are traveling for $z > z_g$ but evanescent for $0 \leq z < z_g$. Evanescent modes are solutions with imaginary momentum. While legitimate solutions of the wave equation, these modes are forbidden in vacuo because of exponential growth at large z and hence non-normalizability. The presence of glass at finite z cuts off this unboundedness and renders the mode permissible. Reconstructing the field inside the glass from measurements at $z = 0$ is hopeless; a small mistake will get exponentially amplified. It would be like trying to measure the electromagnetic field inside a waveguide while standing a kilometer away. In fact, even reconstruction of the field for $0 \leq z < z_g$ from the boundary data at $z = 0$ is no longer straightforward. If we do not measure the evanescent modes (in the absence of assumptions on the form of the solution), we can not reconstruct the field anywhere. We

can, however, measure the evanescent modes at $z = 0$ to some extent without resorting to exponential precision. If we do then the field, coarse-grained in x over a scale σ , can be reconstructed, but only for $z < \sigma$.

We next address the reconstruction question in the context of the AdS/CFT correspondence: we consider a free scalar field in a fixed asymptotically AdS background. A background which is pure AdS, or a perturbation thereof, is like the electromagnetic field in the vacuum - exact reconstruction using boundary data is possible. Thus, local CFT operators give a probe of the bulk on the shortest of scales. A background with a small AdS black hole is like putting in a region with glass. The geometry in the black hole atmosphere, the region from $2M$ to $3M$, changes the boundary conditions and permits evanescent modes. Analogously to the case with glass, reconstructing the field in the atmosphere is hopeless. Reconstructing the field far from the black hole is possible to some extent, and like in the case of the electromagnetic field, requires measuring the evanescent data. The condition $z < \sigma$ translates into the ability to resolve the bulk on AdS scales and no shorter. The measurement of evanescent modes forms the basis of the functioning of a Scanning Tunneling Microscope (STM). In this sense, the CFT is acting like an STM for macroscopic scales.

It may seem puzzling that a small black hole deep in the bulk should have any impact on our ability to reconstruct the bulk close to the boundary. One may pretend the evanescent modes don't exist and work only with the propagating modes. This might be a good approximation for some states and for regions close to the boundary, but it is one that is violated by legitimate finite energy solutions having a significant coefficient for an evanescent mode. Even in the Hartle-Hawking vacuum, the CFT Green's function $G_2(\omega, \mathbf{k})$ (given by the Fourier transform of the finite temperature two-point correlator $\langle O(t, \mathbf{x})O(0, \mathbf{0}) \rangle_T$) is nonzero but exponentially small $\sim \exp(-\tilde{\alpha}|\mathbf{k}|/T)$ in the evanescent regime $k \gg \omega$. The evanescent modes are part of the spectrum of states, so even if one might opt to ignore them for two-point correlators, they necessarily contribute to any finite-temperature n -point function as intermediate states.

We have organized the chapter as follows. In Section 6.2, we formulate in Minkowski spacetime the problem of spacelike reconstruction from timelike boundary data. We show that reconstruction is exact in situations with only propagating modes but requires exponential precision in situations with evanescent modes. In the latter situation, reconstruction without exponential precision is possible but only at the cost of averaging over directions parallel to the boundary, and only out to a distance set by this averaging scale. In Section 6.3, we formulate the reconstruction problem in AdS space and demonstrate that the situation is exactly parallel to the Minkowski spacetime counterpart. In pure AdS, all modes are propagating and the reconstruction is exact. In an AdS black hole background, evanescent modes open up near the black hole horizon and reconstruction requires exponential precision. Here again, reconstruction without exponential precision is possible but only at the cost of an AdS scale averaging over directions parallel to the boundary. In Section 6.4, via the AdS/CFT correspondence, we discuss the impact of evanescent modes on bulk reconstruction from the CFT viewpoint. We show that in general precise determination of a Green's functions at finite temperature requires exponential precision. In the Appendix 6.6 we review evanescent modes in optics, the principles of a microscope, and scanning tunneling (optical) microscopy (STM).




Figure 6.1: The setup of bulk reconstruction in $(2+1)$ -dimensional Minkowski half-space $R_+^{2,1}$. We have a bulk field $\phi(x, t, z)$ obeying the wave equation, which we wish to reconstruct from the data $\phi(x, t, 0)$ on a timelike hypersurface at $z = 0$.

6.2 Bulk reconstruction in Minkowski space

In this section we pose the question: for a field $\phi(x, t, z)$ satisfying the wave equation in $(2+1)$ -dimensional Minkowski spacetime¹, is the boundary data $\phi(x, t, 0)$ specified at a timelike hypersurface $z = 0$ sufficient to reconstruct the field $\phi(x, t, z)$ everywhere? (See the setup shown in Fig. 6.1.) The wave equation admits solutions with both real and complex momentum. The solutions with real momentum are the propagating waves, $e^{\pm ik_z z}$. If the space is everywhere homogeneous, such as pure Minkowski space, then these are the only admissible delta-function normalizable modes. In this case, reconstruction of $\phi(x, t, z)$ from the boundary data $\phi(x', t', 0)$ works perfectly — we show in Sec. 6.2.1 that the smearing function $K(x, t, z|x', t')$ whose convolution with $\phi(x', t', 0)$ yields $\phi(x, t, z)$ is well-defined everywhere in the bulk. On the other hand, if the space is inhomogeneous by, for instance, having a spatially varying index of refraction, then modes with imaginary momentum can become permissible. Instead of propagating, these modes grow exponentially in the z direction: $e^{\pm \kappa_z z}$. They are known as “evanescent modes”. Our goal is to point out that the evanescent modes cause serious difficulties in reconstructing the field anywhere from given boundary data.

6.2.1 Success with propagating modes

We consider the wave equation in $(2+1)$ -dimensional Minkowski spacetime $R^{2,1}$:

$$\left(-\partial_t^2 + \partial_x^2 + \partial_z^2\right) \phi(x, t, z) = 0. \quad (6.1)$$

We will be interested in reconstructing $\phi(x, t, z)$ from boundary data $\phi(x, t, 0)$ specified on a timelike hypersurface b at $z = 0$. To accomplish this, we decompose the solutions to (6.1) in terms of the propagating wave basis:

$$\phi(x, t, z) = \int \int dk_x d\omega \phi(k_x, \omega) e^{-ik_x x - ik_z z - i\omega t}, \quad (6.2)$$

¹Extension to higher dimensions is straightforward and does not reveal any new physics.

where the delta-function normalizability condition puts

$$k_z \equiv \sqrt{\omega^2 - k_x^2} \in R^+ \quad \rightarrow \quad |\omega| \geq |k_x|. \quad (6.3)$$

We chose k_z to be positive, as we will for simplicity assume there are only left-movers.² The decomposition (6.2) is slightly nonstandard as ω is one of the independent variables as opposed to k_z . This is a convenient choice here as the boundary data is specified on a timelike hypersurface. The Fourier transform of $\phi(x, t, z = 0)$ yields

$$\phi(k_x, \omega) = \int \int_b dx' dt' e^{ik_x x' + i\omega t'} \phi(x', t', 0). \quad (6.4)$$

Inserting (6.4) into (6.2),

$$\phi(x, t, z) = \int \int_b dx' dt' K(x, t, z|x', t') \phi(x', t', 0). \quad (6.5)$$

Here, K is the smearing function given by

$$\begin{aligned} K(x, t, z|x', t') &= \int_{-\infty}^{\infty} d\omega \int_{|k_x| \leq |\omega|} dk_x e^{-i\omega(t-t')} e^{-ik_x(x-x')} e^{-i\sqrt{\omega^2 - k_x^2} z} \\ &= \int_{-\infty}^{\infty} dk_x \int_{|\omega| \geq |k_x|} d\omega e^{-i\omega(t-t')} e^{-ik_x(x-x')} e^{-i\sqrt{\omega^2 - k_x^2} z}. \end{aligned} \quad (6.6)$$

The integral in (6.6) is convergent, so (6.5) realizes our goal of reconstructing $\phi(x, t, z)$ in terms of the boundary data $\phi(x, t, 0)$. This is not surprising, as for every mode the field at some value of z is related to the field at $z = 0$ by the phase-factor $e^{ik_z z}$ associated with the translation, where k_z is given by (6.3).

It will be useful for later to first reconstruct per each monochromatic ω mode and then combine them together. We then use

$$\phi_\omega(x, 0) \equiv \int_{-\infty}^{\infty} dt e^{i\omega t} \phi(x, t, z = 0) \quad (6.7)$$

to reconstruct $\phi_\omega(x, z)$,

$$\phi_\omega(x, z) = \int_{-\infty}^{\infty} dx' K_\omega(x, z|x') \phi_\omega(x', 0) \quad (6.8)$$

where

$$K_\omega(x, z|x') = \int_{|k_x| \leq |\omega|} dk_x e^{-ik_x(x-x')} e^{-i\sqrt{\omega^2 - k_x^2} z}. \quad (6.9)$$

The full smearing function is then recovered through the spectral sum

$$K(x, t, z|x', t') = \int_{-\infty}^{\infty} d\omega K_\omega(x, z|x') e^{-i\omega(t-t')}. \quad (6.10)$$

²Assuming there are only left-movers allows us to connect more directly with the analogous problem in AdS space. The more general case in Minkowski space for which right-movers are also included is a straightforward extension and requires specifying $\partial_z \phi(x, t, z = 0)$ in addition to $\phi(x, t, z = 0)$ at the boundary b .

6.2.2 Reality with evanescent modes

In the above discussion we restricted k_z to be real-valued as a consequence of the delta-function normalizability condition for the modes. If one were to consider $k_x > \omega$ such that k_z is imaginary: $k_z \equiv i\kappa_z$, a mode would take the form

$$e^{-ik_x x} e^{-i\omega t} e^{\kappa_z z}. \quad (z \geq 0) \quad (6.11)$$

This solution diverges at large z , and is therefore not admissible. We were thus correct to discard it. However, if the Minkowski space is inhomogeneous in the z -direction, often solutions such as (6.11) become admissible. In Appendix 6.6, we give two situations which generate evanescent modes: (1) a wave traveling in a medium with a z dependent index of refraction, and (2) a wave scattering off a material that has an x -dependent transmission coefficient and is located at some fixed z .

If evanescent modes are present, then they pose a serious challenge for reconstructing the field from the $z = 0$ boundary data. Unlike propagating modes, evanescent modes will have an exponentially suppressed imprint on the $z = 0$ boundary compared to their value at $z > 0$. This exponential behavior renders the reconstruction procedure of Sec. 6.2.1 inapplicable: the smearing function K (6.6) is ill-defined because of the divergence from the $k_x \gg \omega$ region of integration. For the rest of this section, we simply assume that evanescent modes are present and do not inquire as to their origin. An exemplary situation to keep in mind is the one where the index of refraction of the background changes at some large value of $z > 0$. For simplicity, we will assume the change made at large z is such that all possible evanescent modes are produced so that the mode solutions now contain all values of k_x regardless of ω . We would like to understand the impact this has on reconstructing the field from boundary data.

We can take two approaches in dealing with evanescent modes. We may ignore the modes from the outset, contending that they are not propagating. Or we may include the modes. We now argue that in both cases we will face unavoidable limitations on the resolving power of the reconstruction.

6.2.3 Ignoring evanescent modes

Consider first the approach of ignoring the evanescent mode data at the boundary: we will not try to extract their coefficient from the boundary data provided at $z = 0$. In the procedure of Sec. 6.2.1, we are now in the situation that the Fourier decomposition of $\phi(x, t, z)$ (6.2) extends over all k_x and ω , but we truncate the Fourier decomposition of the smearing function K in (6.6) to $|\omega| \geq |k_x|$. Reconstruction of $\phi(x, t, z)$ can be done monochromatically per each frequency ω , first reconstructing $\phi_\omega(x, z)$ and then combining them to get $\phi(x, t, z)$. Reconstruction of $\phi_\omega(x, z)$ is the same problem as the one encountered in optics when one tries to resolve features of a sample by shining monochromatic light on it. (For those who need to refresh optics, consult Appendix 6.6.) Ignoring the evanescent modes is the standard assumption in optical microscopy: the detector (playing the role of the boundary) is far from the sample, so the magnitude of the evanescent modes at the screen is exponentially small and is zero for all practical purposes. For this reason, we only have knowledge of features of the sample, $T(k_x)$, for $|k_x| \leq |\omega|$. This is the standard result

we would expect: detecting light of frequency ω , we can probe the features of a sample but only on scales larger than the resolution power set by ω^{-1} .

Specifically, we assume the boundary data specified at the $z = 0$ hypersurface tells us nothing about the coefficients $\phi_\omega(k_x)$ for $|k_x| \geq |\omega|$. Clearly, we can then not hope to confidently reconstruct $\phi_\omega(x, z)$, as the coefficient $\phi_\omega(k_x)$ for the large k_x could be arbitrarily large. To ameliorate this, we may ask for reconstructing the field $\phi_\omega(x, z)$ coarse-grained over x with the Gaussian window function of resolution scale σ . We denote this coarse-grained field as $\phi_\omega^\sigma(x, z)$:

$$\phi_\omega^\sigma(x, z) = \int_{-\infty}^{\infty} dx' e^{-(x-x')^2/\sigma^2} \phi_\omega(x', z). \quad (6.12)$$

We would expect that reconstructing $\phi_\omega^\sigma(x, z)$ should only require knowledge of modes with $|k_x| \lesssim |\sigma|^{-1}$. To verify this, we rewrite (6.12) as

$$\phi_\omega^\sigma(x, z) = \int dk_x \phi_\omega(k_x, z) e^{-k_x^2 \sigma^2} e^{-ik_x x}. \quad (6.13)$$

Unless $\phi_\omega(k_x, z)$ grows exponentially in k_x^2 for large k_x , its value is irrelevant for $|k_x| \gtrsim |\sigma|^{-1}$. Therefore, if we are only interested in reconstructing the field $\phi_\omega(x, z)$ coarse-grained over x with resolution scale σ , then we can reconstruct it using only the propagating mode data at the $z = 0$ boundary as long as $|\omega| \gtrsim 1/\sigma$ (and provided a reasonable assumption is made about the behavior of the $|k_x| \gg |\omega|$ modes).

All seems well, but there is a problem. To reconstruct $\phi^\sigma(x, t, z)$ for *all* time t , we need $\phi_\omega^\sigma(x, z)$ for all ω . Yet, there is no σ for which this condition will be satisfied: for any σ , there is an interval of missing ω , $|\omega| < 1/\sigma$. This means that, with the assumption we made that we ignore the evanescent mode data at $z = 0$, we can not possibly reconstruct the temporal evolution of the field at any z location, regardless of how large a coarse-graining resolution σ in x we are willing to compromise for. In short, for reconstruction relying only on propagating mode data, coarse-graining over x achieved reconstructability over z but sacrificed reconstructability over t .

6.2.4 Dealing with evanescent modes

Consider next the approach of retaining all the evanescent mode data at the boundary. If we wish to resolve the field on a scale $|\Delta x| \simeq \sigma$, we expect to require knowledge of the field profile with $k_x \lesssim \sigma^{-1}$. But then, evanescent modes with $\omega = 0$ and $k_x = \sigma^{-1}$ behave like $e^{z/\sigma}$, and we would expect we can determine $\phi^\sigma(x, t, z)$ only for $z \lesssim \sigma$. At bulk locations z larger than σ , the evanescent modes grow exponentially large compared to what we have access to on the boundary. We will show below that one would require exponential accuracy in the knowledge of the $z = 0$ boundary data to reconstruct the field at bulk regions as deep as $z \gg \sigma$. To foreshadow considerations in Sec. 6.3 for AdS space, let us mention that the criterion $z \lesssim \sigma$ will turn into $\sigma_{\text{proper}}(z) \gtrsim L_{AdS}[1 + \varepsilon(z)]$, where $\sigma_{\text{proper}}(z)$ is the proper distance at a bulk location z corresponding to the coordinate distance σ , up to a correction factor ε that depends on details of the bulk.

We now redo the computation of Section 6.2.1 but with evanescent modes taken into account. The monochromatic field $\phi_\omega(x, z)$ is expressed in terms of $\phi_\omega(x, 0)$ through

K_ω of (6.9). In the expression (6.9) for K_ω , we must integrate over all $|k_x| < \infty$. We split this integral into contributions of the propagating modes and the evanescent modes, respectively,

$$K_\omega(x, z|x') = \int_{-\omega}^{\omega} dk_x e^{-ik_x(x-x')} e^{-i\sqrt{\omega^2-k_x^2} z} + \int_{|k_x|>\omega} dk_x e^{-ik_x(x-x')} e^{\sqrt{k_x^2-\omega^2} z}. \quad (6.14)$$

The second integral is badly divergent. Thus, the evanescent modes have obstructed our ability to reconstruct the field in the bulk. As before, we instead ask for reconstructing the more physical quantity: the field coarse-grained over x with a Gaussian window function. We have

$$\phi_\omega^\sigma(x, z) = \int_b dx' K_\omega^\sigma(x, z|x') \phi_\omega(x', 0) \quad (6.15)$$

where

$$K_\omega^\sigma(x, z|x') = \int_b d\bar{x} e^{-(x-\bar{x})^2/\sigma^2} K_\omega(\bar{x}, z|x'). \quad (6.16)$$

Inserting (6.9) into (6.16) yields

$$K_\omega^\sigma(x, z|x') = \int dk_x e^{-ik_x(x-x')} e^{-i\sqrt{\omega^2-k_x^2} z} e^{-k_x^2\sigma^2}. \quad (6.17)$$

Asking for the field smeared over σ in the x direction thus amounts to coarse-graining the smearing function (6.14) over momentum k_x with the Gaussian window function $\exp(-k_x^2\sigma^2)$ to suppress large k_x . The $k_x \gg \omega$ part of the integral in (6.17) gives

$$[e^{-2i(x-x')z/\sigma^2} \cdot e^{z^2/\sigma^2}] e^{-(x-x')^2/\sigma^2}. \quad (6.18)$$

Having smeared out in the x -direction, let's examine the behavior in the bulk z -direction. As measured at deeper bulk regions $z \gg \sigma$, this function oscillates rapidly in phase and grows exponentially in amplitude. So, to determine the σ -grained field at a z greater than σ , one needs to measure the boundary data $\phi_\omega(x, z=0)$ with exponential precision. In short, for reconstruction retaining evanescent mode data, coarse-graining over x ameliorated reconstructability over z , for a depth of order the coarse-graining scale or so.

6.2.5 Window function

In both approaches to dealing with evanescent modes, bulk reconstruction requires a choice of a window function for the x -direction to regularize the divergence in (6.14). In (6.17), we chose a Gaussian window function to achieve the regularization. What about other choices? One might try a hard-wall window function, $\Theta(\sigma^{-1}-|k_x|)$, which gives perfect regularization, and whose x -space window function takes the form

$$\frac{\sin\left(\frac{x-x'}{\sigma}\right)}{x-x'}. \quad (6.19)$$

One might also try a Laplace window function, $e^{-k_x\sigma}$, but this function would be insufficient to regulate the divergence. Details of reconstruction certainly depend on the choice of the window function, but the fact we are limited by resolution bounds do not depend on the choice.

The choice of the window function is also a practical matter for the Scanning Tunneling Microscopy (STM). The basis of STM is the measurement of evanescent modes. In the problem of resolving the features of a sample (Appendix 6.6), the location of the sample is held fixed and bringing the STM probe needle a distance σ close to the sample allows image resolution on a scale σ . From the optics perspective, determining the spatial features of a sample regardless of the frequency with which it is illuminated, as an STM allows, is an enormous achievement. We have shown that the ability to do this in the STM context translates into our ability to reconstruct the bulk field $\phi^\sigma(x, t, z)$ for $z < \sigma$ from boundary data. While the precise depth to which reconstruction is possible would certainly depend on the chosen window function, the fact that the STM probes to a depth set by the image resolution scale, is independent of the choice.

One should note that the coarse-grained smearing function, for any coarse graining other than the hard wall choice (6.19), makes little distinction between σ less than or greater than $1/\omega$. This appears in tension with the standard assumption (reviewed in Sec. 6.2.3) that for $\sigma > 1/\omega$, reconstruction should still be possible even while ignoring the evanescent modes. However, regardless of ω , K_ω^σ has the same behavior coming from $k_x \gg \omega$ and leading to the same complications with reconstruction for $z \gtrsim \sigma$. The resolution is that the smearing function makes no assumption regarding the high k_x behavior of the field we aim to reconstruct, whereas our previous argument that modes with $k_x > \sigma$ are not needed for the σ -grained field relied on a (very reasonable) assumption about the high k_x behavior of the field.

6.2.6 Conclusions

Let's summarize what we have learned so far. Our question of interest has been whether we can reconstruct the field $\phi(x, t, z)$ using the data $\phi(x, t, z = 0)$ at the $z = 0$ timelike hypersurface. We assumed we can do this reconstruction monochromatically, reconstructing $\phi_\omega(x, z)$ from $\phi_\omega(x, 0)$. If the medium is homogeneous and only propagating modes are present, then the reconstruction works flawlessly and is given by (6.5, 6.6). If, however, the medium is inhomogeneous and evanescent modes are present, then reconstructing $\phi(x, t, z)$ point by point in the bulk is not achievable. We could instead reconstruct $\phi_\omega(x, z)$ coarse-grained in x with a σ -sized resolution, $\phi_\omega^\sigma(x, z)$. This reconstruction can be done, while not measuring any of the evanescent modes, but only for $\sigma \gtrsim 1/\omega$. Since reconstructing $\phi^\sigma(x, t, z)$ for all time t requires reconstructing $\phi_\omega^\sigma(x, z)$ for all ω , we are unable to reconstruct $\phi^\sigma(x, t, z)$ for any z . If, on the other hand, we do include the evanescent modes, *the reconstruction of $\phi^\sigma(x, t, z)$ can be done but only for the bulk depth $z \lesssim \sigma$* . This is because reconstructing for $z \gtrsim \sigma$ would require exponential accuracy in measuring the value of $\phi(x, t, 0)$.

6.3 Bulk reconstruction in AdS space

We now turn to AdS space and repeat the bulk reconstruction analysis. In short, we will find that the conclusion is exactly the same as in the flat Minkowski space case, except that we now need to understand the resolving power in proper distances. In Sec. 6.3.1

we make some general remarks regarding the relation between restricting the boundary data to local CFT operators and reconstructing a bulk AdS field from this data. In Sec. 6.3.2 we focus on reconstruction in the near-boundary region of the bulk. In Sec. 6.3.3, we consider evolving deeper into the bulk and find how the resolving power is modified.

6.3.1 Local CFT operators as boundary data

The AdS/CFT dictionary in extrapolate form [6] relates the boundary limit of a bulk operator $\hat{\phi}$ to a local CFT operator O :

$$\lim_{z \rightarrow 0} z^{-\Delta} \hat{\phi}(x, t, z) = O(x, t). \quad (6.20)$$

Hereafter, the only boundary CFT data we will consider is that of local CFT operators. We also hold the CFT Hamiltonian fixed, corresponding to the restriction that all non-normalizable modes of the bulk field ϕ are turned off. So, different states of the bulk field ϕ correspond to exciting different normalizable modes. The CFT operator $O(x, t)$ dual to the bulk field ϕ has a scaling dimension Δ set by the mass of the bulk field ϕ , $\Delta(\Delta - d) = m^2$. In non-vacuum states, the CFT operators acquire nonzero expectation values.

Defining the boundary tail through $\lim_{z \rightarrow 0} \phi(x, t, z) = \phi_0(x, t) z^\Delta$, one would like to express the field $\phi(x, t, z)$ in the bulk in terms of the boundary tail $\phi_0(x', t')$. This question is a nonstandard boundary-value problem of evolving boundary hyperbolic data at a timelike hypersurface into the bulk along a spacelike ‘radial’ direction. This is precisely the sort of problem we addressed in section 6.2 in the simpler context of flat Minkowski space. In normal Cauchy evolution of an initial-value problem, we take a Fourier transform with respect to the spatial direction x, z of elliptic data on the $t = 0$ Cauchy surface to obtain their spectral initial values. Here, we must work with hyperbolic data on the $z = 0$ timelike hypersurface. So, as in section 2, instead of doing a Fourier transform with respect to z , we do one with respect to boundary time t and obtain spectral boundary values.

Suppose complete reconstructability of the bulk is possible. Then, the bulk operator $\hat{\phi}(x, t, z)$ at a bulk location z is a linear combination of the CFT operator $O(x, t)$, smeared over the boundary hypersurface b :

$$\hat{\phi}(x, t, z) = \int_b dx' dt' K(x, t, z | x', t') O(x', t') + \dots, \quad (6.21)$$

where the ellipses refer to nonlinear interactions in the bulk. Although (6.21) is an operator statement, as a result of (6.20), finding K is reduced to a classical field theory problem. In the limit of weakly interacting bulk dynamics (corresponding to arbitrarily large CFT central charge), the nonlinear bulk interactions denoted by the ellipses in (6.21) are negligible and the reconstruction can be done mode by mode. In what follows, we will only be interested in the leading-order term of (6.21). In the bulk, we then have a fixed background on which the non-interacting bulk field $\phi(x, t, z)$ evolves according to the AdS wave equation.

Indeed, we can construct a smearing function for pure AdS space, as was done in [90, 89]. Using perturbation theory, we can also construct a smearing function for an asymptotically AdS space connected perturbatively to pure AdS space (such as one with a planet). In cases when a smearing function exists, through (6.21), the operators O provide us with a probe of bulk locality on as short of scales as the semiclassical equations of motion

are valid. In cases when a well-defined smearing function K can be constructed, (6.21) can be used to express bulk n -point correlators of $\hat{\phi}$ in terms of boundary CFT correlators of O .

We stress that having a well-defined smearing function is a more stringent requirement than simply having an algorithm for determining the bulk field $\phi(x, t, z)$ from given conformal boundary data $\phi_0(x', t')$. For instance, one might suppose that for any particular bulk solution, even if $\phi_0(x', t')$ is extremely small, one can pick an appropriate resolution for one's measuring device so as to see it and reconstruct $\phi(x, t, z)$. However, it could be that no matter how good a resolution one picks, there always exist field configurations having a near-boundary imprint $\phi_0(x', t')$ that is below the resolution power of one's measuring apparatus. In such a case, one does not have a well-defined smearing function as its existence implies a state-independent way of reconstruction. In a sense, one has to pick *a priori* the resolution power without preconceived knowledge of which field configurations will be under consideration. Indeed, in some cases, there is no smearing function [118]. A black hole background is such a case – there are modes with $\omega \ll l$ (where l is the linear or angular momentum measured in units of the AdS-scale) whose relative boundary imprint e^{-l} is arbitrarily small.

A remark is in order regarding our working assumption on the boundary data. Nothing we have said so far regarding reconstruction has made use of there actually existing a dual CFT, except that we assumed from the outset that the boundary conformal data is spanned by the set of local CFT operators. We can consider a collection of near-boundary observers who reconstruct the bulk field, their limitation of being confined to $z = 0$ is overcome by making measurements over an extended period of time. This setup is the equivalent of the statement that a complete set of local CFT operators $O(x, t)$ smeared over space and time is sufficient to reconstruct the bulk. It is important to note that this is a working assumption. The CFT is equipped with many quantities that local near-boundary observers are not. For instance, one could reconstruct the bulk by measuring CFT Wilson loops [150, 129, 148] or entanglement entropy [154]. An important question is in what situations data provided by expectation values of local CFT operators $O(x, t)$ is not sufficient to reconstruct the bulk and these other CFT quantities must be invoked. In situations when local operators $O(x, t)$ are insufficient to exactly probe bulk locality, we would like to understand if the $O(x, t)$'s are at least sufficient to probe the bulk fields coarse-grained over some scale.

6.3.2 Bulk reconstruction near the boundary

Our goal is to study bulk reconstructability for general bulk spacetimes that asymptote to pure AdS space near the timelike boundary. Let's first consider the case of reconstructing the bulk field at locations close to the boundary. There, the metric is approximated by that of pure AdS space, so we should be able to make universal statements concerning the reconstruction. The mode solutions on a general asymptotic AdS space can be classified into propagating and evanescent, depending on whether the bulk radial momentum-like quantum number is real or imaginary. If the bulk supports only propagating modes such as in pure AdS space, one can show that $\phi(x, t, z)$ can be reconstructed exactly. On the other hand, if somewhere in the bulk there is a significant change in the

geometry, then evanescent modes may become admissible.

In this section, in situations where evanescent modes are present, we will show that: (1) if we completely ignore boundary data from the evanescent modes, the bulk can not be reconstructed anywhere or on any scale, and (2) if we include boundary data from the evanescent modes, but not to exponential precision, *we are able to reconstruct the bulk, but only on AdS-size scales and not shorter; viz. we can reconstruct $\phi^\sigma(x, t, z)$ for $\sigma_{proper} > L_{AdS}$.*

AdS smearing function

Let's first work out the explicit functional form of the smearing function. Consider the massive scalar wave equation

$$\frac{1}{\sqrt{g}} \partial_\mu (\sqrt{g} g^{\mu\nu} \partial_\nu \phi) - m^2 \phi = 0, \quad (6.22)$$

in the Poincare patch of $(d+1)$ -dimensional AdS space³:

$$ds^2 = \frac{-dt^2 + d\mathbf{x}^2 + dz^2}{z^2}. \quad (6.23)$$

The modes are

$$\phi_{\omega, k_x} = 2^\nu \Gamma(\nu + 1) q^{-\nu} z^{d/2} e^{-ik_x x - i\omega t} J_\nu(qz), \quad (6.24)$$

where $q = \sqrt{\omega^2 - k_x^2}$ is like a bulk radial momentum and $\nu^2 = m^2 + (d/2)^2$. In (6.24), we normalized the modes so that they approach plane waves of unit amplitude near the boundary:

$$z^{-\Delta} \phi_{\omega, k_x} \rightarrow e^{-ik_x x - i\omega t} \quad \text{as } z \rightarrow 0. \quad (6.25)$$

We would like to express the bulk field $\phi(x, t, z)$ in terms of the conformal boundary data $\phi_0(x, t)$ at the $z = 0$ hypersurface b ,

$$\phi(x, t, z) = \int_b dx' dt' K(x, t, z|x', t') \phi_0(x', t'). \quad (6.26)$$

This equation is essentially identical to the one used in the context of reconstruction in Minkowski space (6.5). The only difference is that since in AdS space all modes universally die off near the boundary, on the right of (6.26), we used the conformal data of the field $\phi_0(x', t') \equiv \lim_{z \rightarrow 0} z^{-\Delta} \phi(x', t', z)$. We construct the smearing function K in the same way as in Sec. 6.2.1: a Fourier transform of the hyperbolic boundary b allows us to extract from $\phi_0(x, t)$ the boundary amplitudes of all modes. This yields the smearing function $K(x, t, z|x', t')$ given by

$$K(x, t, z|x', t') = \int_{-\infty}^{\infty} d\omega K_\omega(x, z|x') e^{-i\omega(t-t')}, \quad (6.27)$$

where

$$K_\omega(x, z|x') = 2^\nu \Gamma(\nu + 1) \int_{-\infty}^{\infty} dk_x e^{-ik_x(x-x')} z^{d/2} q^{-\nu} J_\nu(qz). \quad (6.28)$$

We reemphasize that (6.28) is the counterpart of (6.9), except that AdS modes have different behavior in the radial z -direction (6.24) as compared to Minkowski modes. The smearing function (6.28) accounts for this difference.

³Hereafter, we shall suppress the AdS scale L_{AdS} , and reinstate it in some of the final expressions.




Figure 6.2: The bulk profile of the modes is found from solving the Schrödinger equation in the above effective potential $V(z)$. Modes with $\omega^2 > k^2$ are propagating, whereas those with $\omega^2 < k^2$ are evanescent. In pure AdS space, evanescent modes are forbidden due to their exponential divergence at large z . However, evanescent modes become allowed if there is a change in the effective potential in the large z interior so as to cause the potential to dip below k^2 .

Propagating modes

Let's analyze in what circumstances the smearing function is well-defined. We will find that, as in the Minkowski space analysis of section 2, the AdS smearing function (6.28) is well-defined if there are only propagating modes. From (6.22), we see that the z -component of the AdS modes, $\phi_{\omega k_x} = u(z) z^{\frac{d-1}{2}} e^{-ik_x x - i\omega t}$, satisfies

$$-\frac{d^2 u(z)}{dz^2} + V(z)u(z) = \omega^2 u(z), \quad (6.29)$$

where

$$V(z) = k_x^2 + \frac{\nu^2 - 1/4}{z^2}. \quad (6.30)$$

The effective potential is plotted in Fig. 6.2. Toward the $z = 0$ boundary, the potential is confining. Toward $z \rightarrow \infty$, the potential becomes flat and we recover the behavior of Minkowski modes. Indeed, in (6.24), the Bessel function asymptotes for $qz \gg \nu$ to a plane wave:

$$J_\nu(qz) \sim \frac{1}{\sqrt{qz}} e^{iqz}. \quad (6.31)$$

It now remains to understand when the radial z -component of the momentum, q , is real-valued. In pure AdS space, the effective potential is bounded from below $V(z) \geq k_x^2$, so the frequency is bounded by $\omega^2 \geq k_x^2$. This translates into the statement that $q^2 = \omega^2 - k_x^2 \geq 0$. We find that, as for homogeneous Minkowski space, pure AdS space has only propagating modes. If the bulk space deviates from the AdS space, so long as the effective potential $V(z) > k_x^2$ for all z , there can only exist propagating modes. With propagating modes only, the integrals that define the smearing function (6.28) and (6.10) are convergent and the reconstruction is perfect. We conclude the the simple criteria for exact reconstruction of the bulk near the boundary is that there are only propagating modes.

At this point, we point out the need to exercise caution. While the reconstruction for the near-boundary region appears independent of the geometry deep in the bulk, this is actually incorrect. For instance, as we will see below, the presence of a black hole deep in the bulk drastically modifies our ability to reconstruct the field anywhere in the bulk. This is a consequence of the black hole changing the boundary conditions deep in the bulk and permitting evanescent modes to appear.

Evanescent modes

Let us now assume that there is a change in the potential deep inside the bulk such that the effective potential dips to $V(z) < k_x^2$ and evanescent modes can occur. For instance, if there is a black brane in the bulk of the Poincaré patch AdS space, then the effective potential $V(z)$ will vanish at the horizon and a continuum of evanescent modes will be present just outside the black brane horizon. We will discuss in Sec. 6.3.3 if there are other geometries for which evanescent modes can occur. For now, we simply assume that evanescent modes are present and study their impact on the reconstruction.

Choosing $k_x > \omega$ leads the z -component of the modes (6.24) to become

$$J_\nu(qz) = J_\nu(i\kappa_z z) = e^{i\nu\pi/2} I_\nu(\kappa_z z), \quad (6.32)$$

where $\kappa_z \equiv \sqrt{k_x^2 - \omega^2}$. For large $\kappa_z z$,

$$I_\nu(\kappa_z z) \sim e^{\kappa_z z}, \quad (6.33)$$

and so this solution is precisely an evanescent mode. It is clear that the monochromatic smearing function (6.28) is divergent if evanescent modes are included, and we can no longer exactly reconstruct the field. This situation is completely analogous to what happened in the context of Minkowski space in Sec. 6.2.3.

The reader may find it confusing that we can not reconstruct a local field even close to the boundary, where the bulk asymptotes to pure AdS space. Indeed, the AdS/CFT dictionary (6.20) tells us that knowledge of the local CFT operator $O(x, t)$ should give the local bulk field $\phi(x, t, z)$ at the boundary $z \rightarrow 0$. If we were to regulate this expression at $z = \epsilon$ so as to say that $O(x, t) = \epsilon^{-\Delta} \phi(x, t, \epsilon)$, then we would appear to have a contradiction. The resolution is that the proximity to the boundary, in the sense of the z -location at which modes encounter the confining barrier of the AdS space, depends on the momentum of the mode. We see from (6.31) that the mode at $z \gg \nu q^{-1}$ behaves as if it is in flat space, while the mode at $z \lesssim \nu q^{-1}$ experiences the AdS confining barrier and decays as z^Δ . It is the latter region for which (6.20) applies. In other words, the ϵ for which (6.20) starts to become applicable is momentum-dependent.

AdS locality

If the evanescent modes are present, the spectral integral for the smearing function (6.28) is divergent and we can not reconstruct $\phi(x, t, z)$. The best we can hope for is to reconstruct a coarse-grained bulk field. As in the Minkowski case, we would like to

reconstruct the field coarse-grained in the x -direction with resolution scale σ (6.12). Thus, we should compute the σ -grained smearing function

$$K_\omega^\sigma(x, z|x') = \int_b d\bar{x} e^{-(x-\bar{x})^2/\sigma^2} K_\omega(\bar{x}, z|x'), \quad (6.34)$$

where K_ω appearing in the integrand is given by (6.28). Hence,

$$K_\omega^\sigma(x, z|x') = 2^\nu \Gamma(\nu + 1) \int_{-\infty}^{\infty} dk_x e^{-ik_x(x-x')} z^{d/2} q^{-\nu} J_\nu(\sqrt{\omega^2 - k_x^2} z) e^{-k_x^2 \sigma^2}. \quad (6.35)$$

In the evanescent regime ($k_x > \omega$) and at bulk radial position $z \gtrsim \nu/\sqrt{k_x^2 - \omega^2}$, the Bessel function undergoes exponential growth (6.33). Since ν is assumed to be of order 1, evaluating (6.35) leads to the same behavior as in the analogous expression for the Minkowski case (6.18). Thus, we conclude that we can only reconstruct the σ -grained field in the bulk to the depth

$$\sigma \gtrsim z. \quad (6.36)$$

Going to a bulk coordinate distance z deeper than σ would require exponential precision in the measurements of $\phi_0(x, t)$ at the timelike boundary hypersurface ($z = 0$).

So far, the conclusions are qualitatively the same as for the Minkowski space problem. In AdS space, we need to go one step further and express the resolution bounds in proper distances. The distance σ in which we have coarse-grained in the x -direction is a coordinate distance, natural from the viewpoint of boundary data of the dual CFT. The proper distance between two points (z_1, x_1, t_1) and (z_2, x_2, t_2) in AdS space is given by

$$(\Delta s)^2 = \frac{1}{z_1 z_2} [(z_1 - z_2)^2 + (\mathbf{x}_1 - \mathbf{x}_2)^2 - (t_1 - t_2)^2]. \quad (6.37)$$

So, converting the coordinate distance resolution $\Delta x = \sigma$ to the proper one, we have the proper coarse-graining scale at bulk location z

$$\sigma_{\text{proper}}(z) = \sigma \frac{L_{\text{AdS}}}{z}. \quad (6.38)$$

Thus, (6.36) translates at radial depth z to a proper resolution bound:

$$\sigma_{\text{proper}}(z) \gtrsim L_{\text{AdS}}. \quad (6.39)$$

This is one of our main results: with our assumptions, *in an asymptotically AdS space which gives rise to evanescent modes, we can not reconstruct the bulk field any better than the AdS scale.* Moreover, the right-hand side of (6.39) being independent of z , we are able to reconstruct the AdS bulk without limits on the depth, at least for the near-boundary region.⁴

We derived the proper resolution bound (6.39) from analysis of the near-boundary region of the bulk. For regions deeper in, the bound may be further modified. For instance, for AdS black holes, the bound could also depend on the black hole horizon scale R_{BH} as

⁴The result (6.39) was anticipated in [31] in the context of AdS-Rindler.

well as the bulk depth z . On general grounds, we expect that the resolution bound takes the form

$$\sigma_{\text{proper}}(z) \gtrsim L_{\text{AdS}} [1 + \varepsilon(z, R_{\text{BH}})], \quad (6.40)$$

where ε is background-specific correction factor at the bulk depth z . In the next section, we shall extract this correction for a large AdS black hole.

Once again, we have a clear analogy with Scanning Tunneling Microscopy (STM). The invention of STM revolutionized microscopy: by placing a needle at a distance σ from a sample, some of the evanescent modes could be captured, allowing resolution on a scale σ . Remarkably, in AdS space, the resolution in the bulk is set by the AdS scale and the reconstruction can be done to arbitrary bulk depth.

6.3.3 Bulk reconstruction deeper in

In this section we evolve deeper into the bulk. Specifically, we take an AdS-Schwarzschild background and classify the types of modes present, and identify criteria for the presence of evanescent modes in other backgrounds. We also find the correction ε in (6.40) that the resolution bound (6.39) receives at locations deeper in the bulk.

Effective potential for modes

Consider a scalar field ϕ in a general spherically symmetric background that asymptotes to AdS space,

$$ds^2 = -f(r)dt^2 + \frac{dr^2}{f(r)} + r^2 d\Omega_{d-1}^2. \quad (6.41)$$

The field ϕ satisfies the wave equation (6.22). Separating ϕ as

$$\phi(r, t, \Omega) = \varphi(r)Y(\Omega)e^{-i\omega t} \quad (6.42)$$

gives for the radial field $\varphi(r)$,

$$\frac{\omega^2}{f}\varphi + \frac{1}{r^{d-1}}\partial_r(f r^{d-1}\partial_r\varphi) - \frac{l(l+d-2)}{r^2}\varphi - m^2\varphi = 0. \quad (6.43)$$

Letting $\varphi(r) = u(r)/r^{\frac{d-1}{2}}$ and changing variables to the tortoise coordinate $dr_* = f^{-1}dr$ turns (6.43) into a Schrödinger-like equation

$$\frac{d^2u}{dr_*^2} + (\omega^2 - V(r))u = 0, \quad (6.44)$$

with the effective potential

$$V(r) = f \left[\frac{(d-1)}{2} \frac{f'}{r} + \frac{(d-1)(d-3)}{4} \frac{f}{r^2} + \frac{l(l+d-2)}{r^2} + m^2 \right]. \quad (6.45)$$




Figure 6.3: (a) The effective potential (6.45) for a small AdS black hole. The boundary of AdS space is at tortoise coordinate $r_* = \pi/2$, while the horizon is at $r_* = -\infty$. Evanescent modes arise if the potential ever drops below the value at the local minimum present in the pure AdS space ($\sim l^2$). Here, this occurs because the potential approaches zero at the horizon. (b) The effective potential for pure global AdS space.

Small AdS-Schwarzschild Black Hole

The AdS- Schwarzschild black hole metric is given by:

$$ds^2 = -\left(1 + r^2 - \left(\frac{r_0}{r}\right)^{d-2}\right) dt^2 + \frac{dr^2}{1 + r^2 - \left(\frac{r_0}{r}\right)^{d-2}} + r^2 d\Omega_{d-1}^2. \quad (6.46)$$

Here, we will be interested in a small black hole, so $r_0 \ll L_{AdS}$ and $r_0 \approx 2M$. In Fig. 6.3(a), we plotted the effective potential (6.45) for a small black hole. There are four kinds of modes, and can be classified depending on the ratio of ω to l . We characterize the modes going from high ω to low ω .

1. $\omega \gtrsim \frac{l}{r_0}$. These modes are higher than the angular momentum barrier and propagate into the black hole; they correspond to throwing ϕ particles into the black hole.
2. $l \lesssim \omega \lesssim \frac{l}{r_0}$. These modes are directly related to the modes in the pure AdS space (plotted in Fig. 6.3 (b)). They correspond to the particles having sufficient angular momentum so that they stay far away from the black hole and do not notice its presence. They do differ from the pure AdS modes in that they have exponentially suppressed tails which are propagating near the black hole.
3. $l \lesssim \omega \lesssim \frac{l}{r_0}$. This is the same regime as type 2 modes, but these modes are the ones that have most of their support close to the black hole and only an exponentially small amplitude in the asymptotic region. We will call these modes trapped modes.
4. $\omega \lesssim l$. These are the evanescent modes. All possible evanescent modes, with any ω for any value of l , are present as a result of the potential dropping to zero at the horizon.

As discussed in Sec. 6.3.2, it is only the evanescent modes which inhibit reconstruction of the region near the boundary. The trapped modes inhibit reconstruction of the field,

but only for regions close to the black hole ($r < 3r_0/2$). Recall that we normalize all the modes so that their boundary limit is $\phi_{\omega l} r^\Delta \rightarrow 1$. The trapped modes are propagating in most of the AdS bulk. Only when they encounter $r \sim 3r_0/2$, and have to pass under the centrifugal barrier, do these modes begin to undergo exponential growth. Neither modes of type (1) or (2) are problematic for reconstruction as they never undergo exponential growth as compared to their boundary value.

We should emphasize that the difficulties with reconstruction in a black hole background are not due to the presence of the horizon, per se. One might think there should be difficulties with reconstruction because there are now things which fall into the black hole and so can't be seen from the boundary (in other words, that modes of type (1) cause difficulties). However, the case of pure Poincaré Patch is a counter-example. There, everything falls into the Poincaré horizon, yet as we showed in Sec. 6.3.2, there are no difficulties with bulk reconstruction. The only necessary criteria the modes need to satisfy for the bulk reconstruction is that their magnitude at the bulk point of interest is not exponentially larger than their boundary value.

It is interesting to ask in which backgrounds, aside from the AdS black hole, are the trapped modes (type 3) present. For a static, spherically symmetric spacetime,

$$ds^2 = -f(r)dt^2 + \frac{dr^2}{g(r)} + r^2 d\Omega^2, \quad (6.47)$$

one could write down an effective potential (similar to (6.45)) and see if it has a barrier (a local maximum). For the question of reconstruction, we are most interested in modes with large l : if there are trapped/evanescent modes it is these that will be hardest to reconstruct. So, we can simplify the analysis by taking the large l limit of the effective potential. For any given r , in the large l limit, the effective potential simplifies to

$$V = f \frac{l^2}{r^2}, \quad (6.48)$$

and the criterion for having trapped modes is that

$$\frac{d}{dr} \left(\frac{f}{r^2} \right) > 0 \quad (6.49)$$

for some value of r . In [118] it was shown that (6.49) implies that there is no smearing function for some region in the bulk. The condition (6.49) can be achieved by having a sufficiently dense planet so that its radius is less than $3M$.

It is also interesting to ask in which backgrounds, aside from the AdS black hole, are the evanescent modes (type 4) present. The criterion for having evanescent modes is that the effective potential has a new global minimum, different from the one present in pure AdS space which has the value $\sim l^2$. From (6.48), this means that

$$\frac{f}{r^2} < 1 \quad (6.50)$$

for some value of r . Outside the planet, $r > R$, the metric takes the form (6.46). Letting r_h denote the would-be horizon of the planet, (6.50) translates into the following condition on




Figure 6.4: The effective potential for a large black as a function of the tortoise coordinate. Note that with tortoise coordinates the region near the boundary (large r) gets compressed near $\pi/2$. Thus, the top corner of the plot is the only portion of the potential that is similar to a portion of the pure AdS potential (Fig. 6.3(b)). Here, the modes are mostly those of type (1) that propagate into the black hole horizon, and type (4) that have low energy and hence are evanescent.

the radius R of the planet:

$$R - r_h < \frac{r_h}{d-2} \left(\frac{r_h}{L_{AdS}} \right)^2. \quad (6.51)$$

Since $r_h/L_{AdS} \ll 1$ by assumption, it is nontrivial for physical matter to be so dense. In particular, if one assumes the density is nonnegative and a monotone decreasing function, then a general argument [180] gives that one must have $R - r_h > r_h/8$ (for 4 spacetime dimensions).

Large AdS-Schwarzschild black hole

We would like to find corrections to the reconstruction resolution bound $\sigma_{\text{proper}} > L_{AdS}$ for locations deeper in the bulk, where the geometry is no longer pure AdS space. To find this, we need to solve for the radial profile of the evanescent modes. The evanescent modes will grow exponentially as one moves from the boundary into the bulk. The distance into the bulk that we can evolve is set by the rate of the growth: when the coefficient in the exponential becomes of order one, we are unable to evolve deeper in.

We focus on a large AdS black hole, which is similar to a black brane. In the limit that the black hole is large, $r_0 \gg 1$, the metric (6.46) can be simplified by dropping the 1 as it is small relative to r^2 , so that $f(r) = r^2 - (r_0/r)^{d-2}$. If in addition to this we also zoom in on a small portion of the solid angle, this yields the black brane metric,

$$ds^2 = - \left(r^2 - \left(\frac{r_0}{r} \right)^{d-2} \right) dt^2 + \frac{dr^2}{r^2 - \left(\frac{r_0}{r} \right)^{d-2}} + r^2 dx^2. \quad (6.52)$$

The horizon location, r_h , and temperature, T , are given by

$$r_h^d = r_0^{d-2}, \quad T = \frac{d r_h}{4\pi}. \quad (6.53)$$

The equation for the effective potential satisfied by the radial modes is given by (6.45) with $l(l+d-2)$ replaced by k_x^2 ,

$$V = f \left[\frac{d^2 - 1}{4} + m^2 + \frac{(d-1)^2}{4} \frac{r_0^{d-2}}{r^d} + \frac{k_x^2}{r^2} \right]. \quad (6.54)$$

To solve for the radial profile of the modes, we use the WKB approximation, only keeping the exponential factor

$$u(r) = \exp \left(\int_r^\infty \frac{dr'}{f(r')} \sqrt{V(r') - \omega^2} \right). \quad (6.55)$$

We are interested in the mode which is the most evanescent, and this is the one with $\omega = 0$. Thus, for (6.55), we have

$$u(r) = \exp \left(\int_r^\infty \frac{dr'}{\sqrt{r'^2 - r_h^d/r'^{d-2}}} \sqrt{\bar{v}^2 + \frac{k_x^2}{r'^2} + \frac{(d-1)^2}{4} \left(\frac{r_h}{r'} \right)^d} \right), \quad (6.56)$$

where we defined $\bar{v}^2 = m^2 + (d^2 - 1)/4$. Since $r_h/r < 1$ and \bar{v} is of order 1, we can drop the last term in (6.56). Converting to the z -coordinate, $z = 1/r$, and setting $\sigma \equiv k^{-1}$, (6.56) becomes

$$u(z) = \exp \left(\int_0^z \frac{dz'}{\sqrt{1 - (z'/z_h)^d}} \sqrt{\frac{1}{\sigma^2} + \frac{\bar{v}^2}{z'^2}} \right). \quad (6.57)$$

In the case of pure AdS space, we have (6.57) but with $z_h = \infty$. In that case, we argued that for small z (such that $z \ll \sigma$), we can drop the $1/\sigma^2$ term in the square root and find that $u(z)$ has power law behavior consistent with the conformal scaling, z^Δ , of the field near the boundary. For $z \gtrsim \sigma$, we instead drop the \bar{v}^2/z^2 term in the square root, leading to the behavior

$$u(z) = \exp \left(\frac{z}{\sigma} \right). \quad (6.58)$$

With an AdS black hole present, this behavior receives only minor corrections. Since

$$\frac{1}{\sqrt{1 - (z/z_h)^d}} > 1, \quad (6.59)$$

$u(z)$ has a faster growth in an AdS black hole background than in pure AdS space. We can find the field at $z = z_h$ and in the limit $\sigma \ll z_h$. This gives⁵,

$$u(z_h) = \exp \left(\frac{1}{\sigma} \int_\sigma^{z_h} \frac{dz}{\sqrt{1 - (z/z_h)^d}} \right). \quad (6.60)$$

Evaluating the integral gives

$$u(z_h) = \exp \left(\alpha \frac{z_h}{\sigma} \right), \quad (6.61)$$

⁵Here, we set the lower limit of integration to be σ since this is around where the approximation of dropping the \bar{v}^2/z^2 term becomes invalid. For $z \lesssim \sigma$, one should instead drop the $1/\sigma^2$ term. However, since by assumption $\sigma/z_h \ll 1$, one can just as well set the lower limit of integration to 0.

where

$$\alpha = \frac{\sqrt{\pi} \Gamma[1 + 1/d]}{\Gamma[2 + 1/d]} . \quad (6.62)$$

The value of α is larger than 1: for instance, in $d = 4$, $\alpha \approx 1.3$, only slightly larger than 1. This fits with our expectation, since if we had been in pure AdS space, then we would have had $u(z) = \exp(z/\sigma)$. The exponent α is universal in that it depends only on d but not on other details of the AdS black hole; it can be considered a nontrivial prediction for a strongly coupled CFT at finite temperature.

Comparing with (6.40), we now obtain the correction factor ε to the resolution bound at the horizon:

$$\varepsilon(z_h) = \alpha - 1 = \frac{\sqrt{\pi} \Gamma[1 + 1/d]}{\Gamma[2 + 1/d]} - 1 \leq 0.7\dots \quad (6.63)$$

In the next section, we will find the physical interpretation of the correction factor ε at the horizon as a nonperturbative effect in the thermal Green's functions in the CFT dual.

6.4 Evanescence in CFT Dual

In this section, we turn to the CFT dual of the AdS space and pose the question: how do evanescent modes manifest themselves in the CFT? The first question is whether the evanescent modes are part of the CFT spectrum. To be specific, consider the large AdS black hole we studied in the last section. The black hole is the holographic dual of the CFT at a high temperature T that equals the black hole's Hawking temperature. The $T = 0$ vacuum is now changed to the $T > 0$ ground state (this ground state is created out of the vacuum by \mathcal{O}_{BH}). Now, the $T > 0$ spectrum includes the continuous, evanescent spectrum in addition to the discrete spectrum. In effect, the continuous spectrum increases the dimension of the Hilbert space.

One might object that the evanescent states do not propagate to the AdS boundary and need not to be included. However, they do contribute to CFT correlators at finite temperature and should be included. For instance, consider the 4-point CFT correlator G_4 , as computed holographically in a large AdS black hole background:

$$\begin{aligned} G_4(1, 2, 3, 4) &\equiv \left\langle T \left| \mathcal{O}_1 \mathcal{O}_2 \mathcal{O}_3 \mathcal{O}_4 \right| T \right\rangle_{\text{CFT}} \\ &= \int \int_B dx dy G_{Bb}(1, x) G_{Bb}(2, x) K_2(x, y) G_{Bb}(3, y) G_{Bb}(4, y) . \end{aligned} \quad (6.64)$$

The point is that the integrals over the bulk x and y in the region close to the black hole horizon contain contributions from evanescent modes. This is because the evanescent modes are genuine propagating modes in this region and the bulk-to-bulk propagator $K_2(x, y)$ contains these modes. More generally, propagating modes from or to the AdS boundary would couple to the evanescent modes whose wave function is localized near the black hole horizon. This implies that, in the dual CFT, we should expect a new class of 3-point correlators that involve two local CFT operators $\mathcal{O}_1, \mathcal{O}_2$ and one effective operator E_T associated with the heat bath:

$$\langle \mathcal{O}_1 \mathcal{O}_2 E_T \rangle \sim \langle T | \mathcal{O}_1 \mathcal{O}_2 | T \rangle . \quad (6.65)$$

To see (6.65) explicitly, we compute the CFT 2-point function in the evanescent regime by computing a bulk 2-point function and taking its boundary limit. Let us canonically quantize the bulk scalar field $\phi(x, t, z)$ and expand in terms of modes $\phi_{\omega k}$,

$$\phi = \iint d\omega dk \left(\phi_{\omega k} a_{\omega k} + \phi_{\omega k}^* a_{\omega k}^\dagger \right), \quad (6.66)$$

where the creation operators satisfy the usual commutation relation. Here, we normalize the modes differently from Sec. 6.3 and use the bulk Klein-Gordon normalization, which is more natural in the present context.⁶

The bulk 2-point function can thus be written as:⁷

$$\langle \phi(x_2, t_2, z_2) \phi(x_1, t_1, z_1) \rangle = \int d\omega dk \phi_{\omega k}(x_2, t_2, z_2) \phi_{\omega k}^*(x_1, t_1, z_1) \quad (6.67)$$

Writing the modes as

$$\phi_{\omega k}(x, t, z) = f_{\omega k}(z) e^{-ikx + i\omega t}, \quad (6.68)$$

and using the AdS/CFT dictionary (6.20), we get

$$\langle T | O(x_2, t_2) O(x_1, t_1) | T \rangle = \lim_{z \rightarrow 0} \iint d\omega dk z^{-2\Delta} |f_{\omega k}(z)|^2 e^{-ik(x_2 - x_1) + i\omega(t_2 - t_1)}. \quad (6.69)$$

Let us focus on the portion of the 2-point function (6.69) coming from the evanescent regime $\omega \ll k$. Since the Klein-Gordon normalization roughly corresponds to having the mode be of order 1 at the horizon, to find the $z \rightarrow 0$ limit of $|f_{\omega k}(z)|^2$, we just need the WKB factor giving the relative suppression at the boundary. This factor was computed in (6.61). We thus find that, in the regime $\omega \ll k$ and $k \gg T$, the bulk 2-point function has the behavior (ignoring prefactors),

$$G(\omega, k) = e^{-\tilde{\alpha} \frac{k}{T}}. \quad (6.70)$$

Here, $\tilde{\alpha}$ is a constant factor given by $\tilde{\alpha} = (d/2\pi)\alpha$, where α was found in (6.62).⁸ The result (6.70) is not surprising. At zero temperature, there is no evanescent mode with $\omega < k$ and so such a correction should vanish. So, at finite temperature, the evanescent modes must generate a contribution which vanishes in the zero temperature limit. This explains the Boltzmann suppression (6.70) of the evanescent mode contribution.

In the previous sections, we argued that a black hole in the bulk changes the boundary conditions so as to permit evanescent modes and that a small error in determining their coefficient will get exponentially amplified when reconstructing the bulk. Here, we have just taken the ground state for the large black hole background and found the boundary imprint of the evanescent modes through CFT 2-point correlators at finite temperature. If

⁶In Sec. 6.3, in for instance (6.24), we normalized the modes so that when rescaled by $z^{-\Delta}$, their z -component approaches 1 at the boundary. This was a natural normalization in that context since we needed to use the boundary limit of the field to then do a Fourier transform over space and time on the boundary to extract the coefficients of the modes. Here, we normalize with respect to the bulk Klein-Gordon norm so that (6.67) takes a simple form.

⁷Note that we are taking the bulk state to be the Hartle-Hawking vacuum.

⁸In [162] the imaginary part of the retarded Green's function was found in the evanescent regime through an alternate method. Their result for $d = 4$ agrees with (6.70).

the bulk state is a slight deviation from this ground state, then to determine it near the horizon one must make a boundary measurement precise enough to detect something as small as (6.70). Of course, we can have an excited state in the black hole background that has a much larger coefficient for an evanescent mode.

6.5 Discussion

The AdS resolution bound we have found in this chapter has an important impact on the dual CFT. The first question concerns the boundary CFT data. We showed that, in an AdS black hole background, bulk fields have evanescent modes. These modes are exponentially suppressed near the boundary. Translated into the CFT description, this means that we are unable to reconstruct the bulk on sub-AdS scales if we only have local CFT data, $\langle O \rangle$, and do not have it to exponential precision. The next question is what CFT data naturally encodes the evanescent modes.

Much in the same way as the evanescent modes trapped to a material are regarded as “part” of the material, the evanescent modes in an AdS black hole background are naturally part of the semiclassical black hole. Thus, a speculative possibility is that there exists a theory which describes the near-horizon atmosphere of a black hole and, as a subsector of the CFT, can be thought of as living on the black hole horizon.⁹

An extremal Reissner-Norstrom AdS black hole is a particularly concrete setting in which to explore this, as the near horizon limit is $\text{AdS}_2 \times S^{d-1}$. From the perspective of the AdS_2 , the evanescent modes of the underlying AdS black hole are propagating modes of AdS_2 . Furthermore, $\text{AdS}_2 \times S^{d-1}$ should be dual to multiple copies of a CFT_1 . We would therefore conjecture that the bulk of the AdS extremal black hole could be reconstructed by a combination of local operator data coming from both the boundary CFT_d and the tower of horizon CFT_1 ’s.¹⁰

6.6 Evanescent Optics

In this appendix, we summarize some central principles of optics that make us evanescent modes.

Total Internal Reflection

To see how evanescent waves can be generated if homogeneity is broken, consider the wave equation in a background which has a z -dependent speed of light,

$$\left(-\frac{1}{c(z)^2} \partial_t^2 + \partial_x^2 + \partial_z^2 \right) \phi = 0. \quad (6.71)$$

⁹Note that the black hole horizon would be used as a holographic screen to describe the black hole atmosphere; we are not discussing anything about the black hole interior.

¹⁰ $\text{AdS}_2 \times S^n$ has been argued to be problematic [132, 1]. This would not be a problem for us, as we do not fully take the near horizon limit. The complete near horizon limit would correspond to keeping only the $\omega = 0$ modes, while we are interested in small but finite ω modes. Alternatively, one could consider our construction for a black p -brane for which the quoted issues are absent.




Figure 6.5: The effective potential for the wave equation in a medium that undergoes a jump in its index of refraction.

The modes are

$$\phi(x, t, z) = e^{-i\omega t - ik_x x} u_{\omega, k}(z), \quad (6.72)$$

where $u(z)$ satisfies a Schrodinger-like equation

$$-u'' + V(z)u = 0 \quad (6.73)$$

with an effective potential

$$V(z) = k_x^2 - \frac{\omega^2}{c^2(z)}. \quad (6.74)$$

As a simple scenario which generates evanescent modes, we take

$$c(z) = \begin{cases} c_1 & z > z_1 \\ c_0 & z \leq z_1. \end{cases} \quad (6.75)$$

The potential (6.74) for this choice of $c(z)$ is shown in Fig. 6.5. We see that modes with

$$\frac{\omega^2}{c_1^2} < k_x^2 < \frac{\omega^2}{c_0^2}$$

are evanescent for $z < z_1$. Indeed, the lower limit $k_x = \omega/c_1$ is precisely the well-known critical angle for total internal refraction: $\sin\theta_c = n_1$. Here θ_c is the angle between the incoming wave at $z > z_1$ and the normal in the z direction and n_1 is the index of refraction of the material, $n_1 = c_0/c_1$.

Microscopes

Another context in which evanescent modes appear is in the use of a microscope. The microscope is trying to resolve the spatial features of a sample which is thin and located at z_0 (see Fig. 6.6). The sample is projected with monochromatic light of frequency ω coming from a source at $z_1 > z_0$. The sample has some space-dependent transmission coefficient $T(x)$, which then determines the wave profile that is received at the detector at




Figure 6.6: Light is shined from the source (z_1), interacts with the sample (z_0), and is received at $z = 0$.

$z = 0$. As long as the sample has spatial features on scales smaller than ω^{-1} , evanescent wave which are trapped near the sample will be generated.

The field consists of a monochromatic wave with frequency ω , $\phi_\omega(x, z)$, which we will simply denote by $E(x, z)$. We wish to relate the field at the source to the field received at the detector, and see how the transmission coefficient enters this expression.

The field at the source is denoted by $E_{\text{source}}(x, z_1)$, and the generated wave propagates freely from z_1 to z_0 ,

$$E_{\text{source}}(q_x, z_0) = E_{\text{source}}(q_x, z_1) e^{-iq_z(z_1 - z_0)}, \quad (6.76)$$

where $q_z = \sqrt{\omega^2 - q_x^2}$. The effect of the sample is to modify the wave at $z = z_0$ by the transmission coefficient. After passing through the sample the field becomes

$$E_{\text{sample}}(x, z_0) = T(x) E_{\text{source}}(x, z_0), \quad (6.77)$$

In Fourier space (6.77) becomes the convolution

$$E_{\text{sample}}(k_x, z_0) = \int dq_x T(k_x - q_x) E_{\text{source}}(q_x, z_0). \quad (6.78)$$

From the sample the wave propagates freely to the detector at $z = 0$,

$$E_{\text{detector}}(k_x, 0) = E_{\text{sample}}(k_x, z_0) e^{-ik_z z_0}. \quad (6.79)$$

Combining (6.76), (6.78), and (6.79) yields

$$E_{\text{detector}}(x, 0) = \int dk_x dq_x E_{\text{source}}(q_x, z_1) e^{-iq_z(z_1 - z_0)} T(k_x - q_x) e^{-ik_z z_0} e^{-ik_x x}. \quad (6.80)$$

Eq. 6.80 is the relation we had been seeking between the emitted field at the source, E_{source} at $z = z_1$, and the received field at the detector, E_{detector} at $z = 0$.

From (6.80) we see that the sample serves to convert the waves impinging on it with x momentum q_x to those with momentum k_x , with the conversion amplitude given by $T(p_x)$ where $p_x \equiv k_x - q_x$. The incident waves on the sample are propagating; thus $|q_x| < \omega$.




Figure 6.7: The evanescent modes can be converted into propagating modes by placing, for example, a piece of glass near them.

For the converted waves to be propagating they need $|k_x| < \omega$. Thus, unless the sample has no features on any spatial scale smaller than ω^{-1} (so that $T(p_x)$ vanishes for all $|p_x| > \omega$), there will be evanescent modes at $z < z_0$. Indeed, to generate evanescent modes one need only choose a sample which is a reflecting metal sheet with a hole. In this case $T(x)$ has compact support, requiring $T(p_x)$ to be nonzero at arbitrarily large p_x .

Scanning Tunneling Optical Microscope

We have found that the frequency ω of the wave shined on the sample sets the limit of the scale on which the sample can be resolved: modes with $k_x > \omega$ are evanescent and too suppressed at the detector location to be measurable. For a long time this was believed to set a fundamental limit on the best achievable resolution of a material [136].¹¹

The way to resolve the structure on scales shorter than ω^{-1} would be to detect the evanescent modes. Since their magnitude is exponentially small at the location of the detector, one must instead detect them close to the sample or equivalently, convert them into propagating waves. This can be achieved by placing some object near the material so as to change the index of refraction. A sketch is shown in Fig. 6.7. The near-field detection of evanescent waves is the basis of how an STOM functions. In an STOM, a pointer is brought close to the sample and some of the evanescent waves hitting it are converted into propagating waves, which are then able to reach the detector.

Suppose we wish to resolve the sample on a scale $\sigma \ll \omega^{-1}$. This requires detecting evanescent modes with $k_x \sim \sigma^{-1}$, and corresponding $k_z = \sqrt{\omega^2 - k_x^2} \approx -ik_x$. Since the evanescent modes decay exponentially (6.11), the pointer tip must be placed no further than a distance $\Delta z \sim \sigma$ from the sample. Thus, an STOM allows resolution of features on arbitrarily short scales σ , regardless of ω . However, the new limitation is set by the distance of the pointer tip from the sample. To resolve on a scale σ requires placing the tip no more than a distance σ away from the sample so as to capture the evanescent modes.

¹¹It may appear this hindrance can be overcome by simply using waves of arbitrarily high frequency. However, sufficiently high frequency waves will damage the sample.

Chapter 7

Entanglement Entropy: A Perturbative Calculation

7.1 Introduction

Entanglement entropy is a rapidly developing technique in condensed matter physics [113, 46] and holography [154]. One of the main theoretical gaps that substantially limits its studies is the paucity of computational tools. In this chapter we construct a perturbative framework for computing entanglement entropy of the vacuum purely within the context of quantum field theory (QFT).

As of today the existing tools for computing entanglement entropy include: the replica trick, conifolds, and the elegant prescription of Ryu and Takayanagi [154]. The replica trick, and its generalizations, is the only generic approach to calculating entanglement entropy within field theory [46, 47]. It rests on evaluating the partition function on an n -folded cover of the background geometry where a cut is introduced throughout the exterior of the entangling surface. However, evaluation of the partition function on a replicated manifold can only be carried out in a limited number of cases. On the other hand, the Ryu-Takayanagi prescription is much easier to implement. It plays a central role in characterizing new properties of holographic field theories, *e.g.*, [114], and provides new insights into the quantum structure of spacetime [60]. Recently, the generalized replica trick was successfully implemented in the bulk AdS space to provide strong evidence for the Ryu-Takayanagi conjecture [119]

In [48] Casini, Huerta and Myers showed that the reduced density matrix for spherical entangling surfaces in flat space is conformally equivalent to a thermal state on the hyperbolic geometry, and that the entanglement entropy equals the thermodynamic entropy of this thermal state. This observation provided an alternative derivation of the holographic entanglement entropy for spherical regions in flat space. However, their construction tightly relies on the conformal symmetry of the boundary CFT and on the (spherical) geometry of the entangling surface. Hence, their work raises a natural question: how does one accommodate small disturbances of their framework within a perturbative approach? In this chapter we propose a Euclidean path integral formalism that addresses this question. In particular, our method paves the way for an alternative approach to calculating entanglement entropy

within quantum field theory.

In Section 7.2 we set aside holography, the replica trick, and other existing methods of calculating entanglement entropy and begin with the ‘standard’ Euclidean path integral definition of the reduced density matrix. Next, we foliate spacetime in the vicinity of the entangling surface in such a way as to encode both the geometric structure of the surface and the geometry of the background. This choice of coordinates is one of the central aspects of our approach, as any deformation can be now thought of as a background deformation. As a result, a perturbative framework around systems with known reduced density matrices is established.

In Section 7.3 we consider the entanglement entropy obtained by dividing the field theory into two (semi-infinite) regions with a single flat plane separating them. In this case the entanglement entropy for any QFT equals the thermal entropy observed by an accelerating Rindler observer [108]. We apply our general formalism to calculate leading order corrections induced by either slight curvature of the background or mild deformations of the flat wall separating the two subsystems. In particular, we evaluate the universal divergence of the entanglement entropy induced by these modifications in four dimensional spacetimes. The results are in complete agreement with the structure of the universal terms in entanglement entropy of 4D conformal field theories originally proposed by Solodukhin [161].

The main focus of Section 7.4 is the analysis of perturbations around spherical entangling surfaces. The unperturbed case in the context of QFT was studied in [48], whereas in this work we implement our formalism to investigate consequences of small perturbations. The resulting corrections to the universal divergence of entanglement entropy in 4D match known results in the literature [161].

7.2 General framework

We start with a general quantum field theory that lives on a d -dimensional Euclidean manifold \mathcal{M} equipped with a Riemannian metric $g_{\mu\nu}$. The action of the field theory is given by $I_0(\phi, g_{\mu\nu})$, where ϕ collectively denotes all the QFT fields. We assume that the system resides in the vacuum state¹. The entangling surface is chosen to be some general $(d-2)$ -dimensional surface Σ . Our notation for the rest of the chapter is summarized in Appendix 7.5.

The degree of entanglement between the QFT degrees of freedom inside and outside of Σ is encoded in the reduced density matrix ρ_0 that can be written as a path integral over \mathcal{M} with a $(d-1)$ -dimensional cut \mathcal{C} , such that $\partial\mathcal{C} = \Sigma$

$$[\rho_0]_{\phi_- \phi_+} \equiv \langle \phi_- | \rho_0 | \phi_+ \rangle = \int_{\substack{\phi(\mathcal{C}_+) = \phi_+ \\ \phi(\mathcal{C}_-) = \phi_-}} \mathcal{D}\phi e^{-I_0(\phi, g_{\mu\nu})} , \quad (7.1)$$

where \mathcal{C}_\pm are the two sides of the cut and ϕ_\pm are some fixed field configurations (see Fig. 7.1).

In general, evaluation of the above path integral is not a tractable problem, but there are exceptions, *e.g.*, planar and spherical surfaces in R^d that we are going to explore

¹For entanglement entropy of excited states in the holographic context see, [104].




Figure 7.1: Abstract sketch of the two dimensional transverse space to the entangling surface Σ . \mathcal{C}_{\pm} are the two sides of the cut \mathcal{C} where the values ϕ_{\pm} of the field ϕ are imposed.

later. For the rest of this section the details of ρ_0 are not crucial, we only need to assume that it is known, since the main purpose is to get a closed form expression for small perturbations of ρ_0 as a consequence of slight deformations of the background metric $g_{\mu\nu}$ and entangling surface Σ , or perturbations of the QFT by, *e.g.*, a relevant operator.

We start with the normalized density matrix,

$$\hat{\rho}_0 = \frac{\rho_0}{\text{Tr}\rho_0} . \quad (7.2)$$

The corresponding modular Hamiltonian, \hat{K}_0 , and the entanglement entropy, S_0 , are given by

$$\begin{aligned} \hat{K}_0 &= -\log \hat{\rho}_0 , \\ S_0 &= -\text{Tr}\hat{\rho}_0 \log \rho_0 . \end{aligned} \quad (7.3)$$

Now let us consider perturbation of $\hat{\rho}_0$ by a small amount $\delta\hat{\rho}$,

$$\hat{\rho} = \hat{\rho}_0 + \delta\hat{\rho} , \quad (7.4)$$

The new density matrix $\hat{\rho}$ is assumed to be normalized, and therefore $\text{Tr}\delta\hat{\rho} = 0$. The corresponding modular Hamiltonian, \hat{K} , and the entanglement entropy, S , can be constructed perturbatively provided that $\hat{\rho}_0$ and $\delta\hat{\rho}$ are known

$$\begin{aligned} \hat{K} &= -\log \hat{\rho} = \hat{K}_0 - \hat{\rho}_0^{-1} \delta\hat{\rho} + \frac{1}{2}(\hat{\rho}_0^{-1} \delta\hat{\rho})^2 + \mathcal{O}(\delta\hat{\rho}^3) , \\ S &= -\text{Tr}\hat{\rho} \log \hat{\rho} = S_0 + \text{Tr}(\delta\hat{\rho} \hat{K}_0) - \frac{1}{2}\text{Tr}(\delta\hat{\rho} \hat{\rho}_0^{-1} \delta\hat{\rho}) + \mathcal{O}(\delta\hat{\rho}^3) . \end{aligned} \quad (7.5)$$

In those examples that we are going to consider, it is possible (but not always necessary) to implement a conformal transformation that maps the background \mathcal{M} , and hence the path integral (7.1), onto $S^1 \times H^{d-1}$ which we will denote as \mathcal{H} . Of course, we implicitly restrict our consideration here to CFTs. Remarkably, under this transformation the entangling surface Σ is mapped onto the conformal boundary of H^{d-1} while fixed states

$|\phi_{\pm}\rangle$ are mapped onto constant slices $\tau_E = 0$ and $\tau_E = \beta$ (see Fig 7.4). The latter condition ensures that under this map the reduced density matrix $\hat{\rho}_0$ transforms into a normalized thermal density matrix $\hat{\rho}_T$ on \mathcal{H} . In particular, S^1 plays the role of Euclidean time, τ_E , and its period is identified with the inverse temperature β . Additionally,

$$\hat{\rho}_T = \hat{U} \hat{\rho}_0 \hat{U}^{-1} , \quad (7.6)$$

where \hat{U} is a unitary CFT operator that implements the conformal transformation. For example, the primary spinless operators, $\hat{\mathcal{O}}$, of the CFT locally transform as²

$$\hat{\mathcal{O}}_{\mathcal{H}} = \Omega^{\Delta} \hat{U} \hat{\mathcal{O}}_{\mathcal{M}} \hat{U}^{-1} , \quad (7.7)$$

where Δ is the scaling dimension of $\hat{\mathcal{O}}$ and Ω is the conformal factor that relates the metrics on the two manifolds

$$ds_{\mathcal{M}}^2 = \Omega^2 ds_{\mathcal{H}}^2 . \quad (7.8)$$

In what follows we consider separately perturbations of the QFT action, and perturbations associated with either slight changes in the background geometry or mild deformations of the entangling surface Σ .

7.2.1 Geometric perturbations

In general, the modular Hamiltonian depends on the background geometry as well as on the geometry of the entangling surface. The same is true about conformal transformations of \mathcal{M} onto \mathcal{H} that relate the density matrices as in (7.6). Such mappings are sensitive to changes in the background geometry as well as to deformations of the entangling surface Σ . While the former sensitivity is obvious, the latter follows from the fact that (7.6) is valid if and only if the field configurations ϕ_+ and ϕ_- are mapped onto constant slices $\tau_E = 0$ and $\tau_E = \beta$, respectively. Therefore the mapping, if it exists, certainly depends on the details of Σ .

These observations lead us to construct a special foliation of \mathcal{M} that encodes both the background geometry as well as the structure of the entangling surface. Such a foliation for a generic \mathcal{M} and Σ can only be found perturbatively in the distance from the entangling surface. Sufficiently far from Σ caustics may be encountered and our coordinate system will break down. However, this region is not relevant for us. We present here the final answer for the foliation, with the details relegated to Appendix 7.6. To second order in the distance from Σ , the metric on \mathcal{M} is given by

$$\begin{aligned} ds_{\mathcal{M}}^2 &= (\delta_{ab} - \frac{1}{3} \mathcal{R}_{abcd}|_{\Sigma} x^c x^d) dx^a dx^b + (A_i + \frac{1}{3} x^b \varepsilon^{de} \mathcal{R}_{ibde}|_{\Sigma}) \varepsilon_{ac} x^a dx^c dy^i \\ &+ (\gamma_{ij} + 2K_{aij} x^a + x^a x^c (\delta_{ac} A_i A_j + \mathcal{R}_{iacj}|_{\Sigma} + K_{cil} K_{aj}^l)) dy^i dy^j + \mathcal{O}(x^3) , \end{aligned} \quad (7.9)$$

where $\{y^i\}_{i=1}^{d-2}$ and $\{x^a\}_{a=1}^2$ parametrize Σ and the 2-dimensional transverse space, respectively. The entangling surface Σ is located at $x^a = 0$ and γ_{ij} is the corresponding induced

²The subscript on $\hat{\mathcal{O}}$ indicates on which manifold the operator has support.

metric, ε_{ac} is the volume form of the transverse space, whereas $\mathcal{R}_{\mu\nu\alpha\beta}$ and K_{ij}^a are the background and extrinsic curvatures, respectively. Finally, A_i is the analog of the Kaluza-Klein vector field associated with dimensional reduction over the transverse space. Note that by construction the structure of Σ is built into the above ansatz.

The ansatz for the metric with a slightly perturbed background and mildly modified entangling surface Σ can be obtained by varying (7.9) around the unperturbed background. In particular, the metric will take the following form

$$g_{\mu\nu} = \bar{g}_{\mu\nu} + h_{\mu\nu} , \quad (7.10)$$

where $\bar{g}_{\mu\nu}$ is the unperturbed metric of the form (7.9) with known coefficients, while $h_{\mu\nu}$ contains all the information about perturbations that occurred in the background and entangling surface geometries.

If Σ is everywhere a small deformation of the original entangling surface, *e.g.*, if it is a plane everywhere except that in some localized region there is a small ‘‘bump’’, then perturbative analysis applies globally on Σ . However, $h_{\mu\nu}$ does not necessarily even need to be small everywhere on the entangling surface. If, for example, the surface does not globally look like a plane by having a low curvature but long turn, then we can implement a cut and paste procedure suggested in [17]. We cut the surface along regions which are sufficiently flat, compute the entanglement entropy for each section, and then paste the results together. Of course, this cut and paste procedure is not straightforward and there are potential computational subtleties that need to be addressed.

Substituting decomposition (7.10) into the path integral representation of the density matrix, (7.1), and expanding the result around $\bar{g}_{\mu\nu}$ yields,

$$\begin{aligned} [\hat{\rho}]_{\phi_-\phi_+} &= \frac{1}{\mathcal{N}} \int_{\substack{\phi(\mathcal{C}_+) = \phi_+ \\ \phi(\mathcal{C}_-) = \phi_-}} \mathcal{D}\phi e^{-I_0(\phi, \bar{g}_{\mu\nu} + h_{\mu\nu})} \\ &= \frac{1}{\mathcal{N}} \int_{\substack{\phi(\mathcal{C}_+) = \phi_+ \\ \phi(\mathcal{C}_-) = \phi_-}} \mathcal{D}\phi e^{-I_0(\phi, \bar{g}_{\mu\nu})} \left(1 + \frac{1}{2} \int_{\mathcal{M}} T_{\mathcal{M}}^{\mu\nu} h_{\mu\nu} + \dots \right) , \end{aligned} \quad (7.11)$$

where $T_{\mathcal{M}}^{\mu\nu}$ is the energy-momentum tensor of the QFT on the unperturbed Euclidean manifold \mathcal{M}

$$T_{\mathcal{M}}^{\mu\nu} = -\frac{2}{\sqrt{\bar{g}}} \frac{\delta I_0}{\delta \bar{g}_{\mu\nu}} . \quad (7.12)$$

The normalization constant \mathcal{N} appearing in (7.11) is given by

$$\begin{aligned} \mathcal{N} &= \int \mathcal{D}\phi_0 \int_{\phi(\mathcal{C}_+) = \phi(\mathcal{C}_-) = \phi_0} \mathcal{D}\phi e^{-I_0(\phi, \bar{g}_{\mu\nu})} \left(1 + \frac{1}{2} \int_{\mathcal{M}} T_{\mathcal{M}}^{\mu\nu} h_{\mu\nu} + \dots \right) \\ &= \mathcal{N}_0 \left(1 + \frac{1}{2} \int_{\mathcal{M}} \langle \hat{T}_{\mathcal{M}}^{\mu\nu} \rangle_0 h_{\mu\nu} + \dots \right) , \end{aligned} \quad (7.13)$$

where $\langle \hat{T}_{\mathcal{M}}^{\mu\nu} \rangle_0$ is the expectation value of the stress tensor in the state $\hat{\rho}_0$, while \mathcal{N}_0 is the normalization constant of the unperturbed density matrix $\hat{\rho}_0$,

$$\mathcal{N}_0 = \int \mathcal{D}\phi_0 \int_{\phi(\mathcal{C}_+) = \phi(\mathcal{C}_-) = \phi_0} \mathcal{D}\phi e^{-I_0(\phi, \bar{g}_{\mu\nu})} . \quad (7.14)$$

If the path integral representation of $\hat{\rho}_0$ is symmetric under rotations in the transverse space around the entangling surface, then it is convenient to think of it as an effective evolution from the slice \mathcal{C}_+ to the slice \mathcal{C}_- [108, 18, 177]. In this case, based on (7.11), (7.13) and (7.14), one can write

$$[\delta\hat{\rho}]_{\phi_-\phi_+} = \langle \phi_- | \delta\hat{\rho} | \phi_+ \rangle = \frac{1}{2} \int_{\mathcal{M}} \langle \phi_-, \theta_f | \hat{\mathcal{U}}(\theta_f, \theta) \hat{T}_{\mathcal{M}}^{\mu\nu}(\theta) \hat{\mathcal{U}}(\theta, \theta_i) | \phi_+, \theta_i \rangle h_{\mu\nu} - \frac{1}{2} [\hat{\rho}_0]_{\phi_-\phi_+} \int_{\mathcal{M}} \langle \hat{T}_{\mathcal{M}}^{\mu\nu} \rangle_0 h_{\mu\nu}(\theta), \quad (7.15)$$

where we have used the definition $\delta\hat{\rho} = \hat{\rho} - \hat{\rho}_0$, θ is the polar angle around the entangling surface such that θ_i and θ_f equal 0 and 2π respectively, and $\hat{\mathcal{U}}$ is the angular evolution operator given by,

$$\hat{\mathcal{U}}(\theta_2, \theta_1) = \exp\left(-\frac{\theta_2 - \theta_1}{2\pi} \hat{K}_0\right). \quad (7.16)$$

Because of rotational symmetry the expectation value $\langle \hat{T}_{\mathcal{M}}^{\mu\nu} \rangle_0$ is independent of θ . Moreover, the above formula for $\delta\hat{\rho}$ can be written as

$$\delta\hat{\rho} = \frac{1}{2} \int_{\mathcal{M}} \hat{\mathcal{U}}(\theta_f, \theta) \left(\hat{T}_{\mathcal{M}}^{\mu\nu}(\theta) - \langle \hat{T}_{\mathcal{M}}^{\mu\nu} \rangle_0 \right) \hat{\mathcal{U}}(\theta, \theta_i) h_{\mu\nu}. \quad (7.17)$$

Now using eqs.(7.17) and (7.5) we get the following correction to the modular Hamiltonian,

$$\hat{K} = \hat{K}_0 - \frac{1}{2} \int_{\mathcal{M}} \hat{\mathcal{U}}(\theta_i, \theta) \left(\hat{T}_{\mathcal{M}}^{\mu\nu}(\theta) - \langle \hat{T}_{\mathcal{M}}^{\mu\nu} \rangle_0 \right) \hat{\mathcal{U}}(\theta, \theta_i) h_{\mu\nu} + \dots \quad (7.18)$$

The entanglement entropy across Σ can be evaluated by substituting (7.17) into (7.5). This gives

$$S = S_0 + \frac{1}{2} \int_{\mathcal{M}} \langle \hat{T}_{\mathcal{M}}^{\mu\nu} \hat{K}_0 \rangle_c h_{\mu\nu} + \dots, \quad (7.19)$$

where $\langle \dots \rangle_c$ is the connected two point function in the state $\hat{\rho}_0$. We should note that our results (7.18), (7.19) are valid for a general field theory, and not necessarily restricted to a CFT.

Moreover, if we restrict our consideration to conformal field theories, then it is possible to generalize the above results to include the case when the state undergoes a conformal mapping as in (7.6), (7.8). We first recall the rule for conformal transformation of the energy-momentum tensor,

$$T_{\mathcal{M}}^{\mu\nu} = \Omega^{-d-2} \frac{\partial x^\mu}{\partial X^\alpha} \frac{\partial x^\nu}{\partial X^\beta} (T_{\mathcal{H}}^{\alpha\beta} + \mathcal{A}^{\mu\nu}), \quad (7.20)$$

where X^μ are coordinates on \mathcal{H} , x^μ collectively denotes (x^a, y^i) and $\mathcal{A}^{\mu\nu}$ is the higher dimensional analog of the Schwarzian derivative. Hence, from (7.11) we obtain

$$[\hat{U}\hat{\rho}\hat{U}^{-1}]_{\tilde{\phi}_+\tilde{\phi}_-} = \frac{1}{\mathcal{N}} \int_{\substack{\phi(\tau_E=0)=\tilde{\phi}_+ \\ \phi(\tau_E=\beta)=\tilde{\phi}_-}} \mathcal{D}\phi e^{-I_0(\phi, \bar{g}_{\mu\nu})} \left(1 + \frac{1}{2} \int_{\mathcal{H}} \Omega^{-2} (T_{\mathcal{H}}^{\mu\nu} + \mathcal{A}^{\mu\nu}) h_{\mu\nu} + \dots \right), \quad (7.21)$$

where $\tilde{\phi}_\pm$ are the conformally transformed field configurations ϕ_\pm ,

$$|\tilde{\phi}_\pm\rangle = \hat{U}|\phi_\pm\rangle. \quad (7.22)$$

Also note that the normalization constant \mathcal{N} in (7.13) can be rewritten as

$$\mathcal{N} = \mathcal{N}_0 \left(1 + \frac{1}{2} \int_{\mathcal{H}} \Omega^{-2} \langle \hat{T}_{\mathcal{H}}^{\mu\nu} \rangle_T h_{\mu\nu} + \frac{1}{2} \int_{\mathcal{H}} \Omega^{-2} \mathcal{A}^{\mu\nu} h_{\mu\nu} + \dots \right), \quad (7.23)$$

where $\langle \hat{T}_{\mathcal{H}}^{\mu\nu} \rangle_T$ is the thermal expectation value of the stress tensor on \mathcal{H} . Combining eqs. (7.28) and (7.23), yields

$$\hat{U} \delta \hat{\rho} \hat{U}^{-1} = \frac{1}{2} \int_{\mathcal{H}} \hat{\mathcal{U}}_T(\beta, \tau_E) \left(\hat{T}_{\mathcal{H}}^{\mu\nu}(\tau_E) - \langle \hat{T}_{\mathcal{H}}^{\mu\nu} \rangle_T \right) \hat{\mathcal{U}}_T(\tau_E, 0) \Omega^{-2} h_{\mu\nu}, \quad (7.24)$$

where we used the transformation rule (7.6), and $\hat{\mathcal{U}}_T$ is the evolution operator on \mathcal{H} ,

$$\hat{\mathcal{U}}_T(\tilde{\tau}_E, \tau_E) = \exp \left(- (\tilde{\tau}_E - \tau_E) \hat{H} \right), \quad (7.25)$$

where \hat{H} is the Hamiltonian that generates τ_E translations. It is related to the modular Hamiltonian on \mathcal{M} by $\hat{K}_0 = \hat{U}^{-1} \beta \hat{H} \hat{U}$.

Since the von Newman entropy is invariant under unitary transformations, the entanglement entropy across Σ can be evaluated using the density matrix on \mathcal{H} . Substituting (7.24) into (7.5), yields

$$S = S_T + \frac{\beta}{2} \int_{\mathcal{H}} \Omega^{-2} \langle \hat{T}_{\mathcal{H}}^{\mu\nu} \hat{H} \rangle_c h_{\mu\nu} + \dots, \quad (7.26)$$

where S_T is the thermal entropy of the CFT in the state $\hat{\rho}_T$, while $\langle \dots \rangle_c$ is the (thermal) connected two point function on \mathcal{H} . This result is simply a conformal transformation (7.8) of (7.19), accompanied by the rule (7.20). Finally, using eqs.(7.5) and (7.24), we get the following correction to the modular Hamiltonian,

$$\hat{K} = \hat{K}_0 - \frac{1}{2} \hat{U}^{-1} \left(\int_{\mathcal{H}} \hat{\mathcal{U}}_T(0, \tau_E) \left(\hat{T}_{\mathcal{H}}^{\mu\nu}(\tau_E) - \langle \hat{T}_{\mathcal{H}}^{\mu\nu} \rangle_T \right) \hat{\mathcal{U}}_T(\tau_E, 0) \Omega^{-2} h_{\mu\nu} \right) \hat{U} + \dots. \quad (7.27)$$

7.2.2 Relevant perturbations

The main goal of this subsection is to investigate the consequences of small perturbations of the QFT by, *e.g.*, relevant operators. The general form of the reduced density matrix (7.11) that undergoes such a perturbation is

$$\begin{aligned} [\hat{\rho}]_{\phi_+ \phi_-} &= \frac{1}{\mathcal{N}} \int_{\substack{\phi(C_+) = \phi_+ \\ \phi(C_-) = \phi_-}} \mathcal{D}\phi e^{-I_0(\phi, \bar{g}_{\mu\nu}) + g \int_{\mathcal{M}} \mathcal{O}} \\ &= \frac{1}{\mathcal{N}} \int_{\substack{\phi(C_+) = \phi_+ \\ \phi(C_-) = \phi_-}} \mathcal{D}\phi e^{-I_0(\phi, \bar{g}_{\mu\nu})} \left(1 + g \int_{\mathcal{M}} \mathcal{O} + \frac{g^2}{2} \left(\int_{\mathcal{M}} \mathcal{O} \right)^2 + \dots \right), \end{aligned} \quad (7.28)$$

where g is the coupling constant, the scaling dimension of $\hat{\mathcal{O}}$ is $\Delta < d$, and we assume that the effect of the deformation is small, *e.g.*, the theory sits sufficiently close to the UV fixed point. The normalization constant this time is given by

$$\mathcal{N} = \mathcal{N}_0 \left(1 + g \int_{\mathcal{M}} \langle \mathcal{O} \rangle_0 + \frac{g^2}{2} \int_{\mathcal{M}} \int_{\mathcal{M}} \langle \hat{\mathcal{O}} \hat{\mathcal{O}} \rangle_0 + \dots \right), \quad (7.29)$$




Figure 7.2: A sketch of a slightly deformed entangling surface (curved line) in three dimensions. (x_1, x_2) span the transverse space to Σ , while y parametrizes Σ . The foliation (7.9) is designed to capture the geometry of the neighborhood of a given entangling surface Σ .

where the expectation values are taken in the vacuum state. Following now the same steps as in the previous subsection, we obtain the leading order correction to S_0 ,

$$\delta S = g \int_{\mathcal{M}} \langle \hat{\mathcal{O}} \hat{K}_0 \rangle_c . \quad (7.30)$$

If the unperturbed theory is a CFT then the leading correction to S_0 vanishes since $\hat{K}_0 \sim \hat{T}_{\mu\nu}$ and therefore $\langle \hat{K}_0 \hat{\mathcal{O}} \rangle_c = 0$. Hence, we have to resort to the second order perturbation. Using (7.5) yields,

$$\delta S = \frac{g^2}{2} \int_{\mathcal{M}} \int_{\mathcal{M}} \left(\langle \hat{K}_0 \hat{\mathcal{O}} \hat{\mathcal{O}} \rangle_c - \langle \hat{\mathcal{O}} \hat{\mathcal{O}} \rangle_c \right) , \quad (7.31)$$

where certain caution should be taken before evaluating the expectation value of the two point function in the above expression, since according to (7.5) this correlator is evaluated using a path integral with an effective interval of evolution that has to be three times bigger than the interval of evolution used to compute the three point function in (7.31).

We finish this section with a comment that it would be interesting to compare holographic predictions made in [105] with the results based on (7.31), and we hope to report on this in a forthcoming publication.

7.3 Perturbations of a planar entangling surface

In this section we explore the leading order correction (7.19) in the case of small perturbations of a planar entangling surface in flat space. These perturbations could arise from the entangling surface being slightly deformed (see Fig. 7.2), or if the background geometry is weakly curved. For simplicity we restrict our discussion to four spacetime dimensions and evaluate the logarithmic divergence of entanglement entropy. This divergence is universal since it is independent of the regularization scheme.

The entanglement entropy of the unperturbed plane in flat space is closely related to the Unruh effect observed by a uniformly accelerating observer in Minkowski space. Indeed, the reduced density matrix for the vacuum for the semi-infinite domain $x_1 > 0$ is obtained by tracing out the region $x_1 < 0$ on a constant zero Minkowski time slice. This is




Figure 7.3: Transverse space to the entangling surface in the analytically continued spacetime. Σ is located at the origin. The reduced density matrix is given by a path integral (7.1) with fixed boundary conditions ϕ_+ (ϕ_-) on the upper (lower) dashed blue lines.

precisely the region hidden by Rindler horizon and the resulting reduced density matrix has a thermal interpretation in the sense of Unruh [177, 108] with a space dependent temperature that scales as x_1^{-1} . A Rindler observer who is confined to the right wedge, and who is passing through x_1 at $t = 0$, finds himself immersed in a thermal bath of Unruh radiation. The sum of thermal entropies observed by all Rindler observers is the entanglement entropy, and the divergence of the temperature as $x_1 \rightarrow 0$ gives rise to the UV divergence of entanglement entropy.

Analytic continuation of the Rindler wedge to Euclidean signature maps it onto the entire Euclidean space with a puncture at the origin. In Minkowski signature, this puncture corresponds to the Rindler horizon. Furthermore, the analytically continued Rindler Hamiltonian, \hat{H}_R , becomes the generator of rotations in the transverse space to Σ , and as shown in [108] the path integral (7.1) can be written as

$$[\rho_0]_{\phi_+\phi_-} = \langle \phi_- | e^{-2\pi\hat{H}_R} | \phi_+ \rangle . \quad (7.32)$$

In particular, we immediately deduce that the modular Hamiltonian is proportional to the Rindler Hamiltonian, $\hat{K}_0 = 2\pi\hat{H}_R$, which plays the role of the angular evolution operator in the transverse space to Σ . (see Fig. 7.3)

What we have said so far is the standard story for flat space. In a general spacetime, since any region locally looks flat, we expect the leading divergence of the entanglement entropy will be insensitive to the background, in so much as that it scales as an area. The subleading terms of the entanglement entropy are dominated by the region near the entangling surface but have sensitivity to regions slightly away from it as well.

Far away from the surface corrections to the background metric induced by perturbations of the system may be large. However, the further away some region is from the surface, the less relevant it is for the entanglement entropy. Stated in the language of accelerated observers: those who are highly accelerated and close to the Rindler horizon are unlikely to notice a large deviation from a thermal spectrum, while those with small acceleration who are far away find little Unruh radiation and the thermal effect is practically zero.

7.3.1 Calculation

The leading order correction to the entanglement entropy of a flat plane is given by (7.19),

$$\delta S = \pi \int_{R^4} \langle T^{\mu\nu} H_R \rangle h_{\mu\nu}. \quad (7.33)$$

Here H_R is the Rindler Hamiltonian in the unperturbed spacetime³,

$$H_R = - \int_A T_{\mu\nu} \xi^\mu n^\nu, \quad (7.34)$$

where $A = \{x \in R^4 \mid x_2 = 0, x_1 > 0\}$, $\xi = x_1 \partial_2 - x_2 \partial_1$ is the Killing vector field associated with rotational symmetry around the plane at $x^a = 0$, while $n = \partial_2$ is normal to A . Thus,

$$H_R = - \int_A dy dx_1 x_1 T_{22}. \quad (7.35)$$

Substituting H_R into (7.33) gives

$$\delta S = -\pi \int_A d^2x d^2y d\bar{x}_1 d^2\bar{y} \bar{x}_1 h^{\mu\nu}(x, y) \langle T_{\mu\nu}(x, y) T_{22}(\bar{x}, \bar{y}) \rangle. \quad (7.36)$$

Here the coordinates are $x^\mu = (x^a, y^i)$ where x^a with $a = 1, 2$ are orthogonal to the entangling surface (see Fig. 7.3) and y^i with $i = 1, 2$ are along the entangling surface. Also, $\bar{x}_2 = 0$. From (7.9) we find that there are two terms in $h_{\mu\nu}$ that are responsible for the logarithmically divergent contribution to δS . They are

$$h_{ij} = x^a x^c \mathcal{R}_{iacj} \quad (7.37)$$

$$h_{ab} = -\frac{1}{3} \mathcal{R}_{abcd} x^c x^d. \quad (7.38)$$

Note that the $\delta\gamma_{ij}$ term in (7.9) is not relevant as it contributes to the ‘area law’ correction. Also, the cross terms $dx dy$ will give vanishing contributions. Finally, terms proportional to the extrinsic curvatures contribute at second order within our perturbative expansion (since the extrinsic curvature of the plane is zero and the contribution of the linear term vanishes identically).

The connected 2-pt function for the stress tensor for a CFT is given in [139],

$$\langle T_{\mu\nu}(x, y) T_{22}(\bar{x}, \bar{y}) \rangle = \frac{C_T \mathcal{I}_{\mu\nu, 22}}{((x - \bar{x})^2 + (y - \bar{y})^2)^4} \quad (7.39)$$

where

$$\mathcal{I}_{\mu\nu, 22} = I_{\mu 2} I_{\nu 2} - \frac{\delta_{\mu\nu}}{4}, \quad (7.40)$$

with

$$I_{\mu 2} = \delta_{\mu 2} - \frac{2(x - \bar{x})_\mu x_2}{(x - \bar{x})^2 + (y - \bar{y})^2}. \quad (7.41)$$

³The minus sign appears due to the definition (7.12) of the energy-momentum tensor in Euclidean signature.

In Appendix 7.7 we perform the integral (7.36) and find

$$\delta S = \frac{c}{6\pi} \int d^2y (\delta^{ac} \delta^{bd} \mathcal{R}_{abcd} + \delta^{ij} \delta^{ac} \mathcal{R}_{iacj}) \log(\ell/\delta) . \quad (7.42)$$

Here ℓ is the characteristic scale of the perturbations, δ is the UV cut-off, and $C_T = 40c/\pi^4$ with c being the central charge of the CFT defined by the trace anomaly,

$$\langle T^\mu_\mu \rangle = \frac{c}{16\pi^2} \int_{\mathcal{M}} C_{\mu\nu\rho\sigma} C^{\mu\nu\rho\sigma} - \frac{a}{16\pi^2} \int_{\mathcal{M}} E_4 , \quad (7.43)$$

where $C_{\mu\nu\rho\sigma}$ is the Weyl tensor and E_4 is the Euler density,

$$E_4 = R_{\mu\nu\rho\sigma} R^{\mu\nu\rho\sigma} - 4R_{\mu\nu} R^{\mu\nu} + R^2 . \quad (7.44)$$

Our correction (7.42) should be compared with Solodukhin's formula [161] for the universal part of entanglement entropy in the case of a four dimensional CFT,

$$S_{CFT} = \frac{1}{2\pi} \int_{\Sigma} \left[c (\delta^{ac} \delta^{bd} C_{abcd} + K_{ij}^a K_a^{ij} - \frac{1}{2} K^a K_a) - a \mathcal{R}^\Sigma \right] \log(R/\delta) , \quad (7.45)$$

where \mathcal{R}^Σ is the intrinsic curvature of the entangling surface. Of course, for the case of a planar surface in flat space S_{CFT} vanishes identically.

Varying (7.45) around the flat plane embedded in R^d , we obtain to linear order in small perturbations

$$\delta S_{CFT} = \frac{c}{2\pi} \int_{\Sigma} \delta^{ac} \delta^{bd} C_{abcd} \log(R/\delta) = \frac{c}{6\pi} \int_{\Sigma} (\delta^{ac} \delta^{bd} \mathcal{R}_{abcd} + \gamma^{ij} \delta^{ac} \mathcal{R}_{i\hat{a}cj} + \gamma^{ij} \gamma^{kl} \mathcal{R}_{ikjl}) , \quad (7.46)$$

where in the second equality we used the definition of the Weyl tensor. This expression matches (7.42) since the last term is a total derivative in this case, and therefore its integral vanishes. Indeed, the first variation of the Gauss-Codazzi relation (7.89) around the flat plane embedded in flat space gives

$$\gamma^{ij} \gamma^{kl} \mathcal{R}_{ikjl}|_{\Sigma} = \partial^i (\partial^j \delta\gamma_{ij} - \gamma^{mn} \partial_i \delta\gamma_{mn}) , \quad (7.47)$$

where we have used the general variational rule

$$\delta \mathcal{R}_{\Sigma} = -\mathcal{R}_{\Sigma}^{ij} \delta\gamma_{ij} + \nabla^i (\nabla^j \delta\gamma_{ij} - \gamma^{mn} \nabla_i \delta\gamma_{mn}) , \quad (7.48)$$

where ∇_i is covariant derivative compatible with the unperturbed induced metric γ_{ij} .

Before closing this section let us make a couple of comments. First, we note that (7.42) and (7.46) are independent of the central charge a . This is a straightforward consequence of the fact that R^Σ is the Euler density of a two-dimensional manifold, and therefore the last term in (7.45) is a topological invariant that does not change under smooth deformations of the entangling surface and background, *i.e.*,

$$\delta \int_{\Sigma} \mathcal{R}^\Sigma = \int_{\Sigma} \left(\frac{1}{2} \gamma^{ij} \mathcal{R}^\Sigma - \mathcal{R}_{\Sigma}^{ij} \right) \delta\gamma_{ij} = 0 , \quad (7.49)$$




Figure 7.4: We conformally transform between \mathcal{H} (left) and R^d (right). We first map from the $\sigma \equiv u + i\tau$ coordinates of \mathcal{H} to $e^{-\sigma}$ (middle); here the origin is $u = \infty$ and the boundary circle is $u = 0$. We then map via (7.53) to R^d . Dashed lines on the left represent $\tau_E = 0^+, \beta^-$ slices of \mathcal{H} that are mapped through an intermediate step onto $t = 0^\pm$ sides of the cut throughout the interior of the sphere $r = R$ on the right

where by assumption the deformed and original setups approach each other at infinity and we used the fact that Σ is a two-dimensional manifold.

Second, it should be noticed that terms in (7.45) that are quadratic in extrinsic curvature do not contribute to the leading order correction to the entanglement entropy since K_{ij}^a of a flat plane vanishes. To see the effect of extrinsic curvatures one has to study second order perturbations within our formalism and this will be addressed in a forthcoming publication. In order to see the effect of extrinsic curvatures at first order, we now turn to spherical entangling surfaces.

7.4 Perturbations of a spherical entangling surface

In this section the background manifold \mathcal{M} will be identified with R^d , and the entangling surface Σ will be a sphere, S^{d-2} , of radius R . We first show that there is a conformal map that transforms between Euclidean path integral representations of $\hat{\rho}_0$ and $\hat{\rho}_T$ and then apply the analysis of Sec. 7.2 to compute the first order corrections to the entanglement entropy due to slight deformations of R^d and S^{d-2} .

Let us recall that the partition function on $\mathcal{H} \equiv S^1 \times H^{d-1}$ can be evaluated by a path integral on the Euclidean background

$$ds_{\mathcal{H}}^2 = d\tau_E^2 + R^2 (du^2 + \sinh^2 u d\Omega_{d-2}^2), \quad (7.50)$$

where the Euclidean time coordinate has period $\Delta\tau_E = \beta = 2\pi R$. In the following, it will be convenient to introduce complex coordinates:

$$\sigma = u + i\tau_E/R \quad \text{and} \quad \omega = r + it_E, \quad (7.51)$$

where the latter will be used below to describe a conformally mapped geometry. Note that both u and r are radial coordinates, and we must have $\text{Re}(\sigma) = u > 0$ and $\text{Re}(\omega) = r > 0$.




Figure 7.5: We show the constant τ_E slices (blue) and constant u slices (red) in the (r, t_E) plane (7.56). The sphere is located at $r/R = 1$, $t_E = 0$ and corresponds to $u \rightarrow \infty$. The vertical line ($r = 0$) corresponds to $u = 0$.

With the first of these new coordinates, the above metric (7.50) can be written as

$$ds_{\mathcal{H}}^2 = R^2 \left(d\sigma d\bar{\sigma} + \sinh^2 \left(\frac{\sigma + \bar{\sigma}}{2} \right) d\Omega_{d-2}^2 \right). \quad (7.52)$$

Now we make the coordinate transformation

$$e^{-\sigma} = \frac{R - \omega}{R + \omega}. \quad (7.53)$$

Since we are considering $d \geq 3$ there is no guarantee that this holomorphic change of coordinates will result in a conformal transformation. However, one can readily verify the above metric (7.52) becomes

$$\begin{aligned} ds_{\mathcal{H}}^2 &= \Omega^{-2} \left[d\omega d\bar{\omega} + \left(\frac{\omega + \bar{\omega}}{2} \right)^2 d\Omega_{d-2}^2 \right] \\ &= \Omega^{-2} \left[dt_E^2 + dr^2 + r^2 d\Omega_{d-2}^2 \right], \end{aligned} \quad (7.54)$$

where

$$\Omega^{-1} = \frac{2R^2}{|R^2 - \omega^2|} = \cosh u + \cos(\tau_E/R). \quad (7.55)$$

Hence, after eliminating the conformal factor Ω^{-2} in the second line of (7.54), we recognize that the final line element is simply the metric on d -dimensional flat space. Written explicitly in terms of real coordinates, (7.53) takes the form (see Fig. 7.5)

$$\begin{aligned} r &= R \frac{\sinh u}{\cosh u + \cos(\tau_E/R)}, \\ t_E &= R \frac{\sin(\tau_E/R)}{\cosh u + \cos(\tau_E/R)}. \end{aligned} \quad (7.56)$$

Note that (7.56) can be obtained by analytic continuation to Euclidean time of the conformal mapping between causal domain of a sphere in Minkowski space and Lorentzian \mathcal{H} [48]. Under this analytic continuation the boundary of the causal domain shrinks to a sphere of

radius R while its interior spans the rest of Euclidean space. Note also that the conformal factor (7.55) is everywhere regular on the Euclidean space excluding the sphere of radius R .

Eq. (7.56) implements a simple bijection between \mathcal{H} and R^d . Furthermore, the conformal boundary of the hyperbolic space H^{d-1} is mapped onto a $(d-2)$ -dimensional sphere of radius R sitting on a $t_E = 0$ slice of R^d . Finally, constant time slices $\tau_E = 0^+$ and $\tau_E = \beta^-$ are mapped respectively onto $t_E = 0^-$ and $t_E = 0^+$ of the cut $\mathcal{C} = \{x^\mu \in R^d | 0 \leq r < R, t_E = 0\}$. Hence we have shown that the conformal map (7.53) transforms between the thermal state on \mathcal{H} and the entangled state $\hat{\rho}_0$ for a spherical region in R^d .

7.4.1 Geometric perturbations

The metric on \mathcal{M} is given by

$$ds_{\mathcal{M}}^2 = dt_E^2 + dr^2 + r^2 d\Omega_{d-2}^2 . \quad (7.57)$$

We rewrite it as

$$ds_{\mathcal{M}}^2 = dx_1^2 + dx_2^2 + \left(1 + \frac{2}{R} x_2 + \frac{x_2^2}{R^2}\right) ds_{\Sigma}^2 , \quad (7.58)$$

where we defined a new set of coordinates $t_E = x_1$, $r = R + x_2$ with $-R \leq x_2 < \infty$, and ds_{Σ}^2 is the line element on a sphere of radius R

$$ds_{\Sigma}^2 = R^2 d\Omega_{d-2}^2 . \quad (7.59)$$

The extrinsic curvatures of Σ in this case are given by

$$K_{ij}^{\hat{1}} = 0, \quad K_{ij}^{\hat{2}} = \frac{\gamma_{ij}}{R} , \quad (7.60)$$

where γ_{ij} is the induced metric on a sphere of radius R .

We assume that the background curvature, induced metric, and extrinsic curvatures acquire corrections $\mathcal{R}_{\mu\nu\alpha\beta}$, $\delta\gamma_{ij}$ and δK_{ij}^c parametrized by some infinitesimal parameter ϵ

$$R^2 \mathcal{R}_{\mu\nu\alpha\beta} \sim R \delta K_{ij}^c \sim \delta\gamma_{ij} \sim \epsilon \quad (7.61)$$

As a result, the slightly perturbed metric can be expressed in the form of (7.10), where $\bar{g}_{\mu\nu}$ is given by (7.58), while $h_{\mu\nu}$ takes the form

$$\begin{aligned} h_{\mu\nu} dx^\mu dx^\nu &= -\frac{1}{3} \mathcal{R}_{abcd}|_{\Sigma} x^c x^d dx^a dx^b + \left(A_i + \frac{1}{3} x^b \varepsilon^{de} \mathcal{R}_{ibde}|_{\Sigma}\right) \varepsilon_{ac} x^a dx^c dy^i \\ &+ \left(\delta\gamma_{ij} + 2\delta K_{aij} x^a + x^a x^c \left(\mathcal{R}_{iacj}|_{\Sigma} + \frac{2}{R} \delta_c^{\hat{2}} \delta K_{aij} - \delta_a^{\hat{2}} \delta_c^{\hat{2}} \frac{\delta\gamma_{ij}}{R^2}\right)\right) dy^i dy^j + \mathcal{O}(\epsilon^2) . \end{aligned} \quad (7.62)$$

Here y^i are just the standard spherical angles multiplied by R . In what follows we use the unperturbed induced metric γ_{ij} to raise and lower the indices on the entangling surface.

To use (7.26) we need the connected correlator $\langle \hat{T}_{\mathcal{H}}^{\mu\nu} \hat{H} \rangle_c$. Since the Hamiltonian is conserved and hyperbolic space is maximally symmetric, the correlator is insensitive to

where the operators are inserted, and therefore it is constant on \mathcal{H} . In particular, it was shown in [142] that

$$\langle \hat{T}_{\mathcal{H}}^{\tau_E \tau_E} \hat{H} \rangle_c = -\frac{(d-1)}{2^{d+2} \pi^2 d} \frac{\Omega_{d+2}}{R^{d+1}} C_T, \quad \Omega_d = \frac{2\pi^{\frac{d+1}{2}}}{\Gamma\left(\frac{d+1}{2}\right)}, \quad (7.63)$$

where C_T is a ‘‘central charge’’ common to CFTs in any number of dimensions. In four dimensions this coefficient is related to the standard central charge c which appears as the coefficient of the (Weyl)² term in the trace anomaly⁴ $C_T = (40/\pi^4)c$.

Since the background geometry is conformally flat, all Weyl invariants of the trace anomaly vanish. Further, the background is the direct product of two lower dimensional geometries which dictates that the Euler density is also zero. Hence, the trace anomaly vanishes in this particular background. Using the tracelessness of the energy-momentum tensor and maximal symmetry of H^{d-1} yields

$$\langle \hat{T}_{\mathcal{H}j}^i \hat{H} \rangle_c = \frac{\delta_j^i}{2^{d+2} \pi^2 d} \frac{\Omega_{d+2}}{R^{d+1}} C_T, \quad (7.64)$$

where indices i, j run over the hyperbolic space H^{d-1} . It follows from (7.26) that the off diagonal elements of (7.62) do not contribute to linear order corrections since the connected correlator $\langle \hat{T}_{\mathcal{H}}^{\mu\nu} \hat{H} \rangle_c$ is diagonal.

Eqs. (7.26), (7.62), (7.63) and (7.64) give a general solution for linear perturbations of spherical regions in flat space. In the next subsection we carry out a particular calculation in $d = 4$ and show that our formula (7.26) agrees with the known results in the literature.

7.4.2 Calculation

Let us evaluate the logarithmic divergence of entanglement entropy for a four dimensional CFT using our result (7.26). This divergence is universal since it is independent of the details of regularization scheme, and it was shown in [48] that for a perfect sphere in flat space it is entirely fixed by the coefficient of the A-type trace anomaly. In particular, in $d = 4$ the universal divergence takes the form

$$S_{univ} = -4a \log(R/\delta), \quad (7.65)$$

Here δ is the UV cut-off and a is the central charge defined in (7.43),

As argued in [48], the leading order term in (7.26) satisfies $S_T = S_{univ}$. The logarithmic divergence within the thermal computation on \mathcal{H} is a result of the divergent volume of hyperbolic space. This IR divergence emerges because we have a uniform entropy density, but the volume of H^3 is infinite. Hence, to regulate the thermal entropy in \mathcal{H} we integrate to some maximum radius, $u = u_{max}$ where $u_{max} \gg 1$. On the other hand, the divergence of entanglement entropy is entirely due to short distance fluctuations in the vicinity of Σ . Thus, in order to regulate this divergence we exclude the δ -neighborhood of the entangling surface Σ , where $\delta/R \ll 1$. These two UV and IR cut-off’s should be

⁴See (7.43) for the definition of the central charges that we use throughout this chapter.

related by the conformal mapping between the two spaces. If we focus on the $t_E = 0$ slice (or equivalently the $\tau_E = 0$ slice), then (7.56), yields the following relation [48]

$$1 - \frac{\delta}{R} = \frac{\sinh u_{max}}{\cosh u_{max} + 1} \Rightarrow u_{max} \simeq \log(R/\delta). \quad (7.66)$$

To get corrections to the leading order result we substitute eqs. (7.62), (7.63) and (7.64) into (7.26) and use eqs. (7.50), (7.55) and (7.56) to carry out the integrals. The final answer for the logarithmically divergent part of the integrals is given by

$$\delta S_{univ} = \frac{c}{6\pi} \int_{\Sigma} \left(\frac{\delta\gamma}{R^2} - \frac{2}{R} \delta_c^2 \delta K^c + \gamma^{ij} \delta^{ac} \mathcal{R}_{i\hat{a}cj}|_{\Sigma} + \delta^{ac} \delta^{bd} \mathcal{R}_{abcd} \right) \log(R/\delta), \quad (7.67)$$

where Σ is a sphere of radius R , $\delta\gamma$ and δK^c are the traces of the perturbations $\delta\gamma_{ij}$ and δK_{ij}^c , and we used (7.56) to evaluate the components of $h_{\mu\nu}$ in coordinates (7.50),

$$\begin{aligned} h_{uu} &= -\frac{R^2 \Omega^4}{6} \delta^{ac} \delta^{bd} \mathcal{R}_{abcd} e^{-2u} \sin^2(\tau_E/R), \\ h_{\tau_E \tau_E} &= -\frac{R^2 \Omega^4}{6} \delta^{ac} \delta^{bd} \mathcal{R}_{abcd} (1 + e^{-u} \cos(\tau_E/R))^2, \\ h_{u\tau_E} &= \frac{R^2 \Omega^4}{24} \delta^{ac} \delta^{bd} \mathcal{R}_{abcd} e^{-u} \sin(\tau_E/R) (1 + e^{-u} \cos(\tau_E/R)). \end{aligned} \quad (7.68)$$

Let us now compare (7.67) with Solodukhin's formula (7.45). For the case of a sphere in flat space, this formula reduces to (7.65). Corrections to (7.65) can be evaluated by varying (7.45) around sphere of radius R embedded into R^d . Provided that variations are small and satisfy (7.61), we get (7.46) again. The latter is not a coincidence, it is a straightforward consequence of the fact that (7.45) is Weyl invariant while the two setups (a plane and a sphere in flat space) are conformally equivalent. To see it explicitly, let us write the metric around flat plane as follows

$$ds^2 = dr^2 + r^2 d\theta^2 + \delta_{ij} dy^i dy^j = \frac{r^2}{R^2} \left(d\tau_E^2 + \frac{R^2}{r^2} (dr^2 + \delta_{ij} dy^i dy^j) \right), \quad (7.69)$$

where we have defined $\tau_E = R\theta$ and used polar coordinates in the transverse space to the plane. Stripping off conformal factor on the right hand side of this expression leaves us with the metric on \mathcal{H} in Poincare patch. Note that conformal factor is everywhere regular in the punctured Euclidean space (or analytically continued Rindler wedge), and the plane at $r = 0$ is mapped onto conformal boundary of \mathcal{H} .

Hence, we have shown that two setups are conformally equivalent to \mathcal{H} with entangling surfaces being mapped onto conformal boundary of the hyperbolic space. Therefore they are conformally equivalent to each other. In particular, it follows that quadratic in extrinsic curvatures term of (7.45),

$$I = \int_{\Sigma} \left(K_{ij}^a K_a^{ij} - \frac{1}{2} K^a K_a \right), \quad (7.70)$$

does not contribute to the first variation of entanglement entropy around spherical region. This claim can be verified by direct computation, however there is a simple argument based

on the Weyl symmetry inherent to the problem. Indeed, this term is separately Weyl invariant and its first variation vanishes in the planar case, therefore the same is true for conformally equivalent spherical region in flat space. In our forthcoming publication we are going to explore the second order perturbation theory to uncover the effect of extrinsic curvatures on the entanglement entropy.

Let us now show that (7.46) agrees with (7.67). Varying the Gauss-Codazzi relation (7.89) around the unperturbed sphere of radius R embedded in flat space gives

$$\gamma^{ij}\gamma^{kl}\mathcal{R}_{ikjl}|_{\Sigma} = \frac{1}{R^2}\delta\gamma - \frac{2}{R}\delta_c^{\hat{2}}\delta K^c + \nabla^i(\nabla^j\delta\gamma_{ij} - \gamma^{mn}\nabla_i\delta\gamma_{mn}) , \quad (7.71)$$

where we have used the variational rule (7.48). Substituting this result into (7.46) gives (7.67)⁵.

7.5 Notation

In this appendix we explain our notation and conventions. Greek indices run over the entire background, whereas Latin letters from the ‘second’ half of the alphabet i, j, \dots represent directions along the entangling surface.

There is a pair of independent orthonormal vectors which are orthogonal to Σ , we denote them by n_a^μ (with $a = 1, 2$), where the letters from the beginning of the Latin alphabet are used to denote the frame or tangent indices in the transverse space. Then delta Kronecker $\delta_{ab} = n_a^\mu n_b^\nu g_{\mu\nu}$ is the metric in the tangent space spanned by these vectors and δ^{ab} is the inverse of this metric.

We also have tangent vectors t_i^μ to Σ , which are defined in the usual way with $t_i^\mu = \partial x^\mu / \partial y^i$, where x^μ and y^i are the coordinates in the full embedding space and along the surface, respectively. The induced metric is then given by $\gamma_{ij} = t_i^\mu t_j^\nu g_{\mu\nu}$. It can also be defined as a bulk tensor with $\gamma_{\mu\nu} = g_{\mu\nu} - g_{\mu\nu}^\perp$, where $g_{\mu\nu}^\perp = \delta_{ab} n_\mu^a n_\nu^b$ is the metric in the space transverse to Σ . The second fundamental forms are defined for the entangling surface with $K_{ij}^a = t_i^\mu t_j^\nu \nabla_\mu n_\nu^a$, where ∇_μ is covariant derivative compatible with $g_{\mu\nu}$. We use this definition to construct the bulk vector $K_{ij}^\mu = n_a^\mu K_{ij}^a$.

Next we define the volume form in the tangent space spanned by the normal vectors

$$\begin{aligned} \varepsilon_{ab} &= -\varepsilon_{ba} , & \varepsilon_{\hat{1}\hat{2}} &= 1 , \\ \varepsilon^{ab} &= \delta^{ac}\delta^{bd}\varepsilon_{cd} = \varepsilon_{ab} . \end{aligned} \quad (7.72)$$

Using this definition the volume form in the transverse space can be written as $\varepsilon_{\mu\nu} = \varepsilon_{ab} n_\mu^a n_\nu^b$. We use $g_{\mu\nu}^\perp$ to raise and lower the indices in the transverse space, while indices along the direction of the entangling surface are raised and lowered with the induced metric $\gamma_{\mu\nu}$. Note that the following useful identity holds,

$$\varepsilon_{\mu\nu}\varepsilon_{\rho\sigma} = g_{\mu\rho}^\perp g_{\nu\sigma}^\perp - g_{\mu\sigma}^\perp g_{\nu\rho}^\perp . \quad (7.73)$$

Finally, our convention for the curvature tensor is given by

$$\mathcal{R}_{\mu\nu\rho\sigma} = \frac{1}{2}(g_{\mu\sigma,\nu\rho} + g_{\nu\rho,\mu\sigma} - g_{\mu\rho,\nu\sigma} - g_{\nu\sigma,\mu\rho}) + \Gamma_{\nu\rho,\chi}^\chi \Gamma_{\mu\sigma}^\chi - \Gamma_{\nu\sigma,\chi}^\chi \Gamma_{\mu\rho}^\chi . \quad (7.74)$$

⁵The total derivative in (7.71) does not contribute since Σ in our case has no boundaries.

7.6 Foliation of \mathcal{M} in the vicinity of the entangling surface

In this appendix we build a particular foliation of \mathcal{M} in the vicinity of Σ . First, we choose some parametrization $\{y^i\}_{i=1}^{d-2}$ for the entangling surface Σ , then for a given point $O \in \Sigma$ we fill the transverse space with geodesics radiating orthogonally out from O . For each point p on the resulting two-dimensional manifold, T_O , we find a geodesic that connects it to O , such that p lies a unit affine parameter from O . Tangent vector to such a geodesic at O can be expanded in terms of a chosen two-dimensional basis n_a^μ . We give its components the names x^a and choose them as coordinates on T_O . Together $\{y^i, x^a\}$ parametrize \mathcal{M} in the vicinity of Σ .

Note that we keep the parametrization of the entangling surface unspecified and therefore the final answer for entanglement entropy will be symmetric with respect to reparametrizations of Σ . On the other hand, choosing a particular foliation of the transverse space does not destroy general covariance of the entanglement entropy since the final answer is obtained by integrating out this space.

By construction, the following relations hold

$$n_\mu^a = \delta_\mu^a, \quad t_i^\mu = \delta_i^\mu, \quad g_{\mu\nu}^\perp = \delta_{ac} \delta_\mu^a \delta_\nu^c, \quad g_{ia} = 0 \text{ on } \Sigma. \quad (7.75)$$

In particular, δ_{ac} plays the role of the transverse metric in this foliation and one can readily evaluate the extrinsic curvatures of Σ ,

$$K_{ij}^a = \nabla_i n_j^a|_\Sigma = \frac{1}{2} \delta^{ac} \partial_c g_{ij}|_\Sigma. \quad (7.76)$$

Hence,

$$g_{ij} = \gamma_{ij} + 2K_{aij} x^a + \mathcal{O}(x^2). \quad (7.77)$$

Furthermore, geodesics radiating orthogonally out from a given point $y \in \Sigma$ take the form $x^a(\tau) = v^a \tau$, where v^a belongs to the two-dimensional tangent space spanned by two normal vectors at y . Substituting this parametrization into the geodesic equation yields

$$\Gamma_{ac}^\mu v^a v^c = 0 \Rightarrow \Gamma_{ac}^\mu = 0 \text{ at } O. \quad (7.78)$$

This identity can be further generalized by differentiating the geodesic equation n times with respect to τ and setting $\tau = 0$. This gives

$$\partial_{(d_1} \partial_{d_2} \cdots \partial_{d_n} \Gamma_{ac}^\mu) = 0 \text{ at } O, \quad (7.79)$$

where as usual (\cdots) denotes symmetrization with respect to the indices within the parenthesis. This result (7.79) with index μ in the transverse space can be used to derive the expansion of the metric on T_y ,

$$g_{ab}(x, y) = \delta_{ab} - \frac{1}{3} \mathcal{R}_{acbd}(y) x^c x^d - \frac{1}{6} \partial_e \mathcal{R}_{acbd}(y) x^c x^d x^e + \mathcal{O}(x^4). \quad (7.80)$$

Moreover, it follows from $\Gamma_{ac}^i|_\Sigma = \partial_{(d} \Gamma_{ac)}^i|_\Sigma = 0$ that Taylor expansion of g_{ic} in the vicinity of Σ can be written as follows

$$g_{ic} = \left(A_i + \frac{1}{3} x^b \varepsilon^{de} \mathcal{R}_{ibde}|_\Sigma \right) x^a \varepsilon_{ac} + \mathcal{O}(x^3), \quad (7.81)$$

where we have introduced a vector field that lives on Σ

$$A_i = \frac{1}{2} \varepsilon^{ac} \partial_a g_{ic} \Big|_{\Sigma} , \quad (7.82)$$

and used the following identity that holds for our foliation

$$\mathcal{R}_{ibac} \Big|_{\Sigma} = \partial_b \partial_{[a} g_{c]i} \Big|_{\Sigma} , \quad (7.83)$$

where $[\dots]$ denotes antisymmetrization with respect to the indices inside the square brackets.

We only need to compute $\mathcal{O}(x^2)$ term in (7.77) to get the expansion of the full metric to second order in the distance from the entangling surface. We first note that Christoffel symbols with at least one index in the transverse space are given by

$$\Gamma_{ac}^{\mu} \Big|_{\Sigma} = 0 , \quad \Gamma_{ic}^a \Big|_{\Sigma} = -\varepsilon^a{}_c A_i , \quad \Gamma_{ij}^a \Big|_{\Sigma} = -K_{ij}^a , \quad \Gamma_{ia}^j \Big|_{\Sigma} = K_{ai}^j . \quad (7.84)$$

Now using (7.74), we obtain

$$\mathcal{R}_{iajb} \Big|_{\Sigma} = \frac{1}{2} \varepsilon_{ba} F_{ij} - \frac{1}{2} \partial_a \partial_b g_{ij} \Big|_{\Sigma} + \delta_{ab} A_i A_j + K_{bil} K_{aj}^l , \quad (7.85)$$

where $F_{ij} = \partial_i A_j - \partial_j A_i$ is the field strength. Symmetrizing this expression with respect to a and b , yields

$$\frac{1}{2} \partial_a \partial_b g_{ij} \Big|_{\Sigma} = \mathcal{R}_{i(ab)j} \Big|_{\Sigma} + \delta_{ab} A_i A_j + \frac{1}{2} (K_{bil} K_{aj}^l + K_{ail} K_{bj}^l) , \quad (7.86)$$

where (\dots) means symmetrization with respect to the indices inside the parenthesis. Hence (7.77) to second order in x^a takes the form

$$g_{ij} = \gamma_{ij} + 2K_{aij} x^a + x^a x^c (\delta_{ac} A_i A_j + \mathcal{R}_{i(ac)j} \Big|_{\Sigma}) + x^a x^c K_{cil} K_{aj}^l + \mathcal{O}(x^3) . \quad (7.87)$$

Altogether eqs. (7.80), (7.81) and (7.87) correspond to the second order expansion of the full metric $g_{\mu\nu}$ in the vicinity of Σ . To linear order in the distance from Σ this metric takes the simple form,

$$ds^2 = \delta_{ac} dx^a dx^c + 2A_i \varepsilon_{ac} x^a dx^c dy^i + (\gamma_{ij} + 2K_{aij} x^a) dy^i dy^j + \mathcal{O}(x^2) . \quad (7.88)$$

Note that using the definition (7.74) and (7.84), one can evaluate various components of the Riemann tensor that were not necessary so far. For instance, considering directions along the entangling surface Σ yields the well known Gauss-Codazzi identity

$$\mathcal{R}_{ijkl} \Big|_{\Sigma} = \mathcal{R}_{ijkl}^{\Sigma} + K_{jk}^a K_{ail} - K_{jl}^a K_{aik} , \quad (7.89)$$

where $\mathcal{R}_{ijkl}^{\Sigma}$ is the intrinsic curvature tensor on Σ .

Furthermore,

$$\mathcal{R}_{ijab} \Big|_{\Sigma} = \varepsilon_{ba} F_{ij} + K_{bil} K_{aj}^l - K_{ail} K_{bj}^l , \quad (7.90)$$

This identity can be used to express the field strength in terms of the background curvature and extrinsic geometry of Σ .

Finally,

$$\begin{aligned} \mathcal{R}_{ijla} \Big|_{\Sigma} &= \nabla_i K_{ajl} - \nabla_j K_{ail} + 2\varepsilon_{ba} A_{[i} K_{j]l}^b , \\ \mathcal{R}_{abcd} \Big|_{\Sigma} &= \mathcal{R}_{abcd}^T \Big|_{\Sigma} , \end{aligned} \quad (7.91)$$

where ∇_i is the covariant derivative compatible with the induced metric on Σ and \mathcal{R}_{abcd}^T is the intrinsic curvature tensor of the transverse space, T_y , at a given point $y \in \Sigma$.

7.7 Intermediate calculations for Sec. 7.3

In this Appendix we evaluate the integral (7.36) appearing in the calculation of the first order correction to the entanglement entropy for a deformed plane in a weakly curved background. First we consider the contribution of the metric perturbation with indices in the direction of the entangling surface, *i.e.*, $h_{ij} = x^a x^c R_{iacj}$. In this case (7.40) is given by,

$$\mathcal{I}_{ij,22} = \frac{4 x_2^2 (y - \bar{y})_i (y - \bar{y})_j}{((x - \bar{x})^2 + (y - \bar{y})^2)^2} - \frac{1}{4} \delta_{ij} \quad (7.92)$$

We begin evaluating (7.36) by first doing the integral over \bar{y} through a change of variables $\bar{y} \rightarrow \bar{y} + y$ giving

$$\delta S_1 = -\frac{\pi^2}{10} C_T \int_{\bar{x}_1 > 0} d^2 x d^2 y d\bar{x}_1 \bar{x}_1 \frac{\delta^{ij} h_{ij}}{((x_1 - \bar{x}_1)^2 + x_2^2)^3} \left(\frac{x_2^2}{(x_1 - \bar{x}_1)^2 + x_2^2} - \frac{5}{6} \right). \quad (7.93)$$

Next, we carry out the \bar{x}_1 integral and introduce polar coordinates in the transverse space, $x_1 = r \cos \theta$, $x_2 = r \sin \theta$,

$$\delta S_1 = \frac{\pi^2}{240} C_T \int d^2 y d\theta \frac{dr}{r^3} \delta^{ij} h_{ij}, \quad (7.94)$$

As expected, the integral over r exhibits logarithmic divergence close to the entangling surface at $r = 0$. Hence, we introduce a UV cut off, δ , to regularize divergence and integrate over r and θ

$$\delta S_1 = \frac{c}{6\pi} \int d^2 y \delta^{ij} \delta^{ac} \mathcal{R}_{iacj} \log(\ell/\delta), \quad (7.95)$$

where ℓ is characteristic scale of small perturbations, and we used the value of $C_T = (40/\pi^4)c$ in four spacetime dimensions.

Next we calculate the contribution of perturbations in the transverse space, *i.e.*, $h_{ab} = -\frac{1}{3} R_{acbd} x^c x^d$. Using $\mathcal{I}_{ab,22}$ from (7.40) and performing the integral over \bar{y} in (7.36) yields

$$\begin{aligned} \delta S_2 &= -\pi^2 C_T \int_{\bar{x}_1 > 0} d^2 x d^2 y d\bar{x}_1 \bar{x}_1 h_{ab}(x, y) \\ &\times \left(\frac{1}{3} \frac{\delta_{a2} \delta_{b2} - \delta_{ab}/4}{((x_1 - \bar{x}_1)^2 + x_2^2)^3} - \frac{x_2 (x - \bar{x})_b \delta_{a2}}{((x_1 - \bar{x}_1)^2 + x_2^2)^4} + \frac{4}{5} \frac{x_2^2 (x - \bar{x})_a (x - \bar{x})_b}{((x_1 - \bar{x}_1)^2 + x_2^2)^5} \right) \end{aligned}$$

As before, we perform the \bar{x}_1 integral, introduce polar coordinates in the transverse space, substitute h_{ab} , carry out θ integral, and finally get

$$\delta S_2 = \frac{c}{6\pi} \int d^2 y \delta^{ac} \delta^{bd} \mathcal{R}_{abcd} \log(\ell/\delta) \quad (7.96)$$

Combined with (7.95), we have thus recovered (7.42).

Part II

Emergent Spacetime with $\Lambda > 0$

In this second part we turn to the harder question of finding the holographic dual to the multiverse. As the correct theory is not yet known, we seek to make progress by extracting general lessons about holography from AdS/CFT. We make a proposal for how to implement the analog of the UV/IR prescription for the multiverse context so as to define a bulk measure. We then test our measure and others in the literature to see how well they match with observation. Finally, we point out a serious difficulty encountered by all measure proposals. This part of the thesis is organized as follows.

In chapter 8, based on [28], we propose to regulate the infinities of eternal inflation by relating a late time cut-off in the bulk to a short distance cut-off on the future boundary. The light-cone time of an event is defined in terms of the volume of its future light-cone on the boundary. We seek an intrinsic definition of boundary volumes that makes no reference to bulk structures. This requires taming the fractal geometry of the future boundary, and lifting the ambiguity of the conformal factor. We propose to work in the conformal frame in which the boundary Ricci scalar is constant. We explore this proposal in the FRW approximation for bubble universes. Remarkably, we find that the future boundary becomes a round three-sphere, with smooth metric on all scales. Our cut-off yields the same relative probabilities as a previous proposal that defined boundary volumes by projection into the bulk along timelike geodesics. Moreover, it is equivalent to an ensemble of causal patches defined without reference to bulk geodesics. It thus yields a holographically motivated and phenomenologically successful measure for eternal inflation.

In chapter 9, based on [29], we test different measures. We show that the geometry of cutoffs on eternal inflation strongly constrains predictions for the timescales of vacuum domination, curvature domination, and observation. We consider three measure proposals: the causal patch, the fat geodesic, and the apparent horizon cutoff, which is introduced here for the first time. We impose neither anthropic requirements nor restrictions on landscape vacua. For vacua with positive cosmological constant, all three measures predict the double coincidence that most observers live at the onset of vacuum domination and just before the onset of curvature domination. The hierarchy between the Planck scale and the cosmological constant is related to the number of vacua in the landscape. These results require only mild assumptions about the distribution of vacua (somewhat stronger assumptions are required by the fat geodesic measure). At this level of generality, none of the three measures are successful for vacua with negative cosmological constant. Their applicability in this regime is ruled out unless much stronger anthropic requirements are imposed.

Finally, in chapter 10, based on [30], we point out a difficulty with all existing measures. Present treatments of eternal inflation regulate infinities by imposing a geometric cutoff. We point out that some matter systems reach the cutoff in finite time. This implies a nonzero probability for a novel type of catastrophe. According to the most successful measure proposals, our galaxy is likely to encounter the cutoff within the next 5 billion years.

Chapter 8

Boundary Definition of a Multiverse Measure

8.1 Introduction

String theory gives rise to an enormous multiverse where the constants of nature are locally constant but vary over extremely large distance scales [41, 109]. In this context, it is natural to make predictions by counting. The relative probability of two events (for example, two different outcomes of a cosmological or laboratory measurement) is defined by their relative abundance,

$$\frac{p_1}{p_2} = \frac{N_1}{N_2}, \quad (8.1)$$

where N_1 is the expectation value of the number of times an event of type 1 occurs in the multiverse.

There is ambiguity in computing the ratio N_1/N_2 : naively, both numbers are infinite. Starting from finite initial conditions, eternal inflation produces an infinite spacetime volume in which everything that can happen does happen an infinite number of times. A procedure for regulating this divergence is called a measure. Different measures can lead to different relative probabilities starting from an otherwise identical theory.

A number of measures have been proposed, including [125, 123, 71, 72, 70, 76, 179, 178, 24, 121, 122, 141, 73, 182, 183, 184, 126, 26], but over the past several years a vigorous phenomenological effort (e.g., [77, 62, 145, 63, 159, 157, 37, 43, 137, 33, 158, 38, 140, 53, 34, 155, 35, 144, 39, 87]) has focused attention on a few simple proposals that remain viable. In this chapter, we will focus on the light-cone time cut-off, which arises from an analogy with the AdS/CFT correspondence. Given a short-distance cut-off on the boundary conformal field theory, the radial position of the corresponding bulk cut-off [165] can be obtained from causality alone [42], without reference to the details of the bulk-boundary correspondence. In the multiverse, the time of a bulk event can similarly be defined in terms of a scale on the future boundary of the multiverse [73]. The simplest causal relation is to associate to each bulk event a boundary scale given by the volume of its future light-cone on the boundary [26] (see Fig. 8.1).¹

¹A different bulk-boundary relation was proposed in Ref. [73]. One of us (RB) has argued that this relation




Figure 8.1: Constant light-cone size on the boundary defines a hypersurface of constant “light-cone time” in the bulk. The green horizontal lines show two examples of such hypersurfaces. They constitute a preferred time foliation of the multiverse. In the multiverse, there are infinitely many events of both type 1 and type 2 (say, two different values of the cosmological constant measured by observers). Their relative probability is defined by computing the ratio of the number of occurrences of each event prior to the light-cone time t , in the limit as $t \rightarrow \infty$.

There is a remaining ambiguity: how is the volume of the future light-cone to be defined? Near the future boundary, physical distances diverge in inflating regions and go to zero near singularities. Ref. [26] defined boundary volumes by erecting a congruence of geodesics orthogonal to a fixed, fiducial bulk hypersurface and projecting future infinity onto the fiducial hypersurface along the geodesics. The definition of Ref. [26] is quite robust²; in particular, it does not matter whether geodesics expand or collapse or cross.

We would like to find an alternative definition of the boundary volume which does not rely on such an elaborate bulk construction. The geodesics used to project the light-cone onto the fiducial hypersurface do only one thing for us: they define boundary volumes. Yet they encode an enormous amount of geometric bulk information (the exact path of each geodesic, the expansion and shear of nearby geodesics, etc.), most of which is never used for any purpose. They bear no apparent relation to any physical system; for example, they do not represent the worldlines of actual particles. Moreover, the construction takes an absurdly classical viewpoint of the bulk geometry: because of the exponential expansion of de Sitter vacua, the projection of a late time light-cone onto the fiducial hypersurface has subplanckian volume.

There is another reason why it would be nice to eliminate bulk geodesics from the definition of the measure. Ultimately, one expects that a fundamental description of the

is less well-defined than light-cone time and is not analogous to the UV/IR relation of AdS/CFT [26]. These concerns aside, it could be combined with the metric we construct on the future boundary. The resulting measure would be different from the one obtained here.

²In this respect, the definition of volumes explored here remains inferior, for now, since it is completely well-defined only for homogeneous bubble universes. However, for the reasons stated below, it may ultimately prove to be more fundamental.




Figure 8.2: Conformal diagram of de Sitter space with a single bubble nucleation. The parent de Sitter space is separated from the daughter universe by a domain wall, here approximated as a light-cone with zero initial radius (dashed line). There is a kink in the diagram where the domain wall meets future infinity. This diagram represents a portion of the Einstein static universe, but the Ricci scalar of the boundary metric is not constant.

multiverse will involve its boundary structure in some way. This was a key motivation for seeking a multiverse analogue of the UV/IR relation of AdS/CFT in the first place. The equivalence of the causal patch and light-cone time cut-off, along with their phenomenological successes, encourages us to take seriously the motivations behind the two measures. In particular, we would like to define the light-cone time cut-off, to the greatest extent possible, in terms of quantities that are intrinsic to the boundary.

In this chapter we will propose an intrinsic definition of volumes on the future boundary of the multiverse, and we will explore the resulting time foliation and measure. Like in AdS, we will render the boundary finite by a conformal rescaling of the bulk metric. However, we face two difficulties that have no direct analogue in AdS/CFT.

The boundary of the multiverse is naturally a fractal. In the bulk, bubbles of different vacua keep being produced at later and later times, leading to boundary features on arbitrarily small scales. If the bulk evolution exhibits attractor behavior (this is common to many global foliations of the multiverse, though the attractor itself depends on the foliation), then the boundary will exhibit self-similarity in the short-distance limit. Because of the abundance of short-distance structure, the boundary must be constructed first on large scales, and then refined to better and better resolution, corresponding to evolving the bulk to later and later times. The details of this procedure have not been carefully formulated and will concern us greatly.

The fractal structure leads to a number of difficulties for defining a metric on the future boundary:

- Generically, when a new bubble forms at late times in the bulk, one would like to

include this information on the boundary, for example by coloring a small disk on the boundary which is enclosed by this bubble. But the new vacuum will also change the metric in the bulk, and thus, the shape of the conformal diagram. In other words, generically, the boundary *metric* will keep changing as we increase the resolution. This is worse than the behavior of most fractals we are familiar with, which are defined by coloring points on a *fixed* background metric.

- A natural way of constructing a conformal diagram depicting the formation of a bubble universe in de Sitter space leads to a future boundary with a “sharp edge” (see Fig. 2), where the first derivative of the induced boundary metric is discontinuous. With more and more bubbles forming at late times, such edges would appear on arbitrarily small scales.

The shape of the boundary is defined only up to finite conformal transformations. Because the boundary metric is obtained from the bulk spacetime by a conformal rescaling, it is only defined up to a finite conformal factor. The shape of many conformal diagrams can be changed by a conformal transformation. The resulting diagram is just as legitimate as the original one, but its boundaries may have a different geometry.

- In the multiverse, this ambiguity in the boundary metric leads to an ambiguity in the bulk foliation. Unlike in AdS, which is asymptotically empty, this leads to a potential ambiguity in the measure. For example, a different choice of conformal factor may change the predicted probability distribution for the cosmological constant.

Outline We introduce the concept of light-cone time in Sec. 8.2.1, emphasizing that it requires a definition of the volume of any future light-cone. In Sec. 8.2.2, we propose that the volume should be defined as the volume enclosed by the light-cone on the future conformal boundary of the multiverse. Moreover, we propose choosing the conformal factor so that the boundary metric has constant scalar curvature,

$$R = \text{constant} , \tag{8.2}$$

where R is the Ricci scalar of the boundary metric³ A recent result in mathematics—the solution to the Yamabe problem—guarantees that a suitable conformal transformation can always be found and is essentially unique on the boundary. This proposal defines what we shall refer to as “new light-cone time”. In Sec. 9.3, we apply our proposal to a simplified landscape model. We approximate bubble universes as homogeneous open FRW cosmologies, we assume that they all have nonzero cosmological constant, and we neglect bubble collisions. Despite these simplifications, all of the difficulties listed above arise in this model, and we show that our proposal succeeds in addressing them. We construct the conformal diagram iteratively, making sure that the boundary condition $R = \text{const.}$

³Garriga and Vilenkin [74] defined the metric on future infinity by foliating the spacetime by surfaces whose induced metric has $R = 0$. In general spacetimes this is different from our proposal, and bulk $R = 0$ surfaces do not always exist. However, in the special case we focus on of homogeneous de Sitter bubbles, their boundary metric is the same as ours.

is satisfied each time a new bubble universe is included in the bulk. We find that this procedure leads to a fixed, everywhere smooth boundary metric that can be taken to be a unit three-sphere.

The boundary metric picked out by the condition $R = \text{const.}$ defines new light-cone time. In Sec. 8.3, we present some of the most important properties of this time foliation. In Sec. 10.2, we derive rate equations that govern the distribution of vacua. We emphasize the boundary viewpoint, in which bubble universes correspond to topological disks on the boundary three-sphere. In Sec. 8.3.2, we find the solution of the rate equations. The late time attractor behavior in the bulk corresponds to a universal ultraviolet scaling of the distribution of disks on the boundary. In Sec. 8.3.3, we derive the crucial expression that underlies the probability measure: the number of events of arbitrary type, as a function of light-cone time.

In Sec. 8.4, we analyze the probability measure. In Sec. 8.4.2, we show that new light-cone time yields the same measure as old light-cone time, i.e., both cut-offs predict the same relative probabilities for any two types of events in our simplified landscape model. In particular, this implies that our “new” light-cone time shares the phenomenological successes of the old one, and of the causal patch measure dual to it.

In Sec. 8.5, we discuss how our approach may extend to the general case, where inhomogeneities, collisions, and vacua with $\Lambda = 0$ are included. We also consider the (likely) possibility that the landscape contains vacua of different dimensionality. In defining a unique boundary metric, several difficulties arise in addition to the ones listed above, and a more general method of implementing Eq. (8.2) is needed. We discuss what phenomenological properties may be expected of the resulting measure. In particular, we expect that in the context of inhomogeneous universes, a boundary definition of light-cone volume will address a problem pointed out by Phillips and Albrecht [144].

8.2 New light-cone time

In this section, we define new light-cone time, and we construct the surfaces of constant light-cone time in a simple multiverse.

8.2.1 Probabilities from light-cone time

Given a time foliation of the multiverse, the relative probability of events of type A and B (e.g., two different outcomes of an experiment) can be defined by

$$\frac{p_A}{p_B} = \lim_{t \rightarrow \infty} \frac{N_A(t)}{N_B(t)}, \quad (8.3)$$

where $N_A(t)$ is the number of times an event of type A has occurred prior to the time t . This measure depends strongly on the choice of t . In this chapter, we focus exclusively on the case where t is of the form

$$t \equiv -\frac{1}{3} \log \frac{V(E)}{4\pi/3}. \quad (8.4)$$

Here, $V(E)$ is the (suitably defined) volume of the causal future of E , $I^+(E)$. Any time variable of this form will be referred to as light-cone time.

To complete the definition of light-cone time, and thus of the measure, one must define $V(E)$. One such definition, which results in what we shall refer to as *old light-cone time*, was given in Ref. [26]: consider a family of geodesics orthogonal to a fiducial hypersurface in the bulk, and let $V(E)$ be the volume occupied on the fiducial hypersurface by the geodesics that eventually enter $I^+(E)$. In this chapter we will explore a different definition, *new light-cone time*.

Often, we will find it convenient to work instead with the variable

$$\eta = e^{-t} = \left(\frac{V(E)}{4\pi/3} \right)^{1/3}, \quad (8.5)$$

which is naturally interpreted as a boundary distance scale. Since hypersurfaces of constant η are hypersurfaces of constant t , they define the same bulk foliation, and thus, the same measure. In the bulk, η decreases towards the future and vanishes on the boundary; t increases towards the future and diverges at the boundary.

8.2.2 A gauge choice for the conformal boundary

Physical distances diverge in the asymptotic future of an eternally inflating multi-verse, except inside black holes and vacua with negative cosmological constant, where they approach zero. Such behavior is found in many other spacetimes, such as (Anti-)de Sitter space or the Schwarzschild solution. However, in many cases a boundary of finite volume can be defined by a conformal transformation.

To a physical spacetime M , with metric $g_{\mu\nu}$, we associate an unphysical spacetime \tilde{M} with metric $\tilde{g}_{\mu\nu}$ which satisfies a number of conditions. There must exist a conformal isometry $\psi : M \rightarrow \psi[M] \subset \tilde{M}$ such that

$$\tilde{g}_{\mu\nu} = \Omega^2(x) \psi^* g_{\mu\nu} = e^{2\phi(x)} \psi^* g_{\mu\nu} \quad (8.6)$$

in $\psi[M]$. Note that Ω should be nowhere-vanishing on $\psi[M]$, and we demand that Ω be sufficiently smooth (say, C^3 [95]) in $\psi[M]$ and extend to a continuous function on the closure, $\overline{\psi[M]}$. We will refer to both Ω and $\phi = \log \Omega$ as the conformal factor. Hereafter we will identify M with its image $\psi[M]$ (called “the bulk”), and refer to the physical ($g_{\mu\nu}$) and unphysical ($\tilde{g}_{\mu\nu} = \Omega^2 g_{\mu\nu}$) metrics defined on M , eliminating ψ from the notation.

As a set, the boundary of M is defined as those points in the closure of $M \subset \tilde{M}$ which are not contained in M itself: $\partial M = \overline{M} - M$. By the “boundary metric”, G_{ab} , we mean the unphysical metric induced on ∂M viewed as a subset of \tilde{M} . The key property we will require of the unphysical spacetime is that for any $E \subset M$, the volume of $I^+(E) \cap \partial M$ be finite in the boundary metric. (Below we will define $V(E)$ to be this volume.) Note that for those cases where the physical volumes and distances diverge as one approaches the boundary, this implies that the conformal factor Ω approaches zero on ∂M , or that ϕ diverges. We will return in Sec. 8.5 to the question of which bulk metrics can be conformally rescaled so that the boundary metric is finite and nonsingular.

There is an ambiguity in the boundary metric which is related to an ambiguity in choosing the conformal factor that makes the rescaled bulk spacetime finite. Consider an additional Weyl rescaling with a conformal factor that is *bounded* on the boundary:

$$\tilde{G}_{ab} = e^{2\tilde{\phi}(x)} G_{ab} , \quad (8.7)$$

with $|\phi(x^\mu)| < K$ for some real number K . This will produce a new unphysical spacetime whose boundary still has finite volume, and which we could have obtained directly from the original spacetime via the conformal factor $\phi + \tilde{\phi}$. So the requirement of finiteness is not sufficient to fix the boundary metric completely.

We propose to fix the ambiguity by demanding that the conformal transformation yield a boundary metric of constant scalar curvature:

$$R = \text{constant} . \quad (8.8)$$

Here, R is the Ricci scalar computed in the boundary metric, not the bulk Ricci scalar restricted to the boundary. The value of the constant is arbitrary. (It can be changed by an overall rescaling of the unphysical spacetime; this shifts light-cone time by a constant but leaves the foliation, and thus the probability measure, invariant.) At a naive level of counting degrees of freedom, these conditions fix the metric.

At a more refined level, the question of whether an arbitrary metric on a smooth, compact manifold can be brought to a metric with constant scalar curvature is a difficult problem in mathematics known as the Yamabe problem. The Yamabe problem has been solved in the affirmative [156], [117]. For a smooth compact Riemannian manifold (M_n, G') of dimension $n \geq 3$, there exists a metric G conformal to G' such that the Ricci scalar of G is constant. The problem of finding the appropriate conformal transformation $G = \phi^{4/(n-2)} G'$ amounts to solving the differential equation

$$\frac{4(n-1)}{n-2} \nabla^2 \phi + R' \phi = R \phi^{\frac{n+2}{n-2}} , \quad (8.9)$$

with R a constant. This is nontrivial since the solution ϕ must be smooth and strictly positive, and must exist globally.

Yamabe [187] attempted to solve the problem by noting that (8.9) can be written as the Euler-Lagrange equation for a certain functional of the metric. This functional is closely related to the average scalar curvature of the metric over the manifold. Since the minimum occurs for the metric with constant curvature, the solution focuses on showing that the minimum of this functional is realized. Yamabe's original proof was later shown to be valid only in certain cases [176]. The proof was extended to other cases in [3] and completed in 1984 by Richard Schoen [156]. A unified and self-contained proof of the Yamabe problem is given in [117].

Having shown that a constant curvature metric exists, can we be sure that it is unique? Generically it is [2]: there is an open and dense set U in the space of conformal classes of metrics such that each element $[G] \in U$ has a unique unit volume metric with constant scalar curvature. In simple cases with symmetries, however, there may be a few-parameter family of solutions to the Yamabe problem. For example, the round three-sphere

has an ambiguity given by the global conformal group $SO(4,1)$. However, this ambiguity can only be used to fix the locations of a few points. We will ultimately be interested in the behavior on the shortest scales where this ambiguity has no effect.

We can now define $V(E)$ in Eq. (8.4) as follows. Let $L(E)$ be the portion of the boundary in the causal future of E , i.e., the intersection of $I^+(E)$ with the boundary. Let $V(E)$ be the volume of $L(E)$, measured using the metric G_{ab} obtained by the (essentially unique) conformal transformation that achieves constant Ricci scalar on the boundary:

$$V(E) \equiv \int_{L(E)} d^3y \sqrt{G}, \quad (8.10)$$

where G is the determinant of G_{ab} . With this choice of V , Eq. (8.4) defines the *new light-cone time*, t , and Eq. (8.3) defines a probability measure for the multiverse, the new light-cone time cut-off.

8.2.3 The multiverse in Ricci gauge

In this section, we carry out the construction of future infinity with $R = \text{const}$, subject to the following approximations and assumptions:

- All metastable de Sitter vacua are long-lived: $\kappa_\alpha \ll 1$, where $\kappa_\alpha = \sum \kappa_{i\alpha}$ is the total dimensionless decay rate of vacuum α , and $\kappa_{i\alpha} = \Gamma_{i\alpha} H_\alpha^{-4}$, where $\Gamma_{i\alpha}$ is the decay rate from the de Sitter vacuum α to the arbitrary vacuum i per unit four volume, and $H_\alpha = (3/\Lambda_\alpha)^{1/2}$ is the Hubble constant associated with the asymptotic de Sitter regime of vacuum α .⁴
- Bubble collisions are neglected. This is equivalent to setting all decay rates to zero in the causal past of a bubble.
- Each bubble is a homogeneous, isotropic, open universe; the only exception are bubbles nucleating inside bubbles, which spontaneously break this symmetry. On a large scale, this assumption essentially follows from the previous two approximations, since a large suppression of the decay implies a large suppression of fluctuations that break the $SO(3,1)$ symmetry of the Coleman-de Luccia geometry. But on a small scale, this assumption implies that we suppress any structure formation that results from small initial perturbations; we treat the bubble universe as completely homogeneous and isotropic.
- All bubbles have the same spacetime dimension, $D = 3 + 1$.
- No vacua with $\Lambda = 0$ are produced. This ensures that the boundary contains no null portions or “hats”, where the boundary metric would be degenerate.

The above assumptions allow us to take an iterative approach to the construction of the conformal diagram, while satisfying the gauge condition (8.2). We will now describe this construction step by step.

⁴We denote vacua with $\Lambda > 0$ by indices α, β, \dots and vacua with $\Lambda \leq 0$ by indices m, n, \dots . If no particular sign of Λ is implied, vacua are denoted by i, j, \dots

The universe begins in some metastable vacuum (it will not matter which one). Our initial step is to construct a conformal diagram satisfying $R = \text{const.}$ for this de Sitter vacuum, *as if it were completely stable*. We will refer to this spacetime as the zero-bubble multiverse. The metric is

$$ds_0^2 = \frac{-d\eta^2 + d\Omega_3^2}{H_0^2 \sin^2 \eta}, \quad (8.11)$$

where $d\Omega_3^2 = d\xi^2 + \sin^2 \xi (d\theta^2 + \sin^2 \theta d\phi^2)$ is the metric on the unit three-sphere. (Since an infinitely old metastable de Sitter space has zero probability, let us restrict to $\eta \leq \pi/2$,⁵ as if the universe came into being at the time $\eta = \pi/2$.) Multiplying this metric by the conformal factor $e^{2\phi} = H_0^2 \sin^2 \eta$, one obtains a conformally rescaled metric

$$d\tilde{s}_0^2 = -d\eta^2 + d\Omega_3^2, \quad (8.12)$$

which can be viewed as a portion of an Einstein static universe [95]. The future boundary corresponds to the points with $\eta = 0$, which were not part of the physical spacetime. Its induced metric is that of a round three-sphere with unit radius. Thus, it satisfies the gauge condition we impose, that the three-dimensional Ricci scalar be constant.

The second step is to construct light-cone time in the 1-bubble multiverse. Consider the causal future I^+ of an arbitrary event E with coordinates $(\eta, \xi, \theta, \phi)$ in the physical spacetime. We note that conformal transformations map light-cones to light-cones, and that the unphysical metric, Eq. (8.12), is spatially homogeneous. Therefore, the volume V occupied by $I^+(E)$ on the future boundary can only depend on the η coordinate of the event E , and not on its spatial coordinates; it is given by

$$V = \int_0^\eta 4\pi \sin^2 \eta d\eta. \quad (8.13)$$

The light-cone time is $t \equiv -\frac{1}{3} \log \frac{V}{4\pi/3}$. At late times, $V = \frac{4\pi}{3} \eta^3$, so $t = -\log \eta$.

The next step is to consider the first bubble that nucleates inside this vacuum, i.e., the nucleation event with the earliest light-cone time. (The time and place of this nucleation is of course governed by quantum chance; as we add more and more bubbles, we shall ensure that on average, their nucleation rate per unit four-volume conforms to the rates $\Gamma_{\alpha m}$ dictated by the landscape potential.) We then replace the causal future of the nucleation event with the new bubble, i.e., with an open FRW universe, which may have positive or negative cosmological constant, and may contain matter and radiation. Aside from this one nucleation event, we treat both the parent and the daughter vacuum as completely stable, so we do not modify the spacetime in any other way. We thus obtain a physical spacetime consisting of a parent de Sitter vacuum and one bubble universe with a different vacuum, which we refer to as the one bubble multiverse. Now we construct a new conformal diagram for this spacetime, as follows (see Fig. 8.2.3):

The conformal diagram is left unchanged outside the future light-cone of the nucleation event, where the physical spacetime is also unchanged. That is, we shall use the same conformal factor as before for this region.

⁵In naming the time coordinate in Eq. (8.11) η , we are anticipating the result below that $\eta = \text{constant}$ will define hypersurfaces of constant light-cone time in accordance with the definition of η in Eq. (8.5). Since this definition requires η to take positive values which *decrease* towards the future boundary, we shall take η to have a positive range in Eq. (8.11). Thus, $\eta \geq \pi/2$ corresponds to times *after* $\eta = \pi/2$.




Figure 8.3: (a) The parent de Sitter space with the future of the nucleation event removed and (b) the bubble universe are shown as separate conformal diagrams which are each portions of the Einstein static universe. After an additional conformal transformation of the bubble universe (shaded triangle), the diagrams can be smoothly matched along the domain wall. The resulting diagram (c) has a round S^3 as future infinity but is no longer a portion of the Einstein static universe.

The future of the nucleation event is an open FRW universe, which either crunches (if $\Lambda < 0$) or becomes asymptotically de Sitter in the future (if $\Lambda > 0$). We show in Appendix 8.6 that this spacetime can be conformally mapped to a finite unphysical spacetime with the following properties:

- The conformal factor is smooth (C^n , with n arbitrarily large) in the physical spacetime and continuous when extended to the future boundary.
- The induced metric along the future light-cone of the nucleation event agrees with the induced metric on the same hypersurface in the outside portion of the diagram, which was left unchanged.
- The future boundary is identical to the piece of the future boundary that was removed from the old diagram.

(Note that this portion of the unphysical spacetime will not be a portion of the Einstein static universe, nor is there any reason to demand that it should be.)

Because of the above properties, we are able to combine the remainder of the old Penrose diagram with the new portion covering the interior of the bubble. This results in a single Penrose diagram which is smooth everywhere⁶ in the physical spacetime, and whose future boundary is a round three-sphere of unit radius. Thus, it satisfies the condition we have imposed, that the Ricci scalar of the boundary metric be constant.

We can now simply iterate this procedure. Complete the following sequence of operations, starting with $i = 0$ (or really with $i = 1$ since we have already gone through one iteration, above):

⁶There is enough freedom in the choice of the conformal factor in the bulk to ensure that the conformal factor is not only continuous but arbitrarily differentiable at the seam where the FRW bubble and the old de Sitter parent are matched, while maintaining the round three-sphere metric of the future boundary. Thick-wall bubbles and non-zero initial bubble radii can similarly be accommodated.

1. Using the conformal diagram you have constructed, with a future boundary metric satisfying $R = \text{const}$, construct light-cone time for the i -bubble multiverse.
2. Identify the earliest bubble (i.e., the nucleation event with the smallest light-cone time) that has not already been included, and replace the future of the nucleation event with an FRW universe, which will be treated as completely stable in this step. The resulting spacetime is the $i + 1$ -bubble multiverse.
3. Construct a new conformal diagram with constant Ricci scalar on the future boundary. Since we have assumed that decay rates are small, decays will happen in asymptotically de Sitter regions, where the construction given above for $i = 0$ can be applied.
4. Increase i by 1 and go back to step 1.

Note that our construction is cumulative: every time a new bubble is included, the conformal factor is modified only inside this new bubble. Moreover, the geometry at future infinity remains a unit round three-sphere throughout the process; the only effect of each step is that some part of future infinity is now in a different vacuum. On the boundary, this allows us to think of bubble nucleation as the insertion of a disk of a corresponding size. We will exploit this fact in the next section, when we describe eternal inflation and its attractor behavior from the boundary point of view.

8.3 Properties of new light-cone time

In this section, we will derive and solve the equation governing the distribution of vacua as a function of light-cone time, and we will find an expression for the number of events of type A as a function of light-cone time.

8.3.1 Rate equation

We will work in the approximation of small decay rates, so that we can neglect bubble nucleations that happen during the early, non-vacuum-dominated phase of each FRW universe. Moreover, we will neglect the small fraction of bubble nucleations that would lead to collisions with existing bubbles. In the spirit of our new definition of light-cone time, however, we will express the rate equations in terms of variables that can be thought of as living on the future boundary: the volume V_α taken up by bubbles of type α on the boundary, and the variable $\eta = \exp(-t)$, which has a bulk interpretation as a time variable but also a boundary interpretation as a distance scale.

Consider a bubble of a de Sitter vacuum α that nucleated at some late time $\eta_{\text{muc}} \ll 1$, so that its size on future infinity is $V_\alpha^{(1)} = \frac{4\pi}{3} \eta_{\text{muc}}^3$. Now consider some later time $\eta/\eta_{\text{muc}} \ll 1$ inside the bubble. The metric is approximately

$$ds^2 = \frac{H_\alpha^{-2}}{\eta^2} (-d\eta^2 + d\vec{x}^2) , \quad (8.14)$$

where the comoving coordinate \mathbf{x} ranges over a volume $V_\alpha^{(1)}$. Here we have neglected the fact that at any finite η , the bubble will not have expanded to its full asymptotic comoving

volume, and that the bubble is not exactly de Sitter but may contain matter and radiation. Both of these corrections decrease like powers of η/η_{nuc} ; in particular, most of the physical volume at small η/η_{nuc} will be empty de Sitter space.

The above metric is valid at zeroth order in the decay rate expansion: we have neglected all decays that might have already taken place in the bubble by the time η . We will now include decays at first order. The proper three-volume of the bubble at the time η is $H_\alpha^{-3}V_\alpha^{(1)}/\eta^3$. One Hubble volume occupies proper volume $\frac{4\pi}{3}H_\alpha^{-3}$, so that the number of Hubble volumes is

$$n_\alpha^{(1)} = \frac{V_\alpha^{(1)}}{\frac{4\pi}{3}\eta^3} = \frac{\eta_{\text{nuc}}^3}{\eta^3} . \quad (8.15)$$

We note the relation $H_\alpha dt_{\text{prop}} = dt = -d\eta/\eta$ between the proper time, t_{prop} , light-cone time, t , and η . We also recall the definition that $\kappa_{i\alpha}$ is the rate at which i -bubbles are nucleated inside the α -vacuum, per Hubble volume and Hubble time. In a single bubble of type α , therefore, the total number of bubbles of type i that are produced during a time $d\eta$ is

$$dN_{i\alpha}^{(1)} = \kappa_{i\alpha} n_\alpha^{(1)} dt = \frac{-d\eta}{\frac{4\pi}{3}\eta^4} \kappa_{i\alpha} V_\alpha^{(1)} . \quad (8.16)$$

The total number of bubbles of type i produced inside α -bubbles during the time $d\eta$ is obtained by summing over all α -bubbles:

$$dN_{i\alpha} = \frac{-d\eta}{\frac{4\pi}{3}\eta^4} \kappa_{i\alpha} V_\alpha(\eta) , \quad (8.17)$$

where $V_\alpha(\eta)$ is the total volume taken up at future infinity by α bubbles nucleated prior to η .

Now consider the total volume dV_i^+ taken up at future infinity by all the i -bubbles produced (in *any* vacuum) during the time $d\eta$. Since we are neglecting bubble collisions, this is equal to the number of bubbles produced, times the volume taken up by each bubble:

$$dV_i^+ = \frac{4\pi}{3}\eta^3 \sum_\alpha dN_{i\alpha} . \quad (8.18)$$

Though some vacua with $\Lambda \leq 0$ can decay, this effect is not important and will also be neglected. Thus the sum runs only over de Sitter vacua, whereas i denotes a vacuum of any type.

As a result of decays into other vacua, the total boundary volume taken up by the de Sitter vacuum α decreases during the time $d\eta$, by a volume $\frac{4\pi}{3}\eta^3$ per decay:

$$dV_\alpha^- = -\frac{4\pi}{3}\eta^3 \sum_i dN_{i\alpha} . \quad (8.19)$$

Focussing now on de Sitter vacua alone, we can find that the total rate of change of the boundary volume occupied by α -bubbles, $dV_\alpha = dV_\alpha^+ + dV_\alpha^-$, by combining the previous three equations:

$$dV_\alpha = \frac{-d\eta}{\eta} \sum_\beta M_{\alpha\beta} V_\beta , \quad (8.20)$$

where

$$M_{\alpha\beta} \equiv \kappa_{\alpha\beta} - \kappa_{\alpha}\delta_{\alpha\beta} . \quad (8.21)$$

This is the rate equation for the volume occupied by metastable de Sitter vacua *on the boundary*. We derived this equation using bulk dynamics, which we understand at the semi-classical level. But we have expressed it in terms of variables that are defined on the boundary: V_{α} is a boundary volume and η is a boundary scale. In this form, the rate equation can be naturally interpreted in terms of the boundary dynamics we seek. In effect, we have derived a procedure for constructing the boundary fractal, working from the largest bubbles to the smallest. Begin with the boundary sphere all in one “color”, corresponding to the initial metastable de Sitter vacuum in the bulk. Now use the bubble distribution (8.16) to include new bubbles, one by one, corresponding to disks of different colors. The disks are smaller and smaller (of radius η) as we progress to smaller scales η . Just as there will be a finite number of bubbles prior to the time t in the bulk, there will be a finite number of disks on the boundary for any UV cut-off η . Their distribution will obey Eq. (8.20).

Eq. (8.16) specifies how many i -disks of a given size should be inserted in an α -disk, but we also should specify *where* they should be inserted. Since the nucleation rate is homogeneous in the bulk metric, Eq. (8.14), new bubbles be inserted with equal probability in any infinitesimal volume occupied by the α -disk. We have excluded bubble collisions; this can be incorporated from the boundary point of view by forbidding the addition of any new disk whose boundary would overlap with an existing disk. This amounts to excluding a zone of width η neighboring all existing disk boundaries. As we discussed in the previous section, the solution to the Yamabe problem on the sphere is not unique, and we still have the freedom to act with the group of global conformal transformations $SO(4, 1)$. However, this has no effect on Eq. (8.16), which is conformally invariant.

This completes our discussion of the boundary interpretation of the rate equation. The boundary process we have described should emerge from a more general boundary theory that remains to be discovered, and one would expect Eq. (8.20) to constrain the construction of such a theory. Meanwhile, in order to compare our result to the rate equation for “old” light-cone time, we would like to rewrite it using bulk variables. In terms of light-cone time, Eq. (8.20) becomes

$$\frac{dV_{\alpha}}{dt} = \sum_{\beta} M_{\alpha\beta} V_{\beta} . \quad (8.22)$$

We can use Eq. (8.15) to eliminate V_{α} , the volume taken up by α -disks on the boundary with UV-cutoff $\eta = \exp(-t)$, in favor of the number of horizon volumes of type α in the bulk at the time t :

$$V_{\alpha}(t) = n_{\alpha}(t) \frac{4\pi}{3} \exp(-3t) . \quad (8.23)$$

Thus, we obtain the bulk rate equation

$$\frac{dn_{\alpha}}{dt} = (3 - \kappa_{\alpha})n_{\alpha} + \sum_{\beta} \kappa_{\alpha\beta} n_{\beta} . \quad (8.24)$$

This is identical to the rate equation for old light-cone time, Eq. (36) in Ref. [44]. Thus, at this level of approximation, the two definitions of light-cone time yield precisely the same rate equations.

8.3.2 Attractor solution

The solutions to the above rate equation exhibit attractor behavior at late times ($\eta \rightarrow 0$):⁷

$$V_\alpha(\eta) = \check{V}_\alpha \eta^q + O(\eta^{\bar{q}}) ; \quad (8.25)$$

or, in terms of bulk variables,

$$n_\alpha(t) = \check{n}_\alpha e^{(3-q)t} + O(e^{(3-\bar{q})t}) . \quad (8.26)$$

Here $-q$ is the largest eigenvalue of the matrix $M_{\alpha\beta}$, \check{V}_α is the corresponding eigenvector, and $\check{n}_\alpha \equiv V_\alpha / \frac{4\pi}{3}$; $-\bar{q}$ is the second-largest eigenvalue, and $0 < q < \bar{q}$ [76, 44].

At short distances on the boundary (late times in the bulk), the distribution of disks (Hubble volumes) is governed by the “dominant eigenvector”, \check{V}_α , which can be thought of as the linear combination of de Sitter vacua that decays most slowly. In a realistic landscape, this eigenvector will have almost exclusive support in a single vacuum, which we shall denote by a star:

$$\check{V}_\alpha \approx v \delta_{\alpha^*} . \quad (8.27)$$

This dominant vacuum is the slowest-decaying de Sitter vacuum of the landscape [159]. The normalization, v , of the eigenvector depends on the infrared boundary configuration (initial conditions in the bulk).

8.3.3 Event counting

So far we have discussed only de Sitter vacua; and even in asymptotically de Sitter bubbles, we have neglected any initial, transitory period during which the bubble universe might have been dominated by, say, matter, radiation, or slow-roll inflation. This was sufficient for deriving the rate equation and the asymptotic attractor distribution for de Sitter vacua. As we now turn to the question of counting events in the multiverse, we must include all bubbles and regimes within them.

Let A be some type of event (e.g., a supernova occurs, or the microwave background temperature is found to be between 2.5 K and 3 K by some experiment). Let $N_A(\eta)$ be the number of times an event of type A has occurred in the bulk prior to the cut-off η . (On the boundary, we expect that this is the number of times it has been encoded in modes of size greater than η . For now, we will use the bulk definition since we know how to compute it.)

To compute $N_A(\eta)$, we substitute the asymptotic distribution of de Sitter vacua, Eq. (8.25), into our result for the number of bubbles of type i produced per unit time $d\eta$,

⁷See Eq. (37) of Ref. [76], who analyzed its solutions. Although our Eq. (8.22) takes the same mathematical form, it should be noted that it is for a different physical variable: the boundary volume fraction occupied by α -disks, which by Eq. (8.23) corresponds to a fraction of Hubble volumes in the bulk. (In Ref. [76], the relevant variable was the fraction of proper bulk volume occupied by vacuum α .)

Eq. (8.17):

$$dN_i = \frac{-d\eta}{\frac{4\pi}{3}\eta^{4-q}} \sum_{\alpha} \kappa_{i\alpha} \check{V}_{\alpha} ; \quad (8.28)$$

then we sum over all bubbles and nucleation times:

$$N_A(\eta) = \sum_i \int_{\eta_0}^{i\eta} \left(\frac{dN_A}{dN_i} \right)_{\eta/\eta_n} \left(\frac{dN_i}{d\eta} \right)_{\eta_n} d\eta_n , \quad (8.29)$$

where dN_A/dN_i is the expected number of events of type A that have occurred by the time η in a single bubble of type i nucleated at the time η_n . The lower limit of integration, η_0 , is an arbitrary early-time cut-off (an infrared cut-off on the boundary) that will drop out of the limit in Eq. (8.3).

At late times, all bubbles of type i are statistically equivalent, because they are produced locally in an empty de Sitter region. Therefore, dN_A/dN_i depends only on the ratio

$$\zeta \equiv \eta/\eta_n . \quad (8.30)$$

To avoid overcounting, the integral should run only over a single bubble of vacuum i , excluding regions of other vacua nucleated inside the i bubble; this restriction is denoted by index i appearing on the upper left of the integration symbol.

Combining the above results, and changing the integration variable from η_n to ζ ,

$$N_A(\eta) = \frac{1}{\frac{4\pi}{3}\eta^{3-q}} \sum_i \sum_{\alpha} N_{Ai} \kappa_{i\alpha} \check{V}_{\alpha} . \quad (8.31)$$

Here we have assumed for simplicity that A cannot occur in the dominant vacuum, $*$ (otherwise, see Ref. [44]); and we have defined

$$N_{Ai} \equiv \int_0^{i\eta} \left(\frac{dN_A}{dN_i} \right)_{\zeta} \zeta^{2-q} d\zeta . \quad (8.32)$$

The lower limit of integration in Eq. (8.32) should strictly be $\zeta = \eta/\eta_0$, so this result is valid only at late times ($\eta \rightarrow 0$), but this is the only regime relevant for computing relative probabilities.

The quantity $\left(\frac{dN_A}{dN_i} \right)_{\zeta}$ depends on the details of the bubble universe i and on how the cut-off surface η intersects this bubble. It can be written as

$$\left(\frac{dN_A}{dN_i} \right)_{\zeta} = \int_{\zeta}^{i\eta} d\zeta' \int d^3\mathbf{x} \sqrt{g} \rho_{Ai}(\zeta', \mathbf{x}) . \quad (8.33)$$

where ρ_{Ai} be the density of events of type A per unit four-volume in a bubble i , and g is the determinant of the metric. The integral is over the interior of a single bubble of type i , from its nucleation up to a light-cone time a factor ζ after nucleation. Substituting into Eq. (8.32) and exchanging the order of integration, we obtain an alternative expression for N_{Ai} :

$$N_{Ai} = \int_0^{i\eta} d\zeta' \int d^3\mathbf{x} \sqrt{g} \rho_{Ai}(\zeta', \mathbf{x}) \int_0^{\zeta'} d\zeta \zeta^{2-q} \quad (8.34)$$

$$= \frac{1}{3-q} \int d^4x \sqrt{g} \rho_{Ai}(x) \zeta(x)^{3-q} , \quad (8.35)$$

	use boundary		use geodesics
Global measures:	New Light Cone Time	\Leftrightarrow	Old Light Cone Time
	\Uparrow	\Downarrow	
Local measures:	New Causal Patch		Old Causal Patch

Table 8.1: Equivalences between measures. The new light-cone time cut-off is equivalent to probabilities computed from three other measure prescriptions (double arrows). This implies that all four measures shown on this table are equivalent in the approximation described at the beginning of Sec. 2.3. (See Ref. [44] for a more general proof of the equivalence between the old light-cone time cut-off and the old causal patch measure.)

where the integral is over the entire four-volume of a single i -bubble, with the exclusion of new bubbles nucleated inside it.

8.4 The probability measure

In Sec. 4.1 we derive the probability measure defined by the new light-cone time cut-off. In the following subsections, we show that it is equivalent to probabilities computed from three other measure prescriptions (see Table 8.4.1).

8.4.1 Probabilities from the new light-cone time cut-off

The relative probability of two events A and B is defined by

$$\frac{p_A}{p_B} = \lim_{\eta \rightarrow 0} \frac{N_A(\eta)}{N_B(\eta)}. \quad (8.36)$$

By substituting Eq. (8.31), we find that the probability of an event of type A is given by

$$p_A \propto \check{N}_A \equiv \sum_i \sum_\alpha N_{Ai} \kappa_{i\alpha} \check{V}_\alpha. \quad (8.37)$$

[We recall here that \check{V}_α is the eigenvector with largest eigenvalue of the matrix $M_{\alpha\beta}$ given in Eq. (8.21); $\kappa_{i\alpha}$ is the dimensionless nucleation rate of i -bubbles in the de Sitter vacuum α , and N_{Ai} can be computed from Eqs. (8.32) or (8.35); below we shall discover a simpler way of computing it, from Eqs. (8.39) and (8.44).] This expression is not normalized, but ratios can be taken to obtain relative probabilities.

We should go on to compute some probability distributions of interest. We should verify that catastrophic predictions such as Boltzmann brains, the youngness paradox (“Boltzmann babies”), and the Q -catastrophe are absent. If so, we should then move on to compute probability distributions over certain quantities of interest, such as the cosmological constant, the primordial density contrast, the spatial curvature, and the dark matter fraction, and we could verify that observed values are not unlikely under these distributions (say, in the central 2σ). Instead, we will now demonstrate that the new light-cone time cut-off is equivalent to other measures that have already been studied.

8.4.2 Equivalence to the old light-cone time cut-off

Probabilities computed from the old light-cone time cut-off are also of the general form (8.37) [44]. We will now show that they are in fact identical, under the assumptions made in the previous sections.

We consider each factor appearing inside the sum in Eq. (8.37) in turn. The decay rate $\kappa_{i\alpha}$ is a property of the landscape vacua unrelated to the definition of light-cone time. Moreover, we have shown that new light-cone time gives rise to the same rate equation and attractor solution as old light-cone time [26, 44], assuming that all metastable de Sitter vacua are long-lived. This implies that the eigenvector \check{V}_α is the same, whether old or new light-cone time is used. Therefore any difference between the two measures would have to come from the quantity N_{Ai} . However, if we approximate bubble interiors as homogeneous, isotropic, open universes, with metric of the form

$$ds^2 = -d\tau^2 + a(\tau)^2(d\chi^2 + \sinh^2 \chi d\Omega_2^2) . \quad (8.38)$$

then no such difference arises. At this level of approximation, old and new light-cone time define the same probability measure.

To demonstrate this, let us rewrite N_{Ai} , starting from Eq. (8.35). Homogeneity implies that ρ_{Ai} depends only on the FRW time, τ . Therefore, we can write

$$N_{Ai} = \frac{1}{3-q} \int_0^\infty d\tau \sigma_{Ai}(\tau) \mathcal{V}_c(\tau) , \quad (8.39)$$

where $\sigma_{Ai}(\tau) \equiv \rho_{Ai}(\tau)a(\tau)^3$ is the density of events of type A per FRW time and per unit comoving volume, and

$$\mathcal{V}_c(\tau) \equiv \int d\tilde{\Omega}_3 \zeta(\tau, \chi)^{3-q} . \quad (8.40)$$

Strictly, the integral should exclude comoving regions that have already decayed. But since the decay probability is homogeneous, we may absorb these excisions into a decrease of the comoving density σ_{Ai} and let the integral run over the entire unit hyperboloid, $d\tilde{\Omega}_3 = 4\pi \sinh^2 \chi d\chi$. The integral converges because the factor ζ^{3-q} rapidly vanishes at large χ . Thus, $\mathcal{V}_c(\tau)$ can be thought of as an effective comoving volume whose events contribute to N_{Ai} at the time τ .

Since q is the decay rate of the longest-lived vacuum, we can assume that it is exponentially small in a realistic landscape, so it can be neglected in Eq. (8.40). Moreover, from the definition of light-cone time in terms of volumes V on the future boundary, it follows that

$$\zeta^3 = e^{-(t-t_n)} = V(t)/V(t_n) . \quad (8.41)$$

Note that these equalities can be taken to hold by definition, *no matter how the volume on the future boundary is defined*. We will now abandon the particular definition made in Sec. 8.2.2 and consider a general metric $G_{\mu\nu}(\mathbf{y})$ on the future boundary of the multiverse. We shall find that \mathcal{V}_c , and thus N_{Ai} , is independent of this choice of metric if the bubble is homogeneous.

Upon substitution of Eq. (8.41), the integrand in Eq. (8.40) is the boundary volume of the future light-cone $L(\tau, \chi)$ of each point, which we may in turn write as an integral

over the region of the boundary enclosed by the light-cone:

$$\mathcal{V}_c(\tau) = V(t_n)^{-1} \int d\tilde{\Omega}_3 \int_{L(\tau,\chi)} d^3y \sqrt{G} \quad (8.42)$$

$$= V(t_n)^{-1} \int_{L(0,\chi)} d^3y \sqrt{G} \int_{CP(y,\tau)} d\tilde{\Omega}_3 . \quad (8.43)$$

In the second line, we exchanged the order of integration. The ranges of integration are the crucial point here. The union of the future light-cones from all points on the hyperboloid covers precisely the entire disk corresponding to the bubble universe at future infinity. Hence, the outer integral ranges over this disk, which we have written as the region enclosed by the future light-cone of the bubble nucleation event, $L(0,\chi)$. But each boundary point is covered by many future light-cones of bulk points. Correspondingly, the inner integral ranges over all points on the FRW slice whose future light-cone includes the boundary point y . But this is simply the set of points in the intersection of the $\tau = \text{const.}$ slice with the causal past of y , i.e., the portion of the constant FRW time slice that lies in the causal patch whose tip is at y . By homogeneity, its volume is independent of y and will be denoted $V_c^{\text{CP}}(\tau)$. Thus, the two integrals factorize. Since $\int_{L(0,\chi)} d^3y \sqrt{G} = V(t_n)$ by definition, we obtain

$$\mathcal{V}_c(\tau) = V_c^{\text{CP}}(\tau) . \quad (8.44)$$

With this result, Eq. (8.39) becomes manifestly independent of the definition of boundary volume. Thus, the quantity N_{Ai} will be the same for any type of “light-cone time”, if the bubbles of type i are homogeneous FRW universes.

This is a remarkably general result, so let us state it very clearly. From Eq. (8.32), N_{Ai} would appear to depend on the definition of ζ , which in turn depends on the definition of the scale η associated with a bulk point E as $V(E) = \frac{4\pi}{3}\eta^3$, where V is the volume taken up by the causal future of E on the future boundary. This volume, of course, can be defined in different ways: for example, in Ref. [26] it was defined by projecting onto a fiducial bulk hypersurface, whereas in the present chapter it was defined in terms of the induced metric of the boundary in a conformal frame with $R = \text{const.}$ These different definitions of V generically do lead to different values of ζ , η , and light-cone time t at the event E . What we have shown is that they cannot lead to any difference in N_{Ai} .

In the homogeneous FRW approximation, then, different definitions of light-cone time can only lead to a different probability measure if they lead to different rate equations. This can certainly happen in principle. However, we have shown earlier that the two definitions given so far—“old” and “new” light-cone time—yield the same rate equations. We conclude that they also lead to the same probability measure, if bubble universes are approximated as homogeneous open FRW universes.

8.4.3 Equivalence to the new causal patch measure

Consider an ensemble of causal patches whose tips are distributed over the future boundary of the multiverse, at constant density δ with respect to the boundary metric defined in Sec. 8.2.2. The probability for an event of type A , according to the causal patch measure, is proportional to the ensemble average of the number of times an event of type

A occurs inside a causal patch, $\langle N_A \rangle^{\text{CP}}$. Let η_0 be an early-time cut-off. Then

$$p_A^{\text{CP}} \propto \langle N_A \rangle^{\text{CP}} \propto \int_0^{\eta_0} d\eta \frac{dN_A}{d\eta} \eta^3 . \quad (8.45)$$

The quantity $d\eta \frac{dN_A}{d\eta}$ is the number of events that happen during the interval $d\eta$. Each of these events is contained in $\frac{4\pi}{3}\eta^3\delta$ causal diamonds, since $\frac{4\pi}{3}\eta^3$ is the volume of the causal future of an event at time η , and the causal diamonds that contain a given event are precisely those which have tips in the event's future [44].

Suppose that η_0 is small enough to lie deep inside the asymptotic attractor regime, where Eq. (8.31) holds. Then the above integral can be evaluated, and we find

$$p_A^{\text{CP}} \propto \tilde{N}_A . \quad (8.46)$$

Comparison with Eq. (8.37) reveals that the causal patch measure gives the same relative probabilities as the new light-cone time cut-off, if initial conditions for the causal patch are chosen in the attractor regime.

This result is similar to the global-local equivalence proven for the old light-cone time cut-off [44]. The local dual in that case was also an ensemble of causal patches beginning on a bulk hypersurface with vacuum distribution in the attractor regime, Eq. (8.26). However, the ensemble itself was selected by erecting geodesics at fixed density per unit Hubble volume on the initial hypersurface. By contrast, the causal patch ensemble we derived above from the new light-cone time is defined in terms of a uniform distribution of *tips* of causal patches per unit boundary volume, not of the starting points of their generating geodesics on a bulk hypersurface.

In fact, the concept of a generating geodesic of a causal patch appears to be unnecessary in the “new” causal patch measure. Instead, the patches are defined more simply as the past of points on the future boundary, or TIPS. This simplification is the local analogue of the elimination, in our definition of new light-cone time, of the family of geodesics that was needed to define old light-cone time. We regard these simplifications as a significant formal advantage of the boundary viewpoint.

8.4.4 Equivalence to the old causal patch measure

We turn to the diagonal arrow in the duality square of Table 8.4.1, the equivalence of the new light-cone time cut-off to the “old” causal patch measure, defined by erecting a family of geodesics orthogonal to a late time attractor hypersurface and constructing the causal patch of each geodesic, i.e., the past of its endpoint. This equivalence follows by combining two other arrows: the equivalence of the old and new light-cone time cut-offs (Sec. 8.4.2), and that of the old light-cone time cut-off and the (old) causal patch measure [44]. (This latter proof, unlike any of the arguments in this section, requires no simplifying assumptions such as homogeneity of the bubbles.) In the interest of a self-contained presentation, we will now present a shortcut that directly establishes the equivalence.

A geodesic starting in a vacuum α with probability \check{V}_α will enter vacuum i an expected number $\sum_\alpha \kappa_{i\alpha} \check{V}_\alpha$ of times [43]. (Since vacua are unlikely to be entered twice along the same geodesic, this expectation value can be thought of as an unnormalized

probability of entering i .) By the assumed homogeneity, all causal patches with tips in the same FRW bubble have statistically identical content, which, by Eq. (8.39), is given by $(3-q)N_{Ai}$. Thus, the probability of an event of type A , according to the the “old” causal patch measure, is

$$p_A^{\text{old CP}} \propto \sum_i \sum_\alpha N_{Ai\kappa_{i\alpha}} \check{V}_\alpha , \quad (8.47)$$

which agrees with the probability computed from the new light-cone time cut-off, Eq. (8.37).

8.5 The general case

At first sight, it seems that our prescription is well-defined for any bulk spacetime. Simply find a conformal transformation which makes the bulk spacetime finite, and then use the constant scalar curvature condition to fix the ambiguity in the boundary metric. The Yamabe theorem would seem to guarantee that this can always be done. But this process is not as simple as it sounds once we go beyond the approximation of homogeneous FRW universes stated in Sec. 9.3.

In this section, we will discuss the additional challenges that arise in the general case, how they might be addressed, and what they may imply for the phenomenology of new light-cone time. We first discuss perturbative inhomogeneities which complicate the story but do not obviously lead to a pathology. Then we discuss situations where the boundary metric is singular. These situations require a generalization of our prescription which we outline, but whose details we leave for the future. Finally, we argue that the probabilities derived from new and old light-cone time will differ in the presence of inhomogeneities, with new light-cone time being favored phenomenologically.

8.5.1 Perturbative inhomogeneities

No bubble universe will be exactly homogeneous and isotropic. Gravitational waves will perturb the metric and break spherical symmetry. Modes will leave the horizon during slow roll inflation and become frozen. We expect that such perturbations will affect the future boundary. Indeed, any metric of the form

$$ds^2 = -dt^2 + e^{2Ht} a_{ij}(x^k) dx^i dx^j , \quad (8.48)$$

is a solution to Einstein’s equation with positive cosmological constant in the limit $t \rightarrow \infty$, where a_{ij} is an *arbitrary* 3-metric that depends only on the spatial coordinates x^k .⁸ (This result is attributed to Starobinski in Ref. [19].) In terms of conformal time, the metric is

$$ds^2 = \frac{1}{\eta^2} (-d\eta^2 + a_{ij} dx^i dx^j) , \quad (8.49)$$

and the future boundary is the surface $\eta = 0$. Rescaling by the conformal factor $\Omega = \eta$, we obtain a (preliminary) boundary metric

$$ds^2 = a_{ij} dx^i dx^j . \quad (8.50)$$

⁸There are global constraints at finite time that may restrict the form of the boundary metric. Though it seems to us implausible, we cannot rule out that such restrictions may prevent the appearance of smaller and smaller geometric features on the boundary.

Recall that a_{ij} is arbitrary, so any metric is possible on future infinity. If Eq. (8.48) describes the late time limit of a de Sitter bubble universe, a_{ij} will need to be further conformally rescaled to finite volume and matched to the outside portion of the boundary. The condition $R = \text{const.}$ can then be achieved by a third conformal transformation. But because not every metric a_{ij} is conformal to a portion of the round S_3 , in general the future boundary will not be a round S_3 , and its geometry will contain features on the scale of the future light-cone of any perturbation.

Bubble nucleations occur at arbitrarily late times, corresponding to small sizes on the boundary. As a result, the boundary metric is perturbed on arbitrarily small scales. The boundary *geometry* is a fractal. It is important to understand the distinction we are drawing here to the homogeneous case. In this case, too, a fractal develops on the future boundary, but that fractal describes the disks corresponding to different vacua distributed over a fixed background metric (a round S^3). This is similar to the kind of fractal we are familiar with. What we are finding in the general case is that in addition, the geometry itself becomes a fractal. It exhibits self-similar features such as curvature on arbitrarily small scales.

By itself, the fractal curvature presents no serious obstacle. It just forces us to construct the background geometry of the boundary the same way we constructed its “matter content” (say, the disks representing different vacua): by starting in the infrared and refining it step by step as we move to short distances.

However, there are additional problems. Suppose we would like to include some bulk perturbation, such as the formation of a new bubble universe, on the boundary. In the simplified multiverse we studied in Sec. 9.3, the conformal transformation that achieved constant Ricci scalar acted nontrivially only in the future of the nucleation event; it did not change the conformal transformation outside the future light-cone. But the Yamabe problem cannot generally be solved subject to such strong boundary conditions. Thus, we expect that for general bulk perturbations, the boundary metric with $R = \text{constant}$ will be modified in regions that are outside the lightcone of the perturbation. So the nucleation of an inhomogeneous bubble generically changes the lightcone time of every event, even events out of causal contact with the nucleation.

In the homogeneous approximation, if we are given some finite region of the bulk spacetime, we have enough information to compute the surfaces of constant light-cone time in that region because adding bubbles outside the region, even in its future, will not change the size of future lightcones. If we had a computer simulating the bulk evolution, we could write a simple algorithm for generating the entire bulk spacetime underneath the cutoff surface. But in the inhomogeneous case, we do not have a simple prescription for the size of a bulk event on the boundary. Events in the future, or even in causally disconnected regions, will change the boundary metric, and thus could change the size of “old” future light-cones on the boundary. So given a computer which can simulate the bulk evolution, we do not have an algorithm for computing the bulk spacetime up to the cutoff, because it is unclear which events will turn out to occur at times after the cutoff.

This poses a challenge to the iterative construction of light-cone time, in Sec. 9.3. This procedure relied on moving forward along an existing light-cone time foliation, the i -bubble multiverse, and including the first new bubble nucleation reached in this manner




Figure 8.4: Conformal diagram of de Sitter space containing a bubble universe with $\Lambda = 0$

to compute the $i + 1$ -bubble multiverse. In this algorithm, the light-cone time was never modified for events prior to the new bubble nucleation. Thus, we could be sure at each stage that we have a final version of the foliation up to the $i + 1$ -th bubble nucleation. As a result, a UV cut-off on the boundary had a simple bulk translation: bubble nucleations after a certain time should be ignored.

In fact, it is unclear whether a discrete iterative process remains a viable approximation scheme in the general case. We were able to focus on bubble nucleation events only, since by our assumptions no other events had the potential to modify the future boundary. In general, one expects that all spacetime events, or at least all events in regions that are not empty de Sitter space, will affect the future boundary. This is a continuous set.

While it is possible that some or all of the above effects are quantitatively negligible, they raise interesting conceptual issues. If we had a boundary description of the theory, presumably it would be well defined to go down to shorter and shorter scales. What is missing at this point is a simple bulk prescription once we go beyond the approximation of homogeneous FRW universes. The fact that the boundary geometry becomes fractal once one includes all of the expected bulk dynamics means that defining volumes on the boundary is trickier. What is needed is a way to construct the UV-cutoff version of the boundary geometry, which is not a fractal. This can perhaps be done by beginning with the homogeneous approximation and then perturbing around it, but we have not attempted to do this.

8.5.2 Singularities in the boundary metric

There is a more severe pathology that can arise in constructing the boundary metric. What if the bulk spacetime does not lead to a boundary metric which is nonsingular and $3 + 1$ dimensional everywhere? There are several ways this can happen. Perhaps the simplest is the nucleation of a $\Lambda = 0$ bubble. Future infinity will now contain a “hat” as

shown in Fig. 8.5.2. The induced metric will be degenerate because the “hat” is a null cone. The null portion of the boundary metric is

$$ds^2 = r^2 d\Omega_2^2, \quad (8.51)$$

which is not 3-dimensional. So a bulk spacetime that includes $\Lambda = 0$ asymptotic regions will have a singular boundary metric. It is then unclear how to define volumes on the boundary.

Another example of a bulk spacetime that leads to a singular boundary metric is a Schwarzschild black hole in de Sitter space. The metric for Schwarzschild-de Sitter in static coordinates is

$$ds^2 = -\left(1 - \frac{2m}{r} - \frac{r^2}{l^2}\right) dt^2 + \frac{1}{1 - \frac{2m}{r} - \frac{r^2}{l^2}} dr^2 + r^2 d\Omega_2^2. \quad (8.52)$$

If we continue the metric behind the cosmological horizon and let $r \rightarrow \infty$, the metric approaches pure de Sitter space. For large r , the metric is

$$ds^2 = -\frac{l^2}{r^2} dr^2 + r^2 \left(\frac{1}{l^2} dt^2 + d\Omega_2^2 \right). \quad (8.53)$$

The boundary is $r \rightarrow \infty$; conformally rescaling the factor of r^2 gives the boundary metric

$$ds^2 = \frac{1}{l^2} dt^2 + d\Omega_2^2. \quad (8.54)$$

The $r = \infty$ part of the future boundary is a spatial cylinder. This is conformal to a sphere, aside from two missing points at the poles.

What about the other part of future infinity, the black hole singularity? Here we take the opposite limit $r \rightarrow 0$ to obtain the metric

$$ds^2 = -\frac{r}{2m} dr^2 + \frac{2m}{r} \left(dt^2 + \frac{r^3}{2m} d\Omega_2^2 \right). \quad (8.55)$$

As $r \rightarrow 0$, the t direction blows up while the S_2 becomes small. Therefore, if we do the conventional conformal rescaling to keep the t direction finite, the S_2 will have zero size on the boundary, and the boundary metric will be one-dimensional,

$$ds^2 = dt^2. \quad (8.56)$$

So the boundary metric for a Schwarzschild-de Sitter black hole can be thought of as a sphere with a string coming out from the north pole. This is not a 3-dimensional manifold and it is not clear how to generalize our construction to this case.

Bubbles with negative cosmological constant end in a spacelike singularity. In the homogeneous approximation these big crunches actually do not deform the boundary away from the round S_3 . However, we know that the homogeneous approximation breaks down at late times. Instead of a homogeneous collapse, any small region near the singularity will be of the Kasner form, with two contracting and one expanding direction, much like the singularity of the Schwarzschild black hole discussed above. The boundary metric will be one-dimensional. The Kasner behavior is uncorrelated between different parts of the

singularity: which direction is expanding changes from place to place. Therefore we expect that the boundary geometry for a big crunch is something like a fishnet.

One might expect that the light-cone time cut-off (and by extension, the causal patch) is simply inapplicable to portions of the multiverse that are in the domain of dependence of future singularities or hats [26], since the analogy with AdS/CFT breaks down entirely in such regions. In this case, the measure may still give a good approximation to the correct probabilities in regions such as ours, which are in the eternally inflating domain (roughly, the de Sitter regions). Assuming that a smoothing procedure can be found along the lines of the previous subsection, one can ignore singularities while they are smaller than the UV cut-off and, if one accepts that the measure is meant to apply only to eternally inflating regions, one can stay away from the singularities once they are larger than the UV cut-off.

However, there is another way in which the boundary can have portions of different dimensionality. In a realistic landscape, one expects that there will be de Sitter vacua with dimension smaller or greater than $3 + 1$. For example, in a landscape that includes both 4-dimensional and 6-dimensional de Sitter vacua, some regions of the future boundary will be 3-dimensional while others will be 5-dimensional. We would need some way of generalizing the construction to boundaries with a variety of dimensionalities.

8.5.3 Inhomogeneities

From the previous subsections we conclude that two things are needed to make our prescription well-defined in a realistic multiverse. We must be able to allow the boundary to change dimensionality from place to place, and we need a procedure for smoothing the fractal nature of the boundary geometry. We will now assume that this can be achieved and turn to the phenomenological implications of giving up the assumption of homogeneity.

In Sec. 8.4.2 we established the equivalence of old and new light-cone time in the homogeneous approximation. It is precisely when homogeneity is broken that old and new light-cone time differ. We will find it convenient to analyze this difference in the local version of each measure, the old and new causal patch measures. A causal patch is always the same type of region: the past of a point on the future boundary. But the old and new causal patch measure consider different ensembles of causal patches, and probabilities are computed as an ensemble average.

In the old causal patch measure, the ensemble is defined by erecting timelike geodesics orthogonal to an initial hypersurface in the dominant vacuum, and following their evolution until they reach a crunch. The resulting patches will contain different decoherent histories. If a geodesic enters a bubble universe like ours, it quickly becomes comoving [33]. With probability presumably of order one, it will become bound to a galaxy or galaxy cluster. After the cosmological constant begins to dominate, at the time τ_Λ , it may well remain bound to this structure for a time much greater than τ_Λ , until it is either ejected by some process or falls into a black hole, or the structure disintegrates. As a consequence, it may be the case that typical causal patches in the old ensemble do *not* become empty at late times but contain a single bound structure or black hole [144]. This could have important consequences for some probability distributions.

For example, consider a class of observers that live at a FRW time τ_{obs} . (Nothing

else is assumed, e.g., one need not assume that they live in a vacuum like ours or are in any way similar to us [36].) What value of the cosmological constant are such observers likely to find? The prior probability distribution in the landscape favors large Λ . But for sufficiently large Λ , one has $\tau_\Lambda \ll \tau_{\text{obs}}$. In the homogeneous approximation, the number of observers inside the causal patch is then diluted by the de Sitter expansion as $\exp(-\sqrt{3\Lambda}\tau_{\text{obs}})$, so that the probability distribution peaks at $\Lambda \sim \tau_{\text{obs}}^{-2}$. Thus the causal patch predicts that observers are likely to find themselves at the onset of vacuum domination [37]. This prediction is both more general and more successful than the original prediction that $\Lambda \sim \tau_{\text{gal}}^{-2}$ [181], where τ_{gal} is the time of virialization; and it is more robust to the variation of other parameters such as the primordial density contrast.

In an inhomogeneous universe, however, the exponential decay is cut off when all but structures to which the generating geodesic was *not* bound have been expelled from the horizon. After this time, the mass inside the causal patch does not decrease rapidly. Of course, this still corresponds to a reduction in mass by a factor 10^{-11} in a universe like ours. It will not spoil the prediction unless the number of observers per galaxy during the era $\tau \gg \tau_\Lambda$ is larger than in the present era by the inverse factor, which may well fail to be the case for observers like us in a universe like ours.

But let us view the landscape as a whole and treat observers more abstractly. The effect of inhomogeneity on the old causal patch ensemble is to remove the pressure for observations to take place not too long after τ_Λ . This allows for a much more efficient use of free energy, by conversion into quanta with wavelength as large as the de Sitter horizon. If we model observers by entropy production as proposed in [24, 37], and if we assume that processes that make such efficient use of free energy do not have small prior probability, then we should be rather surprised that we are not among this class of observers [144].

The new causal patch measure appears to resolve this puzzle. Inhomogeneities such as galaxies and the black holes they contain will decay on a timescale that is power-law and thus negligible compared to typical vacuum decay timescales [56]. We expect, therefore, that they will leave little imprint on the future boundary of the multiverse, in the sense that the metric will not have strong features that favor causal patches containing structure at late times. The fraction of causal patches containing any structure at times $\tau \gg \tau_\Lambda$ will be $\exp(-\sqrt{3\Lambda}\tau_{\text{obs}})$, so the above analysis, which was valid only in the homogeneous approximation for the old causal patch measure, will always be valid in the new causal patch measure after ensemble-averaging.

8.6 Conformal factor for a homogeneous bubble universe

In this appendix we show how to find a conformal transformation which maps a universe with a bubble nucleation event into a finite unphysical spacetime with a round S^3 as future infinity. Our procedure is to begin with the full Penrose diagram of the parent de Sitter space, and then remove the portion which lies within the future light-cone of the nucleation event. The removed piece will be replaced with a new unphysical spacetime which is conformally equivalent to an open FRW universe with nonzero cosmological constant and has the following properties, outlined in Sec. 9.3:

- The conformal factor is smooth (C^n with n arbitrarily large) in the physical spacetime

and continuous when extended to the future boundary.

- The induced metric along the future light-cone of the nucleation event agrees with the induced metric on the same hypersurface in the outside portion of the diagram, which was left unchanged.
- The future boundary is identical to the piece of the future boundary that was removed from the old diagram.

(Note that the new portion of the unphysical spacetime will not be a portion of the Einstein static universe, nor is there any reason to demand that it should be.)

We will accomplish this mapping in three steps. First, we will show that any FRW universe with non-zero vacuum energy is conformally equivalent to a portion of de Sitter space. Second, we will show that this portion of de Sitter space can be conformally mapped to a portion of the Einstein static universe whose future boundary is identical to the piece of the future boundary removed from the parent de Sitter diagram. Third, we show that this portion of the Einstein static universe is conformally equivalent to a new unphysical spacetime which satisfies all of the above properties (and is not a portion of the Einstein static universe).

An open FRW universe, such as the one following the nucleation event, has the metric

$$ds_{\text{FRW}}^2 = a^2(T)(-dT^2 + dH_3^2) , \quad (8.57)$$

where T is the conformal time, a is the scale factor, and

$$dH_3^2 = d\chi^2 + \sinh^2 \chi d\Omega_2^2 . \quad (8.58)$$

All bubble nucleation events are followed by a period of curvature domination in the bubble universe, and hence conformal time is unbounded toward the past. Furthermore, as long as $\Lambda \neq 0$, conformal time is finite toward the future. We are thus free to choose the range of T to be $-\infty < T \leq 0$ for every FRW universe under consideration. The open slicing of de Sitter space (with unit de Sitter length) is one special case of this metric:

$$a_{\text{dS}}(T) = \frac{-1}{\sinh T} . \quad (8.59)$$

Any other FRW universe is conformally equivalent to the de Sitter open slicing, and the conformal factor is given by

$$\Omega_1 = -\frac{1}{a(T) \sinh T} . \quad (8.60)$$

The second step is to map the open slicing of de Sitter space into a portion of the Einstein static universe whose future boundary is identical to the piece removed from the parent vacuum. Recall that the Einstein static universe has a metric given by Eq. 8.12,

$$d\tilde{s}_0^2 = -d\eta^2 + d\Omega_3^2 , \quad (8.61)$$

and that the Penrose diagram of the parent de Sitter space consists of the region $0 > \eta > -\pi$. We write the S^3 metric as

$$d\Omega_3^2 = d\xi^2 + \sin^2 \xi d\Omega_2^2 . \quad (8.62)$$

In these coordinates the nucleation event can be taken to be at $\eta = \eta_{\text{muc}}$ and $\xi = 0$, so that the portion of future infinity we need to reproduce in the bubble universe is the portion of the round S^3 with $0 < \xi < -\eta_{\text{muc}}$.⁹

The open slicing coordinates (T, χ) are not convenient for this task, so first we will change coordinates to $(\eta_{\text{in}}, \xi_{\text{in}})$, in terms of which the metric Eq. 8.57, with $a = a_{\text{dS}}$, looks like

$$ds^2 = \frac{1}{\sin^2 \eta_{\text{in}}} (-d\eta_{\text{in}}^2 + d\Omega_3^2) , \quad (8.63)$$

where we need to ensure that $0 < \xi_{\text{in}} < -\eta_{\text{muc}}$ when $\eta_{\text{in}} = 0$. Then we can simply rescale by $\sin \eta_{\text{in}}$ to finish the job. Now we turn to the task of finding the transformation $(T, \chi) \rightarrow (\eta_{\text{in}}, \xi_{\text{in}})$.

Let X_0, X_1, X_2, X_3, X_4 of be the coordinates of 4+1-dimensional Minkowski space, in which de Sitter space is the hyperboloid

$$-X_0^2 + \sum_{i=1}^3 X_i^2 + X_4^2 = 1 . \quad (8.64)$$

The relationship between the coordinates $(\eta_{\text{in}}, \xi_{\text{in}})$ and the X_μ is

$$\begin{aligned} X_0 &= \cot \eta_{\text{in}} \\ X_i &= -\frac{1}{\sin \eta_{\text{in}}} \sin \xi_{\text{in}} \hat{n}_i \\ X_4 &= -\frac{1}{\sin \eta_{\text{in}}} \cos \xi_{\text{in}}, \end{aligned} \quad (8.65)$$

where \hat{n}_i are unit vectors whose sum equals 1. We need to specify how the (χ, T) coordinates of Eq. 8.57 relate to the X_μ . The standard open slicing of de Sitter space is given by

$$\begin{aligned} X_0 &= -\frac{1}{\sinh T} \cosh \chi \\ X_i &= -\frac{1}{\sinh T} \sinh \chi \hat{n}_i \\ X_4 &= -\coth T , \end{aligned} \quad (8.66)$$

but this is not what we want to do. We have to remember that the nucleation event is at $\chi = 0, T = -\infty$, which would be equivalent to $\eta_{\text{in}} = -\pi/2$ and $\xi_{\text{in}} = 0$ if we used this prescription. To fix the problem, we use the 4+1-dimensional boost symmetry, which is an isometry of the de Sitter space, to move the nucleation event to another position. Let \tilde{X}_μ be given by

$$\tilde{X}_0 = \cosh \beta X_0 + \sinh \beta X_4 \quad (8.67)$$

$$\tilde{X}_4 = \sinh \beta X_0 + \cosh \beta X_4$$

$$\tilde{X}_i = X_i, \quad 1 \leq i \leq 3 , \quad (8.68)$$

⁹In this appendix only, we will find it more convenient for η to take negative values and increase towards the future boundary. The η defined in the main body of the chapter is $\eta_{\text{main}} = |\eta_{\text{appendix}}|$.

and then define the open coordinates (T, χ) by

$$\begin{aligned}\tilde{X}_0 &= -\frac{1}{\sinh T} \cosh \chi \\ \tilde{X}_i &= -\frac{1}{\sinh T} \sinh \chi \hat{n}_i \\ \tilde{X}_4 &= -\coth T .\end{aligned}\tag{8.69}$$

Now the nucleation event is at $\eta_{\text{in}} = \eta_{\text{nuc}}$ and $\xi_{\text{in}} = 0$, where $\cot \eta_{\text{nuc}} = -\sinh \beta$. We can write the relationship between (T, χ) and $(\eta_{\text{in}}, \xi_{\text{in}})$ directly in terms of η_{nuc} as

$$\begin{aligned}\frac{\sinh \chi}{\sinh T} &= \frac{\sin \xi_{\text{in}}}{\sin \eta_{\text{in}}} \\ -\cot \eta_{\text{in}} &= \frac{1}{\sin \eta_{\text{nuc}}} \frac{\cosh \chi}{\sinh T} + \cot \eta_{\text{nuc}} \coth T .\end{aligned}\tag{8.70}$$

Now the metric is in the form Eq. 8.63, and $0 < \xi_{\text{in}} < -\eta_{\text{nuc}}$ when $\eta_{\text{in}} = 0$. Hence we can conformally rescale by

$$\Omega_2 = \sin \eta_{\text{in}} ,\tag{8.71}$$

to map the bubble universe into a portion of the Einstein static universe with future boundary identical to the piece cut out of the old diagram.

The product $\Omega_1 \Omega_2$ succeeds in mapping our original FRW universe to a portion of the Einstein static universe with the correct future boundary. Now we will act with one final conformal rescaling, Ω_3 , which must change the induced unphysical metric along the future light-cone of the nucleation event to match the one in the old de Sitter diagram, thus ensuring that the total conformal transformation is continuous. We must also demand that $\Omega_3 = 1$ on the future boundary $\eta_{\text{in}} = 0$, since we have already fixed that part of the unphysical spacetime.

The induced metric on the future light-cone of the nucleation event has a similar form in both the unphysical bubble coordinates $(\eta_{\text{in}}, \xi_{\text{in}})$ and the unphysical parent coordinates, which we will now denote as $(\eta_{\text{out}}, \xi_{\text{out}})$. In both cases it merely comes from the Einstein static universe metric:

$$ds_{\text{light-cone}}^2 = \sin^2 \xi_{\text{in/out}} d\Omega_2^2 ,\tag{8.72}$$

where $\xi_{\text{in/out}} = \eta_{\text{in/out}} - \eta_{\text{nuc}}$. Hence the conformal factor evaluated along the light-cone is

$$\Omega_3 = \frac{\sin \xi_{\text{out}}}{\sin \xi_{\text{in}}} .\tag{8.73}$$

We just need to find out how ξ_{out} and ξ_{in} are related. We can do this by demanding that the induced *physical* metrics also be identical on the light-cone.

The induced physical metric on the light-cone from the parent de Sitter space is found by restricting Eq. 8.11 to the relevant surface:

$$ds_0^2 \rightarrow \frac{\sin^2 \xi_{\text{out}}}{H_0^2 \sin^2 \eta_{\text{out}}} d\Omega_2^2 ,\tag{8.74}$$

where again $\eta_{\text{out}} - \xi_{\text{out}} = \eta_{\text{nuc}}$.

In terms of the χ and T coordinates inside the bubble, the relevant light-cone is given by $\chi \rightarrow \infty$, $T \rightarrow -\infty$, $\chi + T = \text{const}$. In terms of η_{in} and ξ_{in} , the light-cone is given by $\eta_{\text{in}} - \xi_{\text{in}} = \eta_{\text{nuc}}$. One can see that a fixed value of $\chi + T$ is equivalent to a fixed value of ξ_{in} . Since $a(T) \rightarrow Ce^{-T}$ for some constant C as $T \rightarrow -\infty$, we have from Eq. 8.57 that

$$ds_{\text{FRW}}^2 \rightarrow Ce^{2(\chi+T)} d\Omega_2^2 = C \frac{\sin^2 \xi_{\text{in}}}{\sin^2 \eta_{\text{in}}} d\Omega_2^2 . \quad (8.75)$$

We must determine the constant C , which is related to our convention that $T = 0$ on future infinity. This constant is easily determined by considering the particular light-cone given by $\chi + T = 0$. This is the event horizon of the FRW universe, and its area is given by A_{EH} , which depends on the detailed form of $a(T)$. So $C = A_{\text{EH}}/4\pi$. Finally we have our relationship between the in-coordinates and out-coordinates on the bubble wall hypersurface:

$$\sqrt{\frac{A_{\text{EH}}}{4\pi} \frac{\sin(\eta_{\text{in}} - \eta_{\text{nuc}})}{\sin \eta_{\text{in}}}} = \frac{1}{H_0} \frac{\sin(\eta_{\text{out}} - \eta_{\text{nuc}})}{\sin \eta_{\text{out}}} . \quad (8.76)$$

Notice that η_{in} and η_{out} coincide at the nucleation event (both equal to η_{nuc}) and at future infinity (both equal to zero).

We are free to extend Ω_3 into the rest of the bubble universe in any continuous way we please, so long as it restricts to 1 on future infinity and $\sin \xi_{\text{out}}/\sin \xi_{\text{in}}$ on the boundary light-cone. One might worry about the fact that Ω_3 is multivalued at the point where the bubble wall meets future infinity. This is not a concern, though, because the only function which need be well defined is the product

$$\Omega = \Omega_1 \Omega_2 \Omega_3 = -\frac{\sin \eta_{\text{in}}}{a(T) \sinh T} \Omega_3 . \quad (8.77)$$

In this formula, $\eta_{\text{in}} = \eta_{\text{in}}(\chi, T)$ as determined by Eqs 8.70 . From here it is easy to see that as we approach the point $\chi = \infty$ along any slice of fixed T , including both $T = 0$ (future infinity) and $T = -\infty$ (domain wall), we arrive at $\Omega = 0$ so long as Ω_3 remains finite.

Thus we have accomplished our initial task. The function Ω_3 is very ambiguous, as we have only fixed its behavior at the boundaries of the spacetime. We can use this additional freedom to make Ω arbitrarily smooth at the interface between the parent and the bubble universe. Additionally, one could impose that Ω_3 be identically equal to 1 not just on the future boundary, but in the entire region $T > T_0$ for some particular T_0 . In particular, we can choose T_0 to be the time of vacuum domination in the bubble universe.

Chapter 9

Testing measures

9.1 Introduction

String theory appears to contain an enormous landscape of metastable vacua [41, 109], with a corresponding diversity of low-energy physics. The cosmological dynamics of this theory is eternal inflation. It generates a multiverse in which each vacuum is produced infinitely many times.

In a theory that predicts a large universe, it is natural to assume that the relative probability for two different outcomes of an experiment is the ratio of the expected number of times each outcome occurs. But in eternal inflation, every possible outcome happens infinitely many times. The relative abundance of two different observations is ambiguous until one defines a measure: a prescription for regulating the infinities of eternal inflation.

Weinberg's prediction [181] of the cosmological constant [151, 143] was a stunning success for this type of reasoning. In hindsight, however, it was based on a measure that was ill-suited for a landscape in which parameters other than Λ can vary. Moreover, the measure had severe phenomenological problems [27]. This spurred the development of more powerful measure proposals in recent years [125, 123, 71, 72, 70, 76, 179, 178, 24, 121, 122, 141, 73, 182, 183, 184, 126, 26, 28, 160]. Surprisingly, some of these measures do far more than to resolve the above shortcomings. As we shall see in this chapter, they obviate the need for Weinberg's assumption that observers require galaxies; and they help overcome the limitation of fixing all parameters but one to their observed values.

In this chapter we will analyze three different measure proposals. Each regulates the infinite multiverse by restricting attention to a finite portion. The causal patch measure [24] keeps the causal past of the future endpoint of a geodesic; it is equivalent to a global cutoff known as light-cone time [26, 44]. The fat geodesic measure [34] keeps a fixed physical volume surrounding the geodesic; in simple situations, it is equivalent to the global scale factor time cutoff [53]. We also introduce a new measure, which restricts to the interior of the apparent horizon surrounding the geodesic.

From little more than the *geometry* of these cutoffs, we are able to make remarkable progress in addressing cosmological coincidence and hierarchy problems. Using each measure, we will predict three time scales: the time when observations are made, t_{obs} , the time of vacuum energy domination, $t_{\Lambda} \equiv \sqrt{3/|\Lambda|}$, and the time of curvature domination,

t_c .¹ We work in the approximation that observations occur in nearly homogeneous FRW universes so that these time scales are well defined.

We will allow all vacuum parameters to vary simultaneously. Parameters for which we do not compute probabilities are marginalized. We will not restrict attention to vacua with specific features such as baryons or stars. We will make some weak, qualitative assumptions about the prior distribution of parameters in the landscape. We will assume that most observers are made of something that redshifts faster than curvature. (This includes all forms of radiation and matter but excludes, e.g., networks of domain walls.) But we will not impose any detailed anthropic requirements, such as the assumption that observers require galaxies or complex molecules; we will not even assume that they are correlated with entropy production [24, 37, 36]. Thus we obtain robust predictions that apply to essentially arbitrary observers in arbitrary vacua.

The probability distribution over all three variables can be decomposed into three factors, as we will explain further in Sec. 9.2:

$$\frac{d^3 p}{d \log t_{\text{obs}} d \log t_{\Lambda} d \log t_c} = \frac{d^2 \tilde{p}}{d \log t_{\Lambda} d \log t_c} \times M(\log t_{\text{obs}}, \log t_c, \log t_{\Lambda}) \times \alpha(\log t_{\text{obs}}, \log t_c, \log t_{\Lambda}) . \quad (9.1)$$

Here \tilde{p} is the probability density that a bubble with parameters $(\log t_{\Lambda}, \log t_c)$ is produced within the region defined by the cutoff. Its form can be estimated reliably enough for our purposes from our existing knowledge about the landscape. The factor $M(\log t_{\text{obs}}, \log t_{\Lambda}, \log t_c)$ is the mass inside the cutoff region at the time t_{obs} . This is completely determined by the geometry of the cutoff and the geometry of an FRW bubble with parameters $(\log t_{\Lambda}, \log t_c)$, so it can be computed unambiguously. The last factor, $\alpha(\log t_{\text{obs}}, \log t_c, \log t_{\Lambda})$, is the one we know the least about. It is the number of observations per unit mass per logarithmic time interval, averaged over all bubbles with the given values $(\log t_{\Lambda}, \log t_c)$.

A central insight exploited in this chapter is the following. In bubbles with positive cosmological constant, the calculable quantity M so strongly suppresses the probability in other regimes that in many cases we only need to know the form of α in the regime where observers live before vacuum energy or curvature become important, $t_{\text{obs}} \lesssim t_{\Lambda}, t_c$. Under very weak assumptions, α must be independent of t_{Λ} and t_c in this regime. This is because neither curvature nor vacuum energy play a dynamical role before observers form, so that neither can affect the number of observers per unit mass. Thus, for positive cosmological constant α is a function only of t_{obs} in the only regime of interest. The success of the measures in explaining the hierarchy and coincidence of the three timescales depends on the form of this function. We will that the causal patch and apparent horizon cutoff succeed

¹In regions where curvature will never dominate, such as our own vacuum, t_c is defined as the time when curvature would come to dominate if there were no vacuum energy. Since our observations are consistent with a flat universe, we can only place a lower bound on the observed value of t_c . We include t_c in our analysis because bubble universes in the multiverse are naturally very highly curved, so the absence of curvature requires an explanation. Moreover, we shall see that some measures select for high curvature in vacua with with negative cosmological constant. This effect is partly responsible for the problems encountered in this portion of the landscape.

well in predicting the three timescales already under very weak assumptions on α . The fat geodesic cutoff requires somewhat stronger assumptions.

For negative cosmological constant, however, the geometric factor M favors the regime where t_{obs} is not the shortest time scale. Thus, predictions depend crucially on understanding the form of α in this more complicated regime. For example, we need to know on average to what extent early curvature domination disrupts the formation of observers. What is clear from our analysis is that all three measures are in grave danger of predicting that most observers see negative Λ , in conflict with observation, unless the anthropic factor α takes on rather specific forms. Assuming that this is not the case, we must ask whether the measures can be modified to agree with observation. Both the causal patch measure and the fat geodesic measure are dual [34, 44] to global time cutoffs, the lightcone time and the scale factor time cutoff, respectively. These global cutoffs, in turn, are motivated by an analogy with the UV/IR relation in AdS/CFT [73]. But this analogy really only applies to positive Λ , so it is natural to suspect that the measure obtained from it is inapplicable to regions with $\Lambda \leq 0$ [26]. (Indeed, the causal patch does not eliminate infinities in $\Lambda = 0$ bubbles [24]. We do not consider such regions in this chapter.)

Outline and summary of results In Sec. 9.2, we will describe in detail our method for counting observations. We will derive an equation for the probability distribution over the three variables $(\log t_\Lambda, \log t_c, \log t_{\text{obs}})$. We will explain how simple qualitative assumptions about $\alpha(\log t_{\text{obs}}, \log t_c, \log t_\Lambda)$, the number of observations per unit mass per unit logarithmic time interval, allow us to compute probabilities very generally for all measures.

We work in the approximation that observations in the multiverse take place in negatively curved Friedmann-Robertson-Walker (FRW) universes. In Sec. 9.3, we will obtain solutions for their scale factor, in the approximation where the matter-, vacuum-, and possibly the curvature-dominated regime are widely separated in time.

In Secs. 9.4–9.6, we will compute the probability distribution over $(\log t_\Lambda, \log t_c, \log t_{\text{obs}})$, using three different measures. For each measure we consider separately the cases of positive and negative cosmological constant. As described above, the results for negative cosmological constant are problematic for all three measures. We now summarize our results for the case $\Lambda > 0$.

In Sec. 9.4, we show that the causal patch measure predicts the double coincidence $\log t_c \approx \log t_\Lambda \approx \log t_{\text{obs}}$. We find that the scale of all three parameters is related to the number of vacua in the landscape. This result is compatible with current estimates of the number of metastable string vacua. Such estimates are not sufficiently reliable to put our prediction to a conclusive test, but it is intriguing that the size of the landscape may be the origin of the hierarchies we observe in Nature (see also [147, 24, 25, 35, 39, 36]). We have previously reported the result for this subcase in more detail [30].

Unlike the causal patch, the new “apparent horizon measure” (Sec. 9.5) predicts the double coincidence $\log t_c \approx \log t_\Lambda \approx \log t_{\text{obs}}$ for any fixed value t_{obs} . When all parameters are allowed to scan, its predictions agree with those of the causal patch, with mild assumptions about the function α . The apparent horizon measure is significantly more involved than the causal patch: it depends on a whole geodesic, not just on its endpoint, and on metric information, rather than only causal structure. If our assumptions about α

are correct, this measure offers no phenomenological advantage over its simpler cousin and may not be worthy of further study.

The fat geodesic cutoff, like the apparent horizon cutoff, predicts the double coincidence $\log t_c \approx \log t_\Lambda \approx \log t_{\text{obs}}$ for any fixed value of t_{obs} . However, it favors small values of all three timescales unless either (a) an anthropic cutoff is imposed, or (b) it is assumed that far more observers form at later times, on average, than at early times, in vacua where curvature and vacuum energy are negligible at the time t_{obs} . (Qualitatively, the other two measures also require such an assumption, but quantitatively, a much weaker prior favoring late observers suffices.) In the latter case (b), the results of the previous two measures are reproduced, with all timescales set by the size of the landscape. In the former case (a), the fat geodesic would predict that observers should live at the earliest time compatible with the formation of any type of observers (in any kind of vacuum). It is difficult to see why this minimum value of t_{obs} would be so large as to bring this prediction into agreement with the observed value, $t_{\text{obs}} \sim 10^{61}$.

9.2 Counting Observations

In this section, we explain how we will compute the trivariate probability distribution over $(\log t_{\text{obs}}, \log t_c, \log t_\Lambda)$. We will clarify how we isolate geometric effects, which can be well-computed for each cutoff, from anthropic factors; and we explain why very few assumptions are needed about the anthropic factors once the geometric effects have been taken into account.

Imagine labelling every observation within the cutoff region by $(\log t_{\text{obs}}, \log t_c, \log t_\Lambda)$. We are interested in counting the number of observations as a function of these parameters. It is helpful to do this in two steps. First, we count bubbles, which are labelled by $(\log t_c, \log t_\Lambda)$ to get the ‘‘prior’’ probability distribution

$$\frac{d^2 \tilde{p}}{d \log t_\Lambda d \log t_c} . \quad (9.2)$$

This \tilde{p} is the probability density to nucleate a bubble with the given values of the parameters inside the cutoff region.

The next step is to count observations within the bubbles. A given bubble of vacuum i will have observations at a variety of FRW times. In the global description of eternal inflation, each bubble is infinite and contains an infinite number of observations if it contains any, but these local measures keep only a finite portion of the global spacetime and hence a finite number of observations. We parameterize the probability density for observations within a given bubble as

$$\frac{dN_i}{d \log t_{\text{obs}}} \sim M(\log t_{\text{obs}}, \log t_c, \log t_\Lambda) \alpha_i(\log t_{\text{obs}}) , \quad (9.3)$$

where M is the mass inside the cutoff region, and α_i is the number of observations per unit mass per logarithmic time interval inside a bubble of type i . In this decomposition, M contains all of the information about the cutoff procedure. For a given cutoff, M depends only on the three parameters of interest. Because we are considering geometric cutoffs,

the amount of mass that is retained inside the cutoff region does not depend on any other details of the vacuum i . On the other hand, α_i depends on details of the vacuum, such as whether observers form, and when they form; but it is independent of the cutoff procedure.

Since we are interested in analyzing the probability distribution over three variables, we now want to average α_i over the bubbles with a given $(\log t_\Lambda, \log t_c)$, to get the average number of observations per unit mass per logarithmic time α . With this decomposition, the full probability distribution over all three variables is

$$\frac{d^3 p}{d \log t_{\text{obs}} d \log t_\Lambda d \log t_c} = \frac{d^2 \tilde{p}}{d \log t_\Lambda d \log t_c} M(\log t_{\text{obs}}, \log t_c, \log t_\Lambda) \alpha(\log t_{\text{obs}}, \log t_c, \log t_\Lambda) . \quad (9.4)$$

To recap, \tilde{p} is the probability for the formation of a bubble with parameters $(\log t_\Lambda, \log t_c)$ inside the cutoff region; $M(\log t_{\text{obs}}, \log t_\Lambda, \log t_c)$ is the mass inside the cutoff region at the FRW time t_{obs} in a bubble with parameters $(\log t_\Lambda, \log t_c)$; and $\alpha(\log t_{\text{obs}}, \log t_c, \log t_\Lambda)$ is the number of observations per unit mass per logarithmic time interval, averaged over all bubbles with parameters $(\log t_\Lambda, \log t_c)$.

This is a useful decomposition because the mass M inside the cutoff region can be computed exactly, since it just depends on geometrical information. We will assume that $d^2 \tilde{p} / d \log t_\Lambda d \log t_c$ can be factorized into contribution from t_Λ and a contribution from t_c . Vacua with $\Lambda \sim 1$ can be excluded since they contain only a few bits of information in any causally connected region. In a large landscape, by Taylor expansion about $\Lambda = 0$, the prior for the cosmological constant is flat in Λ for $\Lambda \ll 1$, $d\tilde{p}/d\Lambda = \text{const}$. Transforming to the variable $\log t_\Lambda$, we thus have

$$\frac{d^2 \tilde{p}}{d \log t_\Lambda d \log t_c} \sim t_\Lambda^{-2} g(\log t_c) . \quad (9.5)$$

The factor $g(\log t_c)$ encodes the prior probability distribution over the time of curvature domination. We will assume that g decreases mildly, like an inverse power of $\log t_c$. (Assuming that slow-roll inflation is the dominant mechanism responsible for the delay of curvature domination, $\log t_c$ corresponds to the number of e -foldings. If g decreased more strongly, like an inverse power of t_c , then inflationary models would be too rare in the landscape to explain the observed flatness.) The detailed form of the prior distribution over $\log t_c$ will not be important for our results; any mild suppression of large $\log t_c$ will lead to similar results.

With these reasonable assumptions, the probability distribution becomes

$$\frac{d^3 p}{d \log t_{\text{obs}} d \log t_\Lambda d \log t_c} = t_\Lambda^{-2} M(\log t_{\text{obs}}, \log t_c, \log t_\Lambda) g(\log t_c) \alpha(\log t_{\text{obs}}, \log t_c, \log t_\Lambda) . \quad (9.6)$$

Because α depends on all three variables, it is very difficult to compute in general. However, it will turn out that we can make a lot of progress with a few simple assumptions about α . First, we assume that in the regime, $t_{\text{obs}} \ll t_\Lambda$, α is independent of t_Λ . Similarly, we assume that in the regime, $t_{\text{obs}} \ll t_c$, α is independent of t_c . By these assumptions, in the regime where the observer time is the shortest timescale, $t_{\text{obs}} \lesssim t_\Lambda, t_c$, the anthropic factor α will only depend on $\log t_{\text{obs}}$:

$$\alpha(\log t_{\text{obs}}, \log t_c, \log t_\Lambda) \approx \alpha(\log t_{\text{obs}}) \text{ for } t_{\text{obs}} \lesssim t_\Lambda, t_c . \quad (9.7)$$

These assumptions are very weak. Because curvature and Λ are not dynamically important before t_c and t_Λ , respectively, they cannot impact the formation of observers at such early times. One could imagine a correlation between the number of e -foldings and the time of observer formation even in the regime $t_{\text{obs}} \ll t_c$, for example if each one is tied to the supersymmetry breaking scale, but this seems highly contrived. In the absence of a compelling argument to the contrary, we will make the simplest assumption.

Second, we assume that when either $t_{\text{obs}} \gg t_c$ or $t_{\text{obs}} \gg t_\Lambda$, α is not enhanced compared to its value when t_{obs} is the shortest timescale. This is simply the statement that early curvature and vacuum domination does not help the formation of observers. This assumption, too, seems rather weak. With this in mind, let us for the time being overestimate the number of observers by declaring α to be completely independent of $\log t_\Lambda$ and $\log t_c$:

$$\alpha(\log t_{\text{obs}}, \log t_c, \log t_\Lambda) \approx \alpha(\log t_{\text{obs}}) . \quad (9.8)$$

This is almost certainly an overestimate of the number of observations in the regime where t_{obs} is not the shortest time scale. However, we will show that the predictions for positive cosmological constant are insensitive to this simplification because the geometrical factor M , which we can compute reliably, suppresses the contribution from this regime. We will return to a more realistic discussion of α when we are forced to, in analyzing negative cosmological constant.

With these assumptions and approximations, the three-variable probability distribution takes a more tractable form,

$$\frac{d^3 p}{d \log t_{\text{obs}} d \log t_\Lambda d \log t_c} = t_\Lambda^{-2} g(\log t_c) M(\log t_{\text{obs}}, \log t_\Lambda, \log t_c) \alpha(\log t_{\text{obs}}) . \quad (9.9)$$

This is the formula that we will analyze in the rest of the chapter. The only quantity that depends on all three variables is the mass, which we can compute reliably for each cutoff using the geometry of open bubble universes, to which we turn next.

9.3 Open FRW universes with cosmological constant

In this section, we find approximate solutions to the scale factor, for flat or negatively curved vacuum bubbles with positive or negative cosmological constant. The landscape of string theory contains a large number of vacua that form a “discretuum”, a dense spectrum of values of the cosmological constant [41]. These vacua are populated by Coleman-DeLuccia bubble nucleation in an eternally inflating spacetime, which produces open Friedmann-Robertson-Walker (FRW) universes [49]. Hence, we will be interested in the metric of FRW universes with open spatial geometry and nonzero cosmological constant Λ . The metric for an open FRW universe is

$$ds^2 = -dt^2 + a(t)^2 (d\chi^2 + \sinh^2 \chi d\Omega_2^2) . \quad (9.10)$$

The evolution of the scale factor is governed by the Friedmann equation:

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{t_c}{a^3} + \frac{1}{a^2} \pm \frac{1}{t_\Lambda^2} . \quad (9.11)$$

Here $t_\Lambda = \sqrt{3/|\Lambda|}$ is the timescale for vacuum domination, and t_c is the timescale for curvature domination. The term $\rho_m \sim t_c/a^3$ corresponds to the energy density of pressureless matter. (To consider radiation instead, we would include a term $\rho_{\text{rad}} \sim t_c^2/a^4$; this would not affect any of our results qualitatively.) The final term is the vacuum energy density, ρ_Λ ; the “+” sign applies when $\Lambda > 0$, and the “-” sign when $\Lambda < 0$.

We will now display approximate solutions for the scale factor as a function of FRW time t . There are four cases, which are differentiated by the sign of Λ and the relative size of t_c and t_Λ . We will compute all geometric quantities in the limit where the three time scales t , t_c , and t_Λ are well-separated, so that some terms on the right hand side of Eq. (9.11) can be neglected. In this limit we obtain piecewise solution for the scale factor. We will *not* ensure that these solutions are continuous and differentiable at the crossover times. This would clutter our equations, and it would not affect the probability distributions we compute later. Up to order-one factors, which we neglect, our formulas are applicable even in crossover regimes.

If $t_c \gg t_\Lambda$, curvature never comes to dominate. (One can still define t_c geometrically, or as the time when curvature would have dominated in a universe with $\Lambda = 0$.) In this limit the metric can be well approximated as that of a perfectly flat FRW universe, and so becomes independent of t_c . We implement this case by dropping the term t_c/a^3 in Eq. (9.11).

Positive cosmological constant We begin with the case $\Lambda > 0$ and $t_c \ll t_\Lambda$. By solving Eq. (9.11) piecewise, we find

$$a(t) \sim \begin{cases} t_c^{1/3} t^{2/3}, & t < t_c \\ t, & t_c < t < t_\Lambda \\ t_\Lambda e^{t/t_\Lambda - 1}, & t_\Lambda < t. \end{cases} \quad (9.12)$$

If $t_c \gg t_\Lambda$, there is no era of curvature domination, and the universe can be approximated as flat throughout. The scale factor takes the form

$$a(t) \sim \begin{cases} t_c^{1/3} t^{2/3}, & t < t_\Lambda \\ t_c^{1/3} t_\Lambda^{2/3} e^{t/t_\Lambda - 1}, & t_\Lambda < t. \end{cases} \quad (9.13)$$

Negative cosmological constant For $\Lambda < 0$, the scale factor reaches a maximum and then begins to decrease. The universe ultimately collapses at a time t_f , which is of order t_Λ :

$$t_f \approx \pi t_\Lambda. \quad (9.14)$$

The evolution is symmetric about the turnaround time, $t_f/2 \approx \pi t_\Lambda/2$.

Again, we consider the cases $t_\Lambda \gg t_c$ and $t_\Lambda \ll t_c$ separately. For $t_c \ll t_\Lambda$, the scale factor is

$$a(t) \sim \begin{cases} t_c^{1/3} t^{2/3}, & t < t_c \\ t_\Lambda \sin(t/t_\Lambda), & t_c < t < t'_c \\ t_c^{1/3} (t')^{2/3}, & t'_c < t. \end{cases} \quad (9.15)$$

We have defined $t' \equiv t_f - t$.

There is no era of curvature domination if $t_c \gtrsim t_f/2$. For $t_c \gg t_f/2$, we treat the universe as flat throughout, which yields the scale factor

$$a(t) \sim t_\Lambda^{2/3} t_c^{1/3} \sin^{2/3}(\pi t/t_f) , \quad (9.16)$$

where t_f here takes on a slightly different value compared to the curved case:

$$t_f = 2\pi t_\Lambda/3 . \quad (9.17)$$

At a similar level of approximation as we have made above, this solution can be approximated as

$$a(t) \sim \begin{cases} t_c^{1/3} t^{2/3} , & t < t_f/2 \\ t_c^{1/3} (t')^{2/3} , & t_f/2 < t . \end{cases} \quad (9.18)$$

where again $t' \equiv t_f - t$.

9.4 The causal patch cut-off

With all our tools lined up, we will now consider each measure in turn and derive the probability distribution. We will treat positive and negative values of the cosmological constant separately. After computing M , we will next calculate the bivariate probability distribution over $\log t_\Lambda$ and $\log t_c$, for fixed $\log t_{\text{obs}}$. In all cases this is a sharply peaked function of the two variables $\log t_\Lambda$ and $\log t_c$, so we make little error in neglecting the probability away from the peak region. Then we will go on to find the full distribution over all three parameters. In this section, we begin with the causal patch measure [24, 32], which restricts to the causal past of a point on the future boundary of spacetime. The causal patch measure is equivalent [44] to the light-cone time cutoff [26, 28], so we will not discuss the latter measure separately.

We may use boost symmetries to place the origin of the FRW bubble of interest at the center of the causal patch. The boundary of the causal patch is given by the past light-cone from the future end point of the comoving geodesic at the origin, $\chi = 0$:

$$\chi_{\text{CP}}(t) = \int_t^{t_f} \frac{dt'}{a(t')} . \quad (9.19)$$

If $\Lambda < 0$, t_f is the time of the big crunch (see Sec. 9.3). For long-lived metastable de Sitter vacua ($\Lambda > 0$), the causal patch coincides with the event horizon. It can be computed as if the de Sitter vacuum were eternal ($t_f \rightarrow \infty$), as the correction from late-time decay is negligible.

9.4.1 Positive cosmological constant

We begin with the case $\Lambda > 0$, $t_c < t_\Lambda$. Using Eq. (9.12) for $a(t)$, we find

$$\chi_{\text{CP}}(t_{\text{obs}}) \sim \begin{cases} 1 + \log(t_\Lambda/t_c) + 3[1 - (t_{\text{obs}}/t_c)^{1/3}] , & t_{\text{obs}} < t_c \\ 1 + \log(t_\Lambda/t_{\text{obs}}) , & t_c < t_{\text{obs}} < t_\Lambda \\ e^{-t_{\text{obs}}/t_\Lambda} , & t_\Lambda < t_{\text{obs}} . \end{cases} \quad (9.20)$$

The comoving volume inside a sphere of radius χ is $\pi(\sinh 2\chi_{\text{CP}} - 2\chi_{\text{CP}})$. We approximate this, dropping constant prefactors that will not be important, by χ^3 for $\chi \lesssim 1$, and by $e^{2\chi}$ for $\chi \gtrsim 1$:

$$V_{\text{CP}} \sim \begin{cases} \exp(2\chi_{\text{CP}}) , & t_{\text{obs}} < t_{\Lambda} \\ \chi_{\text{CP}}^3 , & t_{\Lambda} < t_{\text{obs}} . \end{cases} \quad (9.21)$$

The mass inside the causal patch is $M_{\text{CP}} = \rho a^3 V_{\text{CP}} = t_c V_{\text{CP}}$.

$$M_{\text{CP}} \sim \begin{cases} t_{\Lambda}^2/t_c , & t_{\text{obs}} < t_c < t_{\Lambda} & I \\ t_{\Lambda}^2 t_c/t_{\text{obs}}^2 , & t_c < t_{\text{obs}} < t_{\Lambda} & II \\ t_c e^{-3t_{\text{obs}}/t_{\Lambda}} , & t_c < t_{\Lambda} < t_{\text{obs}} & III \end{cases} \quad (9.22)$$

Next, we consider the case $\Lambda > 0$, $t_{\Lambda} < t_c$. The above calculations can be repeated for a flat universe, which is a good approximation for this case:

$$M_{\text{CP}} \sim \begin{cases} t_{\Lambda} , & t_{\text{obs}} < t_{\Lambda} & V \\ t_{\Lambda} e^{-3t_{\text{obs}}/t_{\Lambda}} , & t_{\Lambda} < t_{\text{obs}} & IV \end{cases} \quad (9.23)$$

The same result could be obtained simply by setting $t_c = t_{\Lambda}$ in (9.22).

The full probability distribution is given by multiplying the mass in the causal patch by the prior distribution and the number of observations per unit mass per unit time to give

$$\frac{d^3 p_{\text{CP}}}{d \log t_c d \log t_{\Lambda} d \log t_{\text{obs}}} \sim g\alpha \times \begin{cases} \frac{1}{t_c} , & t_{\text{obs}} < t_c < t_{\Lambda} & I \\ \frac{t_c}{t_{\text{obs}}^2} , & t_c < t_{\text{obs}} < t_{\Lambda} & II \\ \frac{t_c}{t_{\Lambda}^2} \exp\left(-\frac{3t_{\text{obs}}}{t_{\Lambda}}\right) , & t_c < t_{\Lambda} < t_{\text{obs}} & III \\ \frac{1}{t_{\Lambda}} \exp\left(-\frac{3t_{\text{obs}}}{t_{\Lambda}}\right) , & t_{\Lambda} < t_{\text{obs}}, t_c & IV \\ \frac{1}{t_{\Lambda}} , & t_{\text{obs}} < t_{\Lambda} < t_c & V \end{cases} \quad (9.24)$$

Recall that $g(\log t_c)$ is the prior distribution on the time of curvature domination, and α is the number of observations per unit mass per logarithmic time interval.

We will first analyze this probability distribution for fixed $\log t_{\text{obs}}$. As explained in the introduction, we will for the time being overestimate the number of observers in the regime where t_{obs} is not the shortest timescale by assuming that α is a function of $\log t_{\text{obs}}$ only. We will find that the overestimated observers do not dominate the probability distribution, so this is a good approximation. With these approximations, α is independent of $\log t_{\Lambda}$ and $\log t_c$ and can be ignored for fixed $\log t_{\text{obs}}$.

The probability distribution we have found for $\log t_c$ and $\log t_{\Lambda}$ is a function of powers of t_c and t_{Λ} , i.e., exponential in the logarithms. Therefore the distribution will be dominated by its maximum. A useful way of determining the location of the maximum is to follow the gradient flow generated by the probability distribution. In the language of Ref. [88, 35], this is a multiverse force pushing $\log t_{\Lambda}$ and $\log t_c$ to preferred values. We




Figure 9.1: The probability distribution over the timescales of curvature and vacuum domination at fixed observer timescale $\log t_{\text{obs}}$, before the prior distribution over $\log t_c$ and the finiteness of the landscape are taken into account. The arrows indicate directions of increasing probability. For $\Lambda > 0$ (a), the distribution is peaked along the degenerate half-lines forming the boundary between regions I and II and the boundary between regions IV and V. For $\Lambda < 0$ (b), the probability distribution exhibits a runaway toward the small t_c , large t_Λ regime of region II. The shaded region is excluded because $t_{\text{obs}} > t_f = \pi t_\Lambda$ is unphysical.

could use our formulas to determine the precise direction of the multiverse force, but this would be difficult to represent graphically, and it is not necessary for the purpose of finding the maximum. Instead, we shall indicate only whether each of the two variables prefers to increase or decrease (or neither), by displaying horizontal, vertical, or diagonal arrows in the $(\log t_c, \log t_\Lambda)$ plane (Fig. 9.1). We ignore the prior $g(\log t_c)$ for now since it is not exponential.

We consider each region in (9.24) in turn. In region I, $t_{\text{obs}} < t_c < t_\Lambda$, the probability is proportional to t_c^{-1} . Hence, there is a pressure toward smaller t_c and no pressure on t_Λ . This is shown by a left-pointing horizontal arrow in region I of the figure. In region II ($t_c < t_{\text{obs}} < t_\Lambda$), the probability is proportional to $t_c t_{\text{obs}}^{-2}$. This pushes toward larger $\log t_c$ and is neutral with respect to $\log t_\Lambda$. (Recall that we are holding $\log t_{\text{obs}}$ fixed for now.) In region III ($t_c < t_\Lambda < t_{\text{obs}}$), the probability goes like $t_c t_\Lambda^{-2} e^{-3t_{\text{obs}}/t_\Lambda}$. Since the exponential dominates, the force goes toward larger $\log t_\Lambda$; $\log t_c$ is pushed up as well. In region IV ($t_\Lambda < t_{\text{obs}}$, $t_\Lambda < t_c$) the exponential again dominates, giving a pressure toward large $\log t_\Lambda$. In region V ($t_{\text{obs}} < t_\Lambda < t_c$), the distribution is proportional to t_Λ^{-1} , giving a pressure toward small $\log t_\Lambda$. The dependence of the probability on $\log t_c$ lies entirely in the prior $g(\log t_c)$ in these last two regions because the universe is approximately flat and hence dynamically independent of $\log t_c$.

Leaving aside the effect of $g(\log t_c)$ for now, we recognize that the probability density in Fig. 9.1 is maximal along two lines of stability, $\log t_\Lambda = \log t_{\text{obs}}$ and $\log t_c = \log t_{\text{obs}}$, along which the probability force is zero. These are the boundaries between regions IV/V and I/II, respectively. They are shown by thick lines in Fig. 9.1. The fact that the

distribution is flat along these lines indicates a mild runaway problem: at the crude level of approximation we have used thus far, the probability distribution is not integrable. Let us consider each line in turn to see if this problem persists in a more careful treatment.

Along the line $\log t_\Lambda = \log t_{\text{obs}}$, the prior distribution $g(\log t_c)$ will suppress large $\log t_c$ and make the probability integrable, rendering a prediction for $\log t_c$ possible. (With a plausible prior, the probability of observing a departure from flatness in a realistic future experiment is of order 10% [67, 54]; see also [39]. See [133, 45] for different priors.) We will find the same line of stability for the other two measures in the following section, and it is lifted by the same argument.

The line $\log t_c = \log t_{\text{obs}}$ looks more serious. The prior on $\log t_\Lambda$ follows from very general considerations [181] and cannot be modified. In fact, the probability distribution is rendered integrable if, as we will suppose, the landscape contains only a finite number \mathcal{N} of vacua. This implies that there is a “discretuum limit”, a finite average gap between different possible values of Λ . This, in turn, implies that there is a smallest positive Λ in the landscape, of order

$$\Lambda_{\min} \sim 1/\mathcal{N} . \quad (9.25)$$

This argument, however, only renders the distribution integrable; it does not suffice to bring it into agreement with observation. It tells us that $\log \Lambda$ is drawn entirely at random from values between $\log \Lambda_{\min}$ and $\log t_{\text{obs}}^{-2}$, so at this level of analysis we do not predict the coincidence $\log t_{\text{obs}} \approx \log t_\Lambda$.

Even though we do not predict a coincidence for observers living at a fixed arbitrary $\log t_{\text{obs}}$, it could still be the case that after averaging over the times when observers could live most observers see $\log t_{\text{obs}} \sim \log t_\Lambda$. To address this question, we need to allow $\log t_{\text{obs}}$ to vary. For fixed $\log t_{\text{obs}}$, the maximum of the probability with respect to $\log t_c$ and $\log t_\Lambda$ is obtained along the boundary between region I and region II, as discussed above. Along this line, the probability is

$$\frac{dp}{d \log t_{\text{obs}}} \sim \frac{g(\log t_{\text{obs}})}{t_{\text{obs}}} \alpha(\log t_{\text{obs}}) . \quad (9.26)$$

Having maximized over $\log t_\Lambda$ and $\log t_c$, we can now ask at what $\log t_{\text{obs}}$ the probability is maximized. Note that since the distribution is exponential, maximizing over $\log t_\Lambda$ and $\log t_c$ is the same as integrating over them up to logarithmic corrections coming from the function g .

The location of the maximum depends on the behavior of $\alpha(\log t_{\text{obs}})$. Let us assume that

$$\alpha \sim t_{\text{obs}}^{1+p} , \quad \text{with } p > 0 . \quad (9.27)$$

(We will justify this assumption at the end of this section, where we will also describe what happens if it is not satisfied.) Then the maximum is at the largest value of $\log t_{\text{obs}}$ subject to the constraint defining regions I and II, $\log t_{\text{obs}} < \log t_\Lambda$. It follows that the maximum of the three-variable probability distribution is at

$$\log t_{\text{obs}} \approx \log t_c \approx \log t_\Lambda \approx \log t_\Lambda^{\max} . \quad (9.28)$$

Therefore, in a landscape with $p > 0$ and vacua with $\Lambda > 0$, the causal patch predicts that all three scales are ultimately set by $\log t_\Lambda^{\max}$, and thus, by the (anthropic) vacuum with

smallest cosmological constant. This, in turn, is set by the discreteness limit, i.e., by the number of vacua in the landscape [24, 37, 35, 36], according to

$$t_{\Lambda}^{\max} \sim \bar{\mathcal{N}}^{1/2} . \quad (9.29)$$

This is a fascinating result [30]. It implies that the remarkable scales we observe in Nature, such as the vacuum energy and the current age of the universe, are not only mutually correlated, but that their absolute scale can be explained in terms of the size of the landscape. If current estimates of the number of vacua [41, 55] hold up, i.e., if $\log_{10} \mathcal{N}$ is of order hundreds,² then Eq. (9.28) may well prove to be in agreement with the observed value $t_{\Lambda} \sim 10^{61}$.

Let us go somewhat beyond our order-of-magnitude estimates and determine how precisely $\log t_{\text{obs}}$ and $\log t_{\Lambda}$ can be expected to agree. To that end, we will now calculate $p_{\text{CP}}(f^{-1}t_{\Lambda} < t_{\text{obs}} < ft_{\Lambda})$ as a function of f , i.e., the probability that $\log t_{\text{obs}}$ lies within an interval $\log f$ of $\log t_{\Lambda}$. The probability distribution of Eq. 9.24 is dominated near the boundary of regions IV and V, and the probability in region IV is exponentially suppressed. So we will neglect all regions except region V. (Ignoring region IV means we are eliminating the possibility that $\log t_{\text{obs}} > \log t_{\Lambda}$.)

The probability density in region V is

$$\frac{dp}{d \log t_{\text{obs}} d \log t_c d \log t_{\Lambda}} \propto \frac{t_{\text{obs}}^{1+p}}{t_{\Lambda}} g(\log t_c) . \quad (9.30)$$

We will further restrict to $t_c > t_{\Lambda}^{\max}$, which is reasonable if t_{obs} is pushed to large values and $g(\log t_c)$ does not strongly prefer small values of $\log t_c$. Since we are computing a probability marginalized over $\log t_c$, this restriction on the range of $\log t_c$ means that the exact form of g will not affect the answer. The quantity

$$\int_{\log t_{\Lambda}^{\max}}^{\infty} d \log t_c g(\log t_c) \quad (9.31)$$

will factor out of our computations, and hence we will ignore it. Having eliminated the $\log t_c$ dependence, we continue by computing the normalization factor Z for $\log t_{\Lambda} > \log t_{\text{obs}}$:

$$Z = \int_0^{\log t_{\Lambda}^{\max}} d \log t_{\text{obs}} \int_{\log t_{\text{obs}}}^{\log t_{\Lambda}^{\max}} d \log t_{\Lambda} \frac{t_{\text{obs}}^{1+p}}{t_{\Lambda}} \approx \frac{(t_{\Lambda}^{\max})^p}{p(1+p)} . \quad (9.32)$$

In the last line we have dropped terms negligible for $t_{\Lambda}^{\max} \gg 1$.

Now we will calculate the unnormalized probability for $f^{-1}t_{\text{obs}} < t_{\Lambda} < ft_{\text{obs}}$. We will split the integration region into two subregions according to whether $t_{\text{obs}} < f^{-1}t_{\Lambda}^{\max}$ or $f^{-1}t_{\Lambda}^{\max} < t_{\text{obs}} < t_{\Lambda}^{\max}$. It turns out that each of these subregions is important. First we do

²The number of anthropic vacua, $\bar{\mathcal{N}}$, may be smaller by dozens or even hundreds of orders of magnitude than the total number of vacua, \mathcal{N} , for low-energy reasons that are unrelated to the cosmological constant or curvature and so are not included in our analysis. Hence, $\log_{10} \mathcal{N} \sim O(1000)$ may be compatible with Eq. (9.28).

$t_{\text{obs}} < f^{-1}t_{\Lambda}^{\text{max}}$:

$$\int_0^{\log(f^{-1}t_{\Lambda}^{\text{max}})} d\log t_{\text{obs}} \int_{\log t_{\text{obs}}}^{\log(ft_{\text{obs}})} d\log t_{\Lambda} \frac{t_{\text{obs}}^{1+p}}{t_{\Lambda}} \quad (9.33)$$

$$\approx \frac{(t_{\Lambda}^{\text{max}})^p}{p} (f^{-p} - f^{-1-p}) \quad (9.34)$$

$$= Z(1+p)(f^{-p} - f^{-1-p}) . \quad (9.35)$$

Finally we calculate the case $f^{-1}t_{\Lambda}^{\text{max}} < t_{\text{obs}} < t_{\Lambda}^{\text{max}}$:

$$\int_{\log(f^{-1}t_{\Lambda}^{\text{max}})}^{\log t_{\Lambda}^{\text{max}}} d\log t_{\text{obs}} \int_{\log t_{\text{obs}}}^{\log t_{\Lambda}^{\text{max}}} d\log t_{\Lambda} \frac{t_{\text{obs}}^{1+p}}{t_{\Lambda}} \quad (9.36)$$

$$= (t_{\Lambda}^{\text{max}})^p \left[\frac{1 - f^{-p}}{p} - \frac{1 - f^{-1-p}}{1+p} \right] \quad (9.37)$$

$$= Z \left[1 - (1+p)f^{-p} + pf^{-1-p} \right] . \quad (9.38)$$

Adding together the unnormalized probabilities and dividing by the factor Z we find the result

$$p_{\text{CP}}(f^{-1}t_{\text{obs}} < t_{\Lambda} < ft_{\text{obs}}) \approx 1 - f^{-1-p} . \quad (9.39)$$

In addition to being independent of $g(\log t_c)$, this result is independent of t_{Λ}^{max} , but the validity of our approximation depends on both. In particular, region V contributes more to the probability for larger t_{Λ}^{max} , so the approximation gets better as t_{Λ}^{max} increases. However, even for $t_{\Lambda}^{\text{max}} = 10^{60}$ the result is only off by a few percent compared to a numerical integration.

Let us now return to discussing our assumption, Eq. (9.27). If p were not positive, that is, if α increased at most linearly with t_{obs} , then the maximum of the probability distribution would be located at the smallest value of t_{obs} compatible with observers. In this case the causal patch would predict $t_{\Lambda} \gg t_{\text{obs}}$. This would be in conflict with observation except under the extremely contrived assumption that $t_{\Lambda}^{\text{max}} \sim t_{\text{obs}}^{\text{min}}$.

However, the assumption that $p > 0$ is quite plausible [30]. Recall that we are only discussing the form of α in the regime where t_{obs} is the shortest time scale, $t_{\text{obs}} \lesssim t_c, t_{\Lambda}$, so we do not have to worry that later observations may be disrupted by curvature or vacuum energy. Recall, moreover, that α is defined by averaging over many vacua, so we must consider only how this average depends on t_{obs} . In particular, this means that we should not imagine that in moving from one value of t_{obs} to another, we need to hold fixed the vacuum, or even restrict to only one or two parameters of particle physics and cosmology. Typical vacua with most observers at one value of t_{obs} are likely to differ in many details from vacua in which most observers arise at a different time.

With this in mind, we note two general effects that suggest that $\alpha(\log t_{\text{obs}})$ increases monotonically. First, the spontaneous formation of highly complex structures such as observers relies both on chance and, presumably, on a long chain of evolutionary processes building up increasing complexity. The later the time, the more likely it is that such a chain has been completed. Secondly, for larger t_{obs} , the same amount of mass can be distributed among more quanta, of less energy each. Therefore, less mass is necessary to construct a

system containing a given number of quanta, such as a system of sufficient complexity to function as an observer. These arguments make it very plausible that α grows. Moreover, while they do not prove that it grows more strongly than linearly with t_{obs} , they do make this type of behavior rather plausible.

9.4.2 Negative cosmological constant

We turn to negative values of the cosmological constant, beginning with the case $\Lambda < 0$, $t_c \ll t_\Lambda$. From Eqs. (9.15) and (9.19), we find that the comoving radius of the causal patch is given by

$$\chi_{\text{CP}}(t) \sim \begin{cases} 3 - 2 \log(t_c/2t_\Lambda) + 3 [1 - (t/t_c)^{1/3}] , & t < t_c \\ 3 - \log \tan(t/2t_\Lambda) + \log \tan(t'_c/2t_\Lambda) , & t_c < t < t'_c \\ 3 \left(\frac{t'}{t_c}\right)^{1/3} , & t'_c < t . \end{cases} \quad (9.40)$$

Recall that a prime denotes time remaining before the crunch: $t' \equiv t_f - t$. To better show the structure of the above expressions, we have kept some order-one factors and subleading terms that will be dropped below. We will approximate $\log \tan(t'_c/2t_\Lambda) = -\log \tan(t_c/2t_\Lambda) \approx -\log(t_c/2t_\Lambda)$.

The mass inside the causal patch at the time t_{obs} is

$$M_{\text{CP}} = \rho a^3 V_{\text{CP}}[\chi_{\text{CP}}(t_{\text{obs}})] \sim t_c V_{\text{CP}} . \quad (9.41)$$

We will again approximate the comoving volume inside a sphere of radius χ by χ^3 for $\chi \lesssim 1$ and by $e^{2\chi}$ for $\chi \gtrsim 1$, giving

$$M_{\text{CP}} \sim \begin{cases} t_\Lambda^4/t_c^3 , & t_{\text{obs}} < t_c & I \\ t_\Lambda^2 t_c^{-1} \tan^{-2}(t_{\text{obs}}/2t_\Lambda) , & t_c < t_{\text{obs}} < t'_c & II \\ t'_{\text{obs}} , & t'_c < t_{\text{obs}} & III \end{cases} \quad (9.42)$$

Now let us consider the case $t_c \gtrsim t_f/2$. The comoving radius of the causal patch is given by using Eqs. (9.18) and (9.19):

$$\chi_{\text{CP}}(t) \sim \begin{cases} (t')^{1/3} t_c^{-1/3} , & t_f/2 < t \\ 2(t_f/2t_c)^{1/3} - (t/t_c)^{1/3} , & t < t_f/2 . \end{cases} \quad (9.43)$$

The mass in the causal patch is then given by, up to order one constant factors,

$$M_{\text{CP}} \sim \begin{cases} t'_{\text{obs}} , & t_f/2 < t_{\text{obs}} & IV \\ t_f , & t_{\text{obs}} < t_f/2 & V \end{cases} \quad (9.44)$$

Now we can combine all of the above information to obtain the full probability

distribution,

$$\frac{d^3 p_{\text{CP}}}{d \log t_c d \log t_\Lambda d \log t_{\text{obs}}} \sim g\alpha \times \begin{cases} \frac{t_\Lambda^2}{t_c^3}, & t_{\text{obs}} < t_c < t_\Lambda & I \\ \frac{1}{t_c \tan^2(t_{\text{obs}}/2t_\Lambda)}, & t_c < t_{\text{obs}} < t'_c, t_c < t_\Lambda, & II \\ \frac{t'_{\text{obs}}}{t_\Lambda^2}, & t'_c < t_{\text{obs}}, t_c < t_\Lambda & III \\ \frac{t'_{\text{obs}}}{t_\Lambda^2}, & t_f/2 < t_{\text{obs}} < t_c & IV \\ \frac{1}{t_\Lambda}, & t_{\text{obs}} < t_f/2 < t_c & V \end{cases} \quad (9.45)$$

The analysis of the probability “forces” proceeds as in the positive cosmological constant case discussed in the previous subsection, by identifying and following the directions along which the probability grows in each distinct region of the $(\log t_\Lambda, \log t_c)$ plane. The result, however, is rather different (Fig. 9.1b). For fixed $\log t_{\text{obs}}$, the unnormalized probability density diverges in the direction of small $\log t_c$ and large $\log t_\Lambda$ (region II) like $t_\Lambda^2 t_c^{-1}$. The discrete spectrum of Λ bounds $\log t_\Lambda$ from above, and the Planck scale is a lower limit on $\log t_c$. Recall that so far, we have approximated the rate of observations per unit mass α as independent of $(\log t_c, \log t_\Lambda)$. However, if $t_c \ll t_{\text{obs}}$ ($t_\Lambda \ll t_{\text{obs}}$), then curvature (or vacuum energy) could dynamically affect the processes by which observers form. One would expect that such effects are generally detrimental.

Here, for the first time, we find a distribution that peaks in a regime where $t_c \ll t_{\text{obs}}$. This means that the detailed dependence of α on $\log t_c$ is important for understanding the prediction and must be included. We do not know this function except for special classes of vacua.

Instead of letting $\log t_c \rightarrow 0$ so that t_c becomes Planckian, we will only allow t_c to fall as low as t_c^{min} . We do this because it does not make our analysis any more difficult, and it may capture some aspects of anthropic selection effects if we choose to set $\log t_c^{\text{min}}$ to be some positive quantity.

Thus, within our current approximations the causal patch predicts that most observers in vacua with negative cosmological constant measure

$$\log t_c \rightarrow \log t_c^{\text{min}}, \quad \log t_\Lambda \rightarrow \log t_\Lambda^{\text{max}}, \quad (9.46)$$

where $t_\Lambda^{\text{max}} \equiv |\Lambda|_{\text{min}}^{-1/2} \sim \mathcal{N}^{1/2}$ is the largest achievable value of t_Λ in the landscape. Our result reveals a preference for separating the curvature, observer, and vacuum timescales: a hierarchy, rather than a coincidence.

What happens if $\log t_{\text{obs}}$ is also allowed to vary? After optimizing $\log t_\Lambda$ and $\log t_c$, the probability distribution over $\log t_{\text{obs}}$ is

$$\frac{dp}{d \log t_{\text{obs}}} \sim \left(\frac{t_\Lambda^{\text{max}}}{t_{\text{obs}}} \right)^2 \alpha(\log t_{\text{obs}}) \frac{g(\log t_c^{\text{min}})}{t_c^{\text{min}}}. \quad (9.47)$$

If α grows faster than quadratically in t_{obs} , then large values of $\log t_{\text{obs}}$ are predicted: $\log t_{\text{obs}} \sim \log t_\Lambda \sim \log t_\Lambda^{\text{max}}$, $\log t_c \sim \log t_c^{\text{min}}$, with the maximum probability density given

by $\alpha(\log t_\Lambda^{\max})g(\log t_c^{\min})/t_c^{\min}$. Otherwise, a small value of $\log t_{\text{obs}}$ is predicted: $\log t_\Lambda \sim \log t_\Lambda^{\max}$, $\log t_{\text{obs}} \sim \log t_{\text{obs}}^{\min}$, $\log t_c \sim \log t_c^{\min}$, with maximum probability

$$(t_\Lambda^{\max}/t_{\text{obs}}^{\min})^2 \alpha(\log t_{\text{obs}}^{\min})g(\log t_c^{\min})/t_c^{\min} .$$

(Here we have introduced $\log t_{\text{obs}}^{\min}$ in an analogous way to $\log t_c^{\min}$. The point here is that typical observers live at the earliest possible time.)

Do these predictions conflict with observation? Not so far: We are observers in a vacuum with $\Lambda > 0$, so the relevant probability distribution over $(\log t_\Lambda, \log t_c)$ is the one computed in the previous subsection. This led to the predictions that $\log t_\Lambda \sim \log t_{\text{obs}}$ and $\log t_c \gtrsim \log t_{\text{obs}}$, both of which agree well with observation; and that the scale of $\log t_\Lambda$ is controlled by the number of vacua in the landscape, which is not ruled out.

However, we do get a conflict with observation if we ask about the total probability for each sign of the cosmological constant. The total probability for positive cosmological constant is approximately given by the value of the distribution of the maximum. With our assumption about α (Eq. 9.27), this is $p_+ \sim g(\log t_\Lambda^{\max})\alpha(\log t_\Lambda^{\max})/t_\Lambda^{\max}$. The total probability for negative Λ is also controlled by the probability density at the maximum of the distribution; as mentioned earlier, it is given by $\alpha(\log t_\Lambda^{\max})g(\log t_c^{\min})/t_c^{\min}$ if $p > 1$, and by $(t_\Lambda^{\max}/t_{\text{obs}}^{\min})^2 \alpha(\log t_{\text{obs}}^{\min})g(\log t_c^{\min})/t_c^{\min}$ for $p < 1$.

Dividing these, we find that a negative value of Λ is favored by a factor

$$\frac{p_-}{p_+} = \frac{t_\Lambda^{\max} g(\log t_c^{\min})}{t_c^{\min} g(\log t_\Lambda^{\max})} \text{ for } p > 1 . \quad (9.48)$$

We know that t_Λ^{\max} must be at least as large as the observed value of t_Λ , which is of order t_{obs} : $t_\Lambda^{\max} > t_{\text{obs}} \sim 10^{61}$. Furthermore, we expect that $g(\log t_\Lambda^{\max}) < g(\log t_c^{\min})$. It follows that $p_+ < t_c^{\min}/t_\Lambda^{\max}$: the observed sign of the cosmological constant is extremely unlikely according to the causal patch measure in our simple model unless t_c^{\min} is rather close to t_{obs} . The situation is similarly bad if $p < 1$.

We regard this result as further evidence [155, 39] that the causal patch cannot be applied in regions with nonpositive cosmological constant, or more generally, in the domains of dependence of future spacelike singularities and hats. This is plausible in light of its relation to the light-cone time cutoff [26, 28], which is well-motivated [73] by an analogy to the UV/IR relation [165] of the AdS/CFT correspondence [127], but only in eternally inflating regions.

9.5 The apparent horizon cutoff

This section is structured like the previous one, but we now consider the apparent horizon cutoff, which is introduced here for the first time.

9.5.1 Definition

To define this cutoff, let us begin with a reformulation of the causal patch. We defined the causal patch as the causal past of a point on the future boundary of spacetime. But it can equivalently be characterized in terms of a worldline that ends on that point: the




Figure 9.2: The causal patch can be characterized as the union of all past light-cones (all green lines, including dashed) of the events along a worldline (vertical line). The apparent horizon cutoff makes a further restriction to the portion of each past light-cone which is expanding toward the past (solid green lines). The dot on each light-cone marks the apparent horizon: the cross-section of maximum area, where expansion turns over to contraction.

causal patch is the union of the past lightcones of all events that constitute the worldline. By the past light-cone we mean the null hypersurface that forms the boundary of the causal past.

Each past light-cone can be uniquely divided into two portions. Beginning at the tip, the cross-sectional area initially expands towards the past. But along each null geodesic generator of the light-cone, the expansion eventually becomes negative, and the cross-sectional area begins to decrease. This turnaround does not happen in all spacetimes, but it does happen in any FRW universe that starts from a big bang (Fig. 9.2) or by bubble nucleation in a vacuum of higher energy. The point along each null geodesic where the expansion vanishes and the area is maximal is called the apparent horizon [23]. The causal patch consists of both portions of the past light-cone. The apparent horizon cutoff entails a further restriction: it consists only of the portion of each light-cone which is expanding towards the past.

Our motivation for considering this cutoff comes from the preferred role played by the apparent horizon in understanding the holographic properties of cosmological spacetimes. In the terminology of Refs. [21, 22], the apparent horizon is a *preferred holographic screen*: it possesses two light-sheets going in opposite spacetime directions, which together form an entire light-cone. The covariant entropy bound states that any light-sheet off of a surface of area A contains matter with entropy $S \leq A/4$. Since the past light-cone consists of two different light-sheets off of the same surface of area A_{AH} , the entropy on it cannot exceed $A_{\text{AH}}/4 + A_{\text{AH}}/4 = A_{\text{AH}}/2$. Both the causal patch cutoff and the apparent horizon cutoff can be thought of as a restriction to the information contained on the preferred holographic




Figure 9.3: Conformal diagrams showing the apparent horizon cutoff region. The boundary of the causal patch is shown as the past light-cone from a point on the conformal boundary. The domain wall surrounding a bubble universe is shown as the future light-cone of the bubble nucleation event. The region selected by the cutoff is shaded. For $\Lambda > 0$ (a), the boundary of the causal patch is always exterior to the apparent horizon. For $\Lambda < 0$ (b), the apparent horizon diverges at a finite time. Because the apparent horizon cutoff is constructed from light-cones, however, it remains finite. The upper portion of its boundary coincides with that of the causal patch.

screen. The causal patch keeps information about both sides of the screen; the apparent horizon cutoff only about one side.

The above definition of the apparent horizon cutoff applies to arbitrary worldlines in general spacetimes. To obtain a definite ensemble of cutoff regions that can be averaged, let us specify that we follow geodesics orthogonal to an initial hypersurface specified according to some rule, for example, a region occupied by the longest lived de Sitter vacuum in the landscape [76, 43]. When a geodesic enters a new bubble, it quickly becomes comoving [34].

For a comoving geodesic in an FRW universe, it is convenient to restate the cutoff by specifying what portion of each FRW time slice should be included. The apparent horizon at equal FRW time is defined as the sphere centered on the geodesic whose orthogonal future-directed ingoing light-rays have vanishing expansion. This sphere exists at the FRW time t if and only if the total energy density is positive, $\rho(t) > 0$. Its surface area is given by [21]

$$A_{\text{AH}}(t) = \frac{3}{2\rho(t)} , \quad (9.49)$$

from which its comoving radius can easily be deduced. The apparent horizon cutoff consists of the set of points that are both within this sphere (if it exists), and within the causal patch. The former restriction is always stronger than the latter in universes with positive cosmological constant, where the apparent horizon is necessarily contained within the causal patch [22]. In universes with $\Lambda < 0$, there is an FRW time t_* when the apparent horizon coincides with the boundary of the causal patch. If $t_{\text{obs}} < t_*$, we restrict our attention to observers within the apparent horizon; otherwise we restrict to observers within the causal patch (see Fig. 9.3).

9.5.2 Positive cosmological constant

We begin with the case $\Lambda > 0$, $t_c \ll t_\Lambda$. The scale factor $a(t)$ is given by Eq. (9.12). The energy density of the vacuum, $\rho_\Lambda \sim \Lambda \sim t_\Lambda^{-2}$, begins to dominate over the density of matter, $\rho_m \sim t_c/a^3$, at the intermediate time

$$t_i \sim t_c^{1/3} t_\Lambda^{2/3} . \quad (9.50)$$

Note that $t_c \ll t_i \ll t_\Lambda$ if t_c and t_Λ are well-separated. Thus we can approximate Eq. (9.49) by

$$A_{\text{AH}}(t) = \frac{3}{2(\rho_m(t) + \rho_\Lambda)} \sim \begin{cases} \rho_m^{-1}(t) , & t < t_i \\ \rho_\Lambda^{-1} , & t > t_i \end{cases} . \quad (9.51)$$

The comoving area of the apparent horizon, A_{AH}/a^2 , is initially small and grows to about one at the time t_c . It remains larger than unity until the time t_Λ and then becomes small again. The proper volume within the apparent horizon is $V_{\text{AH}} \sim a A_{\text{AH}}$ when the comoving area is large and $V_{\text{AH}} \sim A_{\text{AH}}^{3/2}$ when it is small. The mass within the apparent horizon is $M_{\text{AH}} = \rho_m V_{\text{AH}} \sim t_c V_{\text{AH}}/a^3$. Combining the above results, we find

$$M_{\text{AH}} \sim \begin{cases} t_{\text{obs}} , & t_{\text{obs}} < t_i < t_\Lambda & I \\ t_c t_\Lambda^2 / t_{\text{obs}}^2 , & t_i < t_{\text{obs}} < t_\Lambda & II \\ t_c e^{-3(t_{\text{obs}}/t_\Lambda - 1)} , & t_i < t_\Lambda < t_{\text{obs}} & III \end{cases} \quad (9.52)$$

For the case $\Lambda > 0$, $t_\Lambda \lesssim t_c$, the mass can be obtained by setting $t_c \sim t_i \sim t_\Lambda$ in the above result:

$$M_{\text{AH}} \sim \begin{cases} t_{\text{obs}} , & t_{\text{obs}} < t_\Lambda < t_c & V \\ t_\Lambda e^{-3(t_{\text{obs}}/t_\Lambda - 1)} , & t_\Lambda < t_{\text{obs}} , t_c & IV \end{cases} \quad (9.53)$$

The full probability distribution is obtained as before by multiplying by $t_{\text{obs}} \alpha(\log t_{\text{obs}})$ and dividing by t_Λ^2 to get

$$\frac{d^3 p_{\text{AH}}}{d \log t_c d \log t_\Lambda d \log t_{\text{obs}}} \sim g \alpha \times \begin{cases} \frac{t_{\text{obs}}}{t_\Lambda^2} , & t_{\text{obs}} < t_i < t_\Lambda & I \\ \frac{t_c}{t_{\text{obs}}^2} , & t_i < t_{\text{obs}} < t_\Lambda & II \\ \frac{t_c}{t_\Lambda^2} \exp \left[-3 \left(\frac{t_{\text{obs}}}{t_\Lambda} - 1 \right) \right] , & t_i < t_\Lambda < t_{\text{obs}} & III \\ \frac{1}{t_\Lambda} \exp \left[-3 \left(\frac{t_{\text{obs}}}{t_\Lambda} - 1 \right) \right] , & t_\Lambda < t_{\text{obs}} , t_c & IV \\ \frac{t_{\text{obs}}}{t_\Lambda^2} , & t_{\text{obs}} < t_\Lambda < t_c & V \end{cases} \quad (9.54)$$

The probability forces are shown in Fig. 9.4. The boundary between regions I and II is given by $\log t_i = \log t_{\text{obs}}$, which corresponds to $\log t_\Lambda = \frac{3}{2} \log t_{\text{obs}} - \frac{1}{2} \log t_c$. In region I, the probability is proportional to t_Λ^{-2} , corresponding to a force toward smaller $\log t_\Lambda$. In region II there is a force toward large $\log t_c$. In region III, the exponential dominates the $\log t_\Lambda$




Figure 9.4: The probability distribution from the apparent horizon cutoff. The arrows indicate directions of increasing probability. For $\Lambda > 0$ (a), the probability is maximal along the boundary between regions IV and V before a prior distribution over $\log t_c$ is included. Assuming that large values of t_c are disfavored, this leads to the prediction $\log t_\Lambda \sim \log t_c \sim \log t_{\text{obs}}$. For $\Lambda < 0$ (b), the distribution is dominated by a runaway toward small t_c and large t_Λ along the boundary between regions II and III.

dependence, giving a preference for large $\log t_\Lambda$; the t_c prefactor provides a force towards large $\log t_c$. In regions IV and V the probabilities are independent of $\log t_c$ except for the prior $g(\log t_c)$. The force is towards large $\log t_\Lambda$ in region IV, while in region V small $\log t_\Lambda$ is preferred.

Following the gradients in each region, we find that the distribution peaks on the boundary between regions IV and V. Along this line, the probability density is constant except for $g(\log t_c)$. As discussed in Sec. 9.4.1, this degeneracy is lifted by a realistic prior that mildly disfavors large values of $\log t_c$. Thus, the apparent horizon cutoff predicts the double coincidence

$$\log t_{\text{obs}} \sim \log t_\Lambda \sim \log t_c . \quad (9.55)$$

This is in good agreement with observation.

What if the observer time scale is allowed to vary? After optimizing $\log t_\Lambda$ and $\log t_c$, the probability distribution over $\log t_{\text{obs}}$ is

$$\frac{dp}{d \log t_{\text{obs}}} \sim g(\log t_{\text{obs}}) \frac{\alpha(\log t_{\text{obs}})}{t_{\text{obs}}} . \quad (9.56)$$

We have argued in Sec. 9.4.1 that α grows faster than t_{obs} ; under this assumption, all three timescales are driven to the discreteness limit:

$$\log t_{\text{obs}} \approx \log t_c \approx \log t_\Lambda \approx \frac{1}{2} \log \bar{\mathcal{N}} . \quad (9.57)$$

9.5.3 Negative cosmological constant

We turn to the case $\Lambda < 0$, $t_c \ll t_\Lambda$. The scale factor is given by (9.15). The total energy density becomes negative at the intermediate time

$$t_i \sim t_c^{1/3} t_\Lambda^{2/3}, \quad (9.58)$$

when the positive matter density is sufficiently dilute to be overwhelmed by the negative vacuum energy, $\rho_\Lambda \sim -t_\Lambda^{-2}$. As discussed in Sec. 9.5.1, the apparent horizon exists on the FRW timeslice t only if the total density at that time is positive. By Eq. (9.49), the apparent horizon diverges when the density vanishes. Slightly earlier, at the time $t_* = (1 - \epsilon)t_i$, the apparent horizon intersects the boundary of the causal patch. For $t < t_*$, the apparent horizon defines the cutoff; for $t > t_*$, the causal patch does (see Fig. 9.3).

To compute t_* , notice that $t_c \ll t_\Lambda$ and Eq. (9.58) imply $t_c \ll t_i \ll t_\Lambda$. This implies that the scale factor can be well approximated by $a(t) \approx t_\Lambda \sin(t/t_\Lambda) \approx t$ in a neighborhood of t_i . This range includes the time t_* if ϵ is small. We will assume this approximation for now, and we will find that $\epsilon \ll 1$ follows self-consistently. By Eq. (9.49), the proper area of the apparent horizon at the time $t_i(1 - \epsilon)$ is $A_{\text{AH}}(t) = t_\Lambda^2/2\epsilon$. From Eq. 9.40, we find that the causal patch has proper area $16\pi e^3 t_\Lambda^4/t_c^2 + O(\epsilon^2)$. Equating these expressions, we find

$$\epsilon = \frac{1}{32\pi e^3} \frac{t_c^2}{t_\Lambda^2}, \quad (9.59)$$

which is much less than unity.

For times $t < t_*$, we compute the mass within the apparent horizon. When $t \lesssim t_c$ we use that $V_{\text{AH}} \sim A_{\text{AH}}^{3/2}$, while for $t_c \lesssim t < t_*$, we have $V_{\text{AH}} \sim aA_{\text{AH}}$. For times $t > t_*$ we use the results for the causal patch from Sec. 9.4.2.

$$M_{\text{AH}} \sim \begin{cases} t_{\text{obs}}, & t_{\text{obs}} < t_c & I \\ t_{\text{obs}} \left(1 - \left(\frac{t_{\text{obs}}}{t_i}\right)^3\right)^{-1}, & t_c < t_{\text{obs}} < t_* & II \\ t_\Lambda^2 t_c^{-1} \tan^{-2}(t_{\text{obs}}/2t_\Lambda), & t_* < t_{\text{obs}} < t'_c & III \\ t'_{\text{obs}}, & t'_c < t_{\text{obs}} & IV \end{cases}, \quad (9.60)$$

Finally, we consider the case $\Lambda < 0$, $t_f/2 < t_c$, for which the universe can be approximated as spatially flat at all times. The scale factor is given by Eq. (9.16). The area of the apparent horizon, $A_{\text{AH}} \sim t_\Lambda^2 \tan^2(\pi t/t_f)$, diverges at the turnaround time. So at a time $t_* < t_f/2$, the apparent horizon and causal patch are equal in size, and after that time we must use the causal patch as our cutoff. The area of the causal patch is $A_{\text{CP}} \sim t_\Lambda^2$ around this time, so the apparent horizon intersects the causal patch at

$$t_*^{\text{flat}} \approx \beta \frac{t_f}{2} \quad (9.61)$$

for β some order one number less than one.

The comoving size of the apparent horizon is given by $\chi \sim (t/t_c)^{1/3}$ for $t < t_*^{\text{flat}}$; for $t > t_*^{\text{flat}}$ we use our formulas from the causal patch in the previous section to obtain

$$M_{\text{AH}} \sim \begin{cases} t'_{\text{obs}}, & t_f/2 < t_{\text{obs}} < t_c & V \\ t_f, & t_*^{\text{flat}} < t_{\text{obs}} < t_f/2 < t_c & VI \\ t_{\text{obs}}, & t_{\text{obs}} < t_*^{\text{flat}} < t_c & VII \end{cases}, \quad (9.62)$$

We can now write the full probability distribution for the apparent horizon cutoff with negative cosmological constant,

$$\frac{d^3 p_{\text{AH}}}{d \log t_c d \log t_\Lambda d \log t_{\text{obs}}} \sim g \alpha \times \begin{cases} \frac{t_{\text{obs}}}{t_\Lambda^2}, & t_{\text{obs}} < t_c < t_\Lambda & I \\ \frac{t_{\text{obs}}}{t_\Lambda^2 \left[1 - \left(\frac{t_{\text{obs}}}{t_i} \right)^3 \right]}, & t_c < t_{\text{obs}} < t_* < t_\Lambda & II \\ \frac{1}{t_c \tan^2 \left(\frac{t_{\text{obs}}}{2t_\Lambda} \right)}, & t_* < t_{\text{obs}} < t'_c & III \\ \frac{t'_{\text{obs}}}{t_\Lambda^2}, & t'_c < t_{\text{obs}} < t_f & IV \\ \frac{t'_{\text{obs}}}{t_\Lambda^2}, & t_f/2 < t_{\text{obs}} < t_c & V \\ \frac{1}{t_\Lambda}, & t_*^{\text{flat}} < t_{\text{obs}} < t_f/2 < t_c & VI \\ \frac{t_{\text{obs}}}{t_\Lambda^2}, & t_{\text{obs}} < t_*^{\text{flat}} < t_c & VII \end{cases} \quad (9.63)$$

The probability force diagram is shown in Fig. 9.4. Just looking at the arrows, it is clear that the maximum of the probability distribution lies somewhere in region III, perhaps at the boundary with region II. Although the formula in region III is already reasonably simple, there is a simpler form that is correct at the same level of approximation as the rest of our analysis,

$$\frac{t_\Lambda^2}{t_c t_{\text{obs}}^2} \alpha(\log t_{\text{obs}}) g(\log t_{\text{obs}}) . \quad (9.64)$$

This is a good approximation for $t_{\text{obs}} \ll t_\Lambda$, but it is only wrong by an order one factor throughout region III, so we will go ahead and use this.

For fixed $\log t_{\text{obs}}$, it is clear that $\log t_\Lambda$ wants to be as large as possible, and $\log t_c$ as small as possible, but we must remain in region III. The condition $t_* < t_{\text{obs}}$ bounding region III is equivalent to

$$\log t_c + 2 \log t_\Lambda < 3 \log t_{\text{obs}} . \quad (9.65)$$

If $\log t_{\text{obs}}$ is big enough so that $t_c^{\min} (t_\Lambda^{\max})^2 < t_{\text{obs}}^3$, then the maximum of the distribution is at $\log t_\Lambda = \log t_\Lambda^{\max}$ and $\log t_c = \log t_c^{\min}$, with probability given by

$$\frac{(t_\Lambda^{\max})^2}{t_c^{\min} t_{\text{obs}}^2} \alpha(\log t_{\text{obs}}) g(\log t_{\text{obs}}) . \quad (9.66)$$

If $\log t_{\text{obs}}$ is smaller, then the maximum is given by $\log t_c = \log t_c^{\min}$ and $2 \log t_\Lambda = 3 \log t_{\text{obs}} - \log t_c^{\min}$, with probability

$$\frac{t_{\text{obs}}}{(t_c^{\min})^2} \alpha(\log t_{\text{obs}}) g(\log t_{\text{obs}}) . \quad (9.67)$$

In either case, we are driven to $t_c \ll t_{\text{obs}}$.

Note that, as in the case of the causal patch cutoff with $\Lambda < 0$, the distribution is peaked in a regime where $t_c \ll t_{\text{obs}}$. So there is some uncertainty in our result coming from the dependence of α on $\log t_c$ when $\log t_c < \log t_{\text{obs}}$. We do not know the form of this function, which depends on details of the nature of observers, and as before we will just continue to assume that α is independent of $\log t_c$.

Now we allow $\log t_{\text{obs}}$ to vary. For small $\log t_{\text{obs}}$ such that (9.67) is valid, $\log t_{\text{obs}}$ wants to grow given very mild assumptions about α . Eventually $\log t_{\text{obs}}$ becomes large enough that we leave the small $\log t_{\text{obs}}$ regime. For larger $\log t_{\text{obs}}$ such that (9.66) is valid, $\log t_{\text{obs}}$ is driven up to $\log t_{\Lambda}^{\text{max}}$ if α increases faster than quadratically with t_{obs} . In this case we predict $\log t_{\Lambda} \sim \log t_{\text{obs}}$. If α grows more slowly with $\log t_{\text{obs}}$, then we predict $\log t_{\text{obs}} \ll \log t_{\Lambda}$.

Let us compare the total probability for negative Λ to the total probability for positive Λ , assuming the form (9.27) for α . For negative Λ , we will assume that the large $\log t_{\text{obs}}$ regime is the relevant one, so that the correct probability distribution over t_{obs} is (9.66). Note that this is the the same as (9.47), the result for negative Λ in the causal patch. Additionally, (9.56) is identical to (9.26), the result for positive Λ in the causal patch. So the total probabilities are identical to those we found previously for the causal patch. Then a negative value of Λ is favored by a factor

$$\frac{p_-}{p_+} = \frac{t_{\Lambda}^{\text{max}} g(\log t_c^{\text{min}})}{t_c^{\text{min}} g(\log t_{\Lambda}^{\text{max}})} \text{ for } p > 1, \quad (9.68)$$

and a similar result for $p < 1$.

9.6 The fat geodesic cutoff

In this section, we compute probabilities using the fat geodesic cutoff, which considers a fixed proper volume ΔV near a timelike geodesic [34]. To compute probabilities, one averages over an ensemble of geodesics orthogonal to an initial hypersurface whose details will not matter. As discussed in the previous section, geodesics quickly become comoving upon entering a bubble of new vacuum. By the symmetries of open FRW universes, we may pick a fat geodesic at $\chi = 0$, without loss of generality.

In the causal patch and apparent horizon measure, the cutoff region is large compared to the scale of inhomogeneities, which average out. The definition of the fat geodesic, however, is rigorous only if ΔV is taken to be infinitesimal. Thus, in this section, we shall neglect the effects of local gravitational collapse. We shall approximate the universe as expanding (and, for $\Lambda < 0$ after the turnaround, contracting) homogeneously. Since the physical 3-volume, ΔV , of a fat geodesic is constant, the mass within the cutoff region is proportional to the matter density:

$$M_{FG} \propto \rho_m \sim \frac{t_c}{a^3}. \quad (9.69)$$

The fat geodesic cutoff is closely related to the scale factor time cutoff, but it is more simply defined and easier to work with. Scale factor time is defined using a congruence of timelike geodesics orthogonal to some initial hypersurface in the multiverse: $dt \equiv H d\tau$,

where τ is the proper time along each geodesic and $3H$ is the local expansion of the congruence. This definition breaks down in nonexpanding regions such as dark matter halos; attempts to overcome this limitation (e.g., Ref. [53]) remain somewhat ad-hoc. In regions where the congruence is everywhere expanding, scale factor time is exactly equivalent to the fat geodesic cutoff with initial conditions in the longest lived de Sitter vacuum [34].

9.6.1 Positive cosmological constant

We begin with the case $\Lambda > 0$, $t_c \ll t_\Lambda$. Combining Eqs. (9.69) and (9.12), we obtain

$$M_{\text{FG}} \sim \begin{cases} 1/t_{\text{obs}}^2, & t_{\text{obs}} < t_c < t_\Lambda & I \\ t_c/t_{\text{obs}}^3, & t_c < t_{\text{obs}} < t_\Lambda & II \\ (t_c/t_\Lambda^3)e^{-3t_{\text{obs}}/t_\Lambda}, & t_c < t_\Lambda < t_{\text{obs}} & III \end{cases} . \quad (9.70)$$

For the flat universe ($\Lambda > 0$, $t_c > t_\Lambda$), we obtain

$$M_{\text{FG}} \sim \begin{cases} 1/t_{\text{obs}}^2, & t_{\text{obs}} < t_\Lambda & V \\ (1/t_\Lambda^2)e^{-3t_{\text{obs}}/t_\Lambda}, & t_\Lambda < t_{\text{obs}} & IV \end{cases} . \quad (9.71)$$

This leads to the probability distribution

$$\frac{d^3 p_{\text{FG}}}{d \log t_c d \log t_\Lambda d \log t_{\text{obs}}} \sim g\alpha \times \begin{cases} \frac{1}{t_\Lambda^2 t_{\text{obs}}^2}, & t_{\text{obs}} < t_c < t_\Lambda & I \\ \frac{t_c}{t_{\text{obs}}^3 t_\Lambda^2}, & t_c < t_{\text{obs}} < t_\Lambda & II \\ \frac{t_c}{t_\Lambda^5} \exp\left[-3\frac{t_{\text{obs}}}{t_\Lambda}\right], & t_c < t_\Lambda < t_{\text{obs}} & III \\ \frac{1}{t_\Lambda^4} \exp\left[-3\frac{t_{\text{obs}}}{t_\Lambda}\right], & t_\Lambda < t_{\text{obs}}, t_c & IV \\ \frac{1}{t_{\text{obs}}^2 t_\Lambda^2}, & t_{\text{obs}} < t_\Lambda < t_c & V \end{cases} \quad (9.72)$$

The probability force diagram is shown in Fig. 9.5. The result is the same as for the apparent horizon cutoff: the distribution peaks on the entire line separating regions IV and V, up to the effects of $g(\log t_c)$. A realistic prior that mildly disfavors large values of $\log t_c$ will tend to make $\log t_c$ smaller. Thus, the fat geodesic cutoff predicts the double coincidence

$$\log t_{\text{obs}} \sim \log t_\Lambda \sim \log t_c , \quad (9.73)$$

in good agreement with observation.

What if we allow $\log t_{\text{obs}}$ to scan? Optimizing $(\log t_\Lambda, \log t_c)$, we find the probability distribution over $\log t_{\text{obs}}$:

$$\frac{dp}{d \log t_{\text{obs}}} \sim \frac{\alpha(\log t_{\text{obs}})g(\log t_{\text{obs}})}{t_{\text{obs}}^4} . \quad (9.74)$$

The denominator provides a strong preference for $\log t_{\text{obs}}$ to be small. To agree with observation, α must grow at least like the fourth power of t_{obs} for values of t_{obs} smaller than the




Figure 9.5: The probability distribution computed from the scale factor (fat geodesic) cutoff. The arrows indicate directions of increasing probability. For $\Lambda > 0$ (a), the probability distribution is maximal along the boundary between regions IV and V; with a mild prior favoring smaller $\log t_c$, this leads to the prediction of a nearly flat universe with $\log t_c \sim \log t_\Lambda \sim \log t_{\text{obs}}$. For $\Lambda < 0$ (b), the probability distribution diverges as the cosmological constant increases to a value that allows the observer timescale to coincide with the big crunch.

observed value $t_{\text{obs}} \sim 10^{61}$. We cannot rule this out, but it is a much stronger assumption than the ones needed for the causal patch and apparent horizon cutoffs.

The preference for early $\log t_{\text{obs}}$ in the fat geodesic cutoff can be traced directly to the fact that the probability is proportional to the matter density. This result has an interesting manifestation [34] in the more restricted setting of universes similar to our own: it is the origin of the strong preference for large initial density contrast, $\delta\rho/\rho$, which allows structure to form earlier and thus at higher average density.

9.6.2 Negative cosmological constant

For $\Lambda < 0$, $t_c \ll t_\Lambda$, we use Eq. (9.15) for the scale factor. The mass in the cutoff region is

$$M_{\text{FG}} \sim \begin{cases} 1/t_{\text{obs}}^2, & t_{\text{obs}} < t_c & I \\ (t_c/t_\Lambda^3) \sin^{-3}(t_{\text{obs}}/t_\Lambda), & t_c < t_{\text{obs}} < t'_c & II \\ 1/t_{\text{obs}}^2, & t'_c < t_{\text{obs}} & III \end{cases} . \quad (9.75)$$

(Recall that a prime denotes the time remaining before the crunch, $t' \equiv t_f - t$.) For the flat universe case, $\Lambda < 0$ and $t_c > t_f/2$, we use Eq. (9.16) for the scale factor and find

$$M_{\text{FG}} \sim t_\Lambda^{-2} \sin^{-2}(\pi t_{\text{obs}}/t_f), \quad t_f/2 < t_c \quad IV. \quad (9.76)$$

The probability distribution is then

$$\frac{d^3 p_{\text{FG}}}{d \log t_c d \log t_\Lambda d \log t_{\text{obs}}} \sim g \alpha \times \begin{cases} \frac{1}{t_\Lambda^2 t_{\text{obs}}^2}, & t_{\text{obs}} < t_c < t_\Lambda & I \\ \frac{t_c}{t_\Lambda^5 \sin^3(t_{\text{obs}}/t_\Lambda)}, & t_c < t_{\text{obs}} < t'_c & II \\ \frac{1}{t_\Lambda^2 (t'_{\text{obs}})^2}, & t'_c < t_{\text{obs}} & III \\ \frac{1}{t_\Lambda^4 \sin^2(\pi t_{\text{obs}}/t_f)}, & t_f/2 < t_c & IV \end{cases} \quad (9.77)$$

The probability force diagram is shown in Fig. 9.5. At fixed $\log t_{\text{obs}}$, the scale factor measure predicts that observers exist just before the crunch ($\log t'_{\text{obs}} \rightarrow \log t_{\text{obs}}^{\text{min}'}$). Recall that $t_{\text{obs}}^{\text{min}}$ was introduced as a proxy for more detailed anthropic selection effects. The point is that the measure provides a pressure which favors observers living as close as possible to the crunch.

We can now find the probability distribution over t_{obs} . In the previous sections, up to logarithmic corrections it did not matter whether we optimized or marginalized over (t_Λ, t_c) because the distribution near the maximum was exponential. Here, we will get different answers for the two procedures, so we choose to marginalize over $(\log t_\Lambda, \log t_c)$, leaving $\log t_{\text{obs}}$ to scan. The resulting distribution is

$$\frac{dp}{d \log t_{\text{obs}}} \sim t_{\text{obs}}^{-3} \alpha(\log t_{\text{obs}}). \quad (9.78)$$

There is no geometric pressure on $\log t_c$ in region III, where Eq. (9.75) peaks, so the value of $\log t_c$ will be determined by the prior distribution and anthropic selection. Assuming that the prior favors small values of $\log t_c$, it seems likely that expected value of $\log t_c$ is much less than $\log t_{\text{obs}}$. As in the apparent horizon and causal patch measures for $\Lambda < 0$, this complicates the computation of α . However, the situation here is not the same. The difference is that here we have observers forming late in the recollapse phase of a crunching universe, where the dominant contribution to the energy density actually comes from matter. The fact that the universe is in a recollapse phase makes it very hard to say what the form of α will be, whether or not there is an era of curvature domination.

Regardless of the form of α , the first factor in Eq. (9.78) has a preference for $\log t_{\text{obs}}$ to be small. If α it grows faster than t_{obs}^3 , then it is favorable for $\log t_{\text{obs}}$ to be large and $\log t_{\text{obs}} \rightarrow \log t_\Lambda^{\text{max}}$. Otherwise, $\log t_{\text{obs}} \rightarrow \log t_{\text{obs}}^{\text{min}}$, which means that some anthropic boundary determines the expected value.

Now we will estimate the preference for negative values of Λ over positive by integrating the distributions in Eqs. (9.78) and (9.74). As mentioned above, to get agreement with the observed value of Λ we need to assume α grows like a fourth power of t_{obs} . Then for both positive and negative Λ , the distribution is sharply peaked at $t_{\text{obs}} \sim t_\Lambda^{\text{max}}$. Then we find

$$p_-/p_+ \sim t_\Lambda^{\text{max}}. \quad (9.79)$$

So negative values of the cosmological constant are favored.

Finally, for $\Lambda < 0$ it is worth noting the behavior of the probability distribution over $\log t_{\text{obs}}$ for fixed $\log t_{\Lambda}$, using for instance Eq. (9.76) and neglecting for simplicity the factor α . Depending on whether t_{obs} is larger or smaller than $t_f/2$, $\log t_{\text{obs}}$ will be driven either to $\log t_{\text{obs}}^{\text{min}'}$ or to $\log t_{\text{obs}}^{\text{min}}$. The former case is reproduced by our above procedure of fixing $\log t_{\text{obs}}$ and letting $\log t_{\Lambda}$ vary. The latter case is the time-reversed case (and we know that the fat geodesic measure respects the time-reversal symmetry of a crunching universe). When both $\log t_{\Lambda}$ and $\log t_{\text{obs}}$ are allowed to vary, we are driven to $\log t_{\Lambda} \sim \log t_{\text{obs}} \sim \log t_{\text{obs}}^{\text{min}}$ regardless of the order of scanning.

Recall that the fat geodesic cutoff is equivalent to the scale factor measure in simple situations. However, our negative conclusions about negative Λ differ from the analysis of the scale factor measure in [53] which found no conflict with observation. There are two reasons for this discrepancy. First, the fat geodesic measure differs from the detailed prescription given in [53] in the recollapsing region. Second, the analysis of [53] made an unjustified approximation [34], computing the scale factor time in the approximation of a homogeneous FRW universe. It remains to be seen if there is a precise definition of the scale factor cutoff that will give the result computed in [53]. The fat geodesic is our best attempt to define a simple measure in the spirit of [53].

Chapter 10

Future inextendability and non-decoupling of the regulator

10.1 Non-decoupling of the regulator

A sufficiently large region of space with positive vacuum energy will expand at an exponential rate. If the vacuum is stable, this expansion will be eternal. If it is metastable, then the vacuum can decay by the nonperturbative formation of bubbles of lower vacuum energy. Vacuum decay is exponentially suppressed, so for a large range of parameters the metastable vacuum gains volume due to expansion faster than it loses volume to decays [86]. This is the simplest nontrivial example of eternal inflation.

If it does occur in Nature, eternal inflation has profound implications. Any type of event that has nonzero probability will happen infinitely many times, usually in widely separated regions that remain forever outside of causal contact. This undermines the basis for probabilistic predictions of local experiments. If infinitely many observers throughout the universe win the lottery, on what grounds can one still claim that winning the lottery is unlikely? To be sure, there are also infinitely many observers who do not win, but in what sense are there more of them? In local experiments such as playing the lottery, we have clear rules for making predictions and testing theories. But if the universe is eternally inflating, we no longer know *why* these rules work.

To see that this is not merely a philosophical point, it helps to consider cosmological experiments, where the rules are less clear. For example, one would like to predict or explain features of the CMB; or, in a theory with more than one vacuum, one might wish to predict the expected properties of the vacuum we find ourselves in, such as the Higgs mass. This requires computing the relative number of observations of different values for the Higgs mass, or of the CMB sky. There will be infinitely many instances of every possible observation, so what are the probabilities? This is known as the “measure problem” of eternal inflation.

In order to yield well-defined probabilities, eternal inflation requires some kind of regulator. Here we shall focus on geometric cutoffs, which discard all but a finite portion of the eternally inflating spacetime. The relative probability of two types of events, 1 and




Figure 10.1: A multiverse populated by infinitely many observers (vertical line segments) who first see 1 o'clock (at events labeled “1”) and then 2 o'clock (“2”). A geometric cutoff selects a finite set of events, whose relative frequency defines probabilities. Events that are not counted are indicated by dashed lines. The left figure shows a global cutoff: all events prior to the time t_0 (curved line) are counted and all later events ignored. (The global time has nothing to do with the observers’ clocks, which come into being at a rate dictated by the dynamics of eternal inflation.) The right figure shows the causal patch cutoff, which restricts to the causal past of a point on the future boundary. In both figures, the cutoff region contains observers who see 1 o'clock but not 2 o'clock. Their number, as a fraction of all observers who see 1 o'clock before the cutoff, defines the probability of reaching the end of time between 1 and 2 o'clock.

2, is then defined by

$$\frac{p_1}{p_2} = \frac{\langle N_1 \rangle}{\langle N_2 \rangle} \quad (10.1)$$

where $\langle N_1 \rangle$ is the expected number of occurrences of the first type of event within the surviving spacetime region. (We will drop the expectation value brackets below for simplicity of notation.) Here, 1 and 2 might stand for winning or not winning the lottery; or they might stand for a red or blue power spectrum in the CMB. The generalization to larger or continuous sets of mutually exclusive outcomes is trivial.

There are different proposals for what spacetime region should be retained. Our basic observation in this chapter applies to all geometric cutoffs we are aware of, and indeed seems to be an inevitable consequence of any simple geometric cutoff: Some observers will have their lives interrupted by the cutoff (Fig. 10.1).

Let events 1 and 2 be the observation of 1 o'clock and 2 o'clock on an observer's watch. For simplicity, we will suppose that local physics can be approximated as deterministic, and we neglect the probability that the observer may die between 1 and 2 o'clock, or that the clock will break, etc. Each observer is born just before his watch shows 1 and dies just after it shows 2, so that no observer can see more than one event of each type.

Conventionally, we would say that every observer sees both 1 o'clock and then 2 o'clock. But the figure shows that for some observers, 2 o'clock will not be under the cutoff even if 1 o'clock is. A fraction $1 - N_2/N_1$ of observers are prevented from surviving to 2 o'clock. The catastrophic event in question is evidently the cutoff itself: the observer might run into the end of time.

One can imagine a situation where the relative number of observations of 1 o'clock

and 2 o'clock is relevant to predicting the results of an experiment. Suppose that the observers fall asleep right after seeing 1 o'clock. They wake up just before 2 o'clock with complete memory loss: they have no idea whether they have previously looked at their watches before. In this case, they may wish to make a prediction for what time they will see. Since $N_2 < N_1$, by Eq. (10.1), $p_2 < p_1$: the observation of 2 o'clock is less probable than that of 1 o'clock. This is possible only if some observers do not survive to 2 o'clock.

The conclusion that time can end obtains whether or not the observers have memory loss. Consider an observer who retains her memory. She is aware that she is about to look at her watch for the first time or for the second time, so the outcome won't be a surprise on either occasion. But this does not contradict the possibility that some catastrophic event may happen between 1 and 2 o'clock. The figure shows plainly that this event *does* happen to a nonzero fraction of observers. The only thing that changes when considering observers who remember is the type of question we are likely to ask. Instead of asking about two alternative events (1 or 2), we may find it more natural to ask about the relative probability of the two different possible histories that observers can have. One history, "1–", consists of seeing 1 o'clock and running into the cutoff. The alternative, "12", consists of seeing 1 o'clock and then seeing 2 o'clock. From Fig. 10.1 we see that $N_{12} = N_2$ and $N_{1-} = N_1 - N_2$. Since $N_1 > N_2$, we have $p_{1-} > 0$: there is a nonzero probability for the history in which the observer is disrupted by the end of time.

Outline and Frequently Asked Questions The probability for the end of time is nonzero for all geometric cutoffs. Its value, however, depends on the cutoff. In Sec. 10.2 we compute the probability, per unit proper time, that we will encounter the end of time.

A number of objections may be raised against our conclusion that time can end.

- Q: Can't I condition on the end of time not happening?¹

A: Certainly. This is like asking what the weather will be tomorrow, supposing that it will not rain. It is a reasonable question with a well-defined answer: The sun will shine with probability x , and it will snow with probability $1 - x$. But this does not mean that it cannot rain. If the end of time is a real possibility, then it cannot be prevented just by refusing to ask about it.

- Q: In some measures, the cutoff is taken to later and later times. In this limit, the probability to encounter the end of time surely approaches 0?

A: No. In all known measures of this type, an attractor regime is reached where the number of all types of events grows at the same exponential rate, including observers

¹In the above example, this would force us to ask a trivial question ("What is the relative probability of seeing 1 or 2, for an observer whose history includes both 1 and 2?"), which yields the desired answer ($p_2/p_1 = 1$). For a more interesting example, consider an experiment that terminates at different times depending on the outcome, such as the Guth-Vanchurin paradox described in Sec. 10.4, or the decay of a radioactive atom. In such experiments it is tempting to propose that the experiment should be regarded to last for an amount of time corresponding to the latest possible (conventional) outcome, regardless of the actual outcome; and that any outcome (other than the end of time) should be counted only if the experiment has been entirely completed before the cutoff, in the above sense. This proposal is not only futile (as described in the answer), but also ill-defined, since any event in the past light-cone of the event P can be regarded as the beginning of an experiment that includes P .

who see 1 o'clock. The fraction of these observers who also see 2 o'clock before the cutoff approaches a constant less than unity, as will be shown in Sec. 10.3.1.

- Q: But as the cutoff is taken to later times, any given observer's entire history is eventually included. Isn't this a contradiction?

A: No. We do not know which observer we are, so we cannot identify with any particular observer. (If we could, there would be no need for a measure.) Rather, we consider all observers in a specified class, and we define probabilities in terms of the relative frequency of different observations made by these observers.

- Q: If I looked at what happened on Earth up to the present time (my "cutoff"), I would find not only records of past clocks that struck both 1 and 2, but also some recently manufactured clocks that have struck 1 but not yet 2. I could declare that the latter represent a new class of clocks, namely clocks whose existence is terminated by my "cutoff". But I know that this class is fake: it wasn't there before I imposed the "cutoff". Surely, the end of time in eternal inflation is also an artifact that ought to be ignored?

A: Only a finite number of clocks will ever be manufactured on Earth. Probabilities are given not in terms of the sample picked out by your "cutoff", but by relative frequencies in the entire ensemble. If every clock ever built (in the past or future) strikes both 1 and 2, then the probability for a randomly chosen clock to undergo a different history vanishes, so we may say confidently that the "cutoff" has introduced an artifact. In eternal inflation, however, the cutoff cannot be removed. Otherwise, we would revert to a divergent multiverse in which relative frequencies are not well-defined. The cutoff defines not a sample of a pre-existing ensemble; it defines the ensemble. This is further discussed in Sec. 10.3.2.

- Q: Why not modify the cutoff to include 2 o'clock?

A: This is a possibility. If we deform the cutoff hypersurface so that it passes through no matter system, then nothing will run into the end of time. It is not clear whether this type of cutoff can be obtained from any well-defined prescription. At a minimum, such a prescription would have to reference the matter content of the universe explicitly in order to avoid cutting through the world volumes of matter systems. In this chapter, we consider only cutoffs defined by a purely geometric rule, which take no direct account of matter.

In Sec. 10.4, we discuss an apparent paradox that is resolved by the nonzero probability for time to end.

Any conclusion is only as strong as the assumptions it rests on. The reader who feels certain that time cannot end may infer that at least one of the following assumptions are wrong: (1) the universe is eternally inflating; (2) we may restrict attention to a finite subset of the eternally inflating spacetime, defined by a purely geometric prescription; and (3) probabilities are computed as relative frequencies of outcomes in this subset, Eq. (10.1). We discuss these assumptions in Sec. 10.5.1.

In Sec. 10.5.2, we discuss whether, and how, the nonzero probability for the end of time may be observed. We point out that known predictions of various measures actually arise from the possibility that time can end. On the problematic side, this includes the famous youngness paradox of the proper time cutoff; on the successful side, the prediction of the cosmological constant from the causal patch cutoff.

In Sec. 10.5.3, we discuss how the end of time fits in with the rest of physics. This depends on the choice of cutoff. With the causal patch cutoff, there may be a relatively palatable interpretation of the end of time which connects with the ideas of black hole complementarity. The boundary of the causal patch is a kind of horizon, which can be treated as an object with physical attributes, including temperature. Matter systems that encounter the end of time are thermalized at this horizon. This is similar to an outside observer's description of a matter system falling into a black hole. What is radically new, however, is the statement that *we* might experience thermalization upon crossing the black hole horizon.

This work was inspired by discussions with Alan Guth, who first described to us the paradox mentioned in section 10.4. We understand that Guth and Vanchurin will be publishing their own conclusions [85]. In taking seriously the incompleteness of spacetime implied by geometric cutoffs, our conclusion resembles a viewpoint suggested earlier by Ken Olum [138].

10.2 The probability to encounter the cutoff

The phenomenon that time can end is universal to all geometric cutoffs. But the rate at which this is likely to happen, per unit proper time τ along the observer's worldline, is cutoff-specific. We will give results for five measures.

Causal patch The causal patch cutoff [24] restricts attention to the causal past of the endpoint of a single worldline (see Fig. 10.1). Expectation values are computed by averaging over initial conditions and decoherent histories in the causal patch. The end of time, in this case, is encountered by systems that originate inside the causal patch but eventually exit from it.

Our universe can be approximated as a flat FRW universe with metric

$$ds^2 = -d\tau^2 + a(\tau)^2(d\chi^2 + \chi^2 d\Omega^2) . \quad (10.2)$$

Observers are approximately comoving ($d\chi/d\tau = 0$). We assume that the decay rate of our vacuum, per unit four-volume, is much less than t_Λ^{-4} . Then the decay can be neglected entirely in computing where the boundary of the causal patch intersects the equal time surfaces containing observers. The boundary is given by the de Sitter event horizon:

$$\chi_E(\tau) = \int_\tau^\infty \frac{d\tau'}{a(\tau')} . \quad (10.3)$$

We consider all observers currently within the horizon: $\chi < \chi_E(\tau_0)$, with $\tau_0 = 13.7$ Gyr. This corresponds to a comoving volume $V_{\text{com}} = (4\pi/3)\chi_E(\tau_0)^3$. Observers located at

χ leave the patch at a time τ' determined by inverting Eq. (10.3); in other words, they reach the end of time at $\Delta\tau \equiv \tau' - \tau_0$ from now. An (unnormalized) probability distribution over $\Delta\tau$ is obtained by computing the number of observers that leave the causal patch at the time $\tau_0 + \Delta\tau$:

$$\frac{dp}{d\Delta\tau} \propto \frac{4\pi\chi_E(\tau_0 + \Delta\tau)^2}{a(\tau_0 + \Delta\tau)}. \quad (10.4)$$

We compute $a(\tau)$ numerically using the best-fit cosmological parameters from the WMAP5 data combined with SN and BAO [115]. From the distribution (10.4), we may obtain both the median and the expectation value for $\Delta\tau$. We find that the expected amount of proper time left before time ends is

$$\langle\Delta\tau\rangle = 5.3 \text{ Gyr} . \quad (10.5)$$

Time is unlikely to end in our lifetime, but there is a 50% chance that time will end within the next 3.7 billion years.

Light-cone time The light-cone time of an event is defined in terms of the volume of its future light-cone on the future boundary of spacetime [26, 44, 28]. The light-cone time cutoff requires that we only consider events prior to some light-cone time t_0 ; then the limit $t_0 \rightarrow \infty$ is taken. It can be shown that the light-cone time cutoff is equivalent to the causal patch cutoff with particular initial conditions [44]. Thus, the probability for an observer to encounter the end of time is the same as for the causal patch cutoff.

Fat geodesic The fat geodesic cutoff considers a fixed proper volume $4\pi d^3/3$ near a timelike geodesic [34]. To compute probabilities, one averages over an ensemble of geodesics orthogonal to an initial hypersurface whose details will not matter. One can show that the geodesics quickly become comoving after entering a bubble of new vacuum. Since our vacuum is homogeneous, we may pick without loss of generality a fat geodesic at $\chi = 0$. We shall neglect the effects of local gravitational collapse and approximate the universe as expanding homogeneously. Equivalently, we take the proper distance d to be small compared to the present curvature scale of the universe but large compared to the scale of clusters. These approximations are not essential, but they will simplify our calculation and save us work when we later consider the scale factor cutoff.

We should only consider observers who are currently ($\tau_0 = 13.7$ Gyr) within the fat geodesic, with $\chi < d/a(\tau_0)$. An observer leaves the geodesic a time $\Delta\tau$ later, with $\chi = d/a(\tau_0 + \Delta\tau)$. The unnormalized probability distribution over $\Delta\tau$ is

$$\frac{dp}{d\Delta\tau} \propto 4\pi \frac{d^3 (da/d\tau)_{\tau_0 + \Delta\tau}}{a(\tau_0 + \Delta\tau)^4} . \quad (10.6)$$

From this distribution, we find numerically that the expected amount of proper time left before the end of time is 5 Gyr. There is a 50% chance that time will end within the next 3.3 billion years.

While the result is similar, there is an important formal difference between the fat geodesic and causal patch cutoffs. The boundary of the fat geodesic is a timelike hypersurface, from which signals can propagate into the cutoff region. Boundary conditions

must therefore be imposed. When a system leaves the fat geodesic, time ends from its own point of view. But an observer who remains within the cutoff region continues to see the system and to communicate with it. The image of the system and its response to any communications are encoded in data specified on the timelike boundary. In practice, the simplest way to determine these boundary conditions is to consider the global spacetime and select a fat geodesic from it. This means that the fat geodesic is not a self-contained description. The content of the causal patch, by contrast, can be computed from its own initial conditions without reference to a larger spacetime region.

Scale factor time Scale factor time is defined using a congruence of timelike geodesics orthogonal to some initial hypersurface in the multiverse: $dt \equiv H d\tau$, where τ is the proper time along each geodesic and $3H$ is the local expansion of the congruence. This definition breaks down in nonexpanding regions such as dark matter halos; attempts to overcome this limitation (e.g., Ref. [53]) remain somewhat ad-hoc. Here we use for H the Hubble rate of a completely homogeneous universe whose density agrees with the average density of our universe. This does not yield a precise and general cutoff prescription, but it allows us to compute an approximate rate at which we are likely to encounter the cutoff: in an everywhere-expanding timelike geodesic congruence, the scale factor time cutoff is equivalent to the fat geodesic cutoff [34]. Hence it gives the same rate for time to end as the fat geodesic cutoff.

Proper time In the proper time cutoff, the characteristic timescale is the shortest Hubble time of all eternally inflating vacua. In a realistic landscape, this is microscopically short, perhaps of order the Planck time [33]. Thus, time would be overwhelmingly likely to end in the next second:

$$\frac{dp}{d\Delta\tau} \approx t_{\text{Pl}}^{-1}. \quad (10.7)$$

This is the famous “youngness paradox” in a new guise. The cutoff predicts that our observations have superexponentially small probability, and that most observers are “Boltzmann babies” who arise from quantum fluctuations in the early universe. Thus, this measure is already ruled out phenomenologically at a high level of confidence [124, 81, 82, 83, 175, 122, 84, 33].

10.3 Objections

Our intuition rebels against the conclusion that spacetime could simply cease to exist. In the introduction, we answered several objections that could be raised against the end of time. In this section, we will discuss two of these arguments in more detail.

10.3.1 The cutoff can not be encountered in a late-time limit

In some measure proposals, such as the proper time cutoff [120, 121], the scale factor time cutoff [53], and the light-cone time cutoff [26], a limit is taken in which the

cutoff region is made larger and larger as a function of a time parameter t_0 :

$$p_1/p_2 = \lim_{t_0 \rightarrow \infty} N_1(t_0)/N_2(t_0) . \quad (10.8)$$

Naively one might expect the cutoff observers to be an arbitrarily small fraction of all observers in the limit $t_0 \rightarrow \infty$. This turns out not to be the case.

One finds that the number of events of type I that have occurred prior to the time t is of the form

$$N_I(t) = \tilde{N}_I \exp(\gamma t) + O(\sigma t) , \quad (10.9)$$

with $\sigma < \gamma \approx 3$. Thus, the growth approaches a universal exponential behavior at late times [76, 33, 44], independently of initial conditions. The ratio N_1/N_2 appearing in Eq. (10.1) remains well-defined in the limit $t_0 \rightarrow \infty$, and one obtains

$$\frac{p_1}{p_2} = \frac{\tilde{N}_1}{\tilde{N}_2} . \quad (10.10)$$

The constants \tilde{N}_I , and thus the relative probabilities, depend on how the time variable is defined; we will discuss some specific choices below.

Suppose that observers live for a fixed global time interval Δt . Then a person dies before time t if and only if he was born before $t - \Delta t$. Therefore the number of births N_b is related to the number of deaths N_d by

$$N_d(t) = N_b(t - \Delta t) \quad (10.11)$$

Using the time dependence of N_b given in (10.9), this can be rewritten

$$\frac{N_d(t)}{N_b(t)} \approx \exp(-\gamma \Delta t) , \quad (10.12)$$

up to a correction of order $e^{(\sigma-\gamma)t}$ which becomes negligible in the late time limit. Thus, the fraction of deaths to births does not approach unity as the cutoff is taken to infinity. The fraction of observers whose lives are interrupted by the cutoff is

$$\frac{N_c}{N_b} = 1 - \exp(-\gamma \Delta t) \quad (10.13)$$

where $N_c = N_b - N_d$ is the number of cutoff observers.

Since (10.13) is true for any time interval Δt , it is equivalent to the following simple statement: any system has a constant probability to encounter the end of time given by

$$\frac{dp}{dt} = \gamma \approx 3 \quad (10.14)$$

This result can be interpreted as follows. Due to the steady state behavior of eternal inflation at late times, there is no way to tell what time it is. The exponential growth (10.9) determines a t_0 -independent probability distribution for how far we are from the cutoff, given by (10.14).

10.3.2 Hitting the cutoff is an artifact

Could it be that observers who run into the cutoff are an artifact, not to be taken seriously as a real possibility? Certainly they would not exist but for the cutoff. Yet, we argue that cutoff observers are a real possibility, because *there is no well-defined probability distribution without the cutoff; in particular, only the cutoff defines the set of allowed events*. In order to convince ourselves of this, it is instructive to contrast this situation with one where a cutoff may introduce artifacts. We will consider two finite ensembles of observers, without reference to eternal inflation. We then restrict attention to a portion of each ensemble, defined by a cutoff. We find that this sample looks the same in both ensembles, and that it contains observers that run into the cutoff. In the first ensemble, these observers are an artifact of sampling; in the second, they are real. We will then explain why eternal inflation with a cutoff is different from both examples.

A cutoff on a finite ensemble defines a sample Consider a civilization which begins at the time $t = 0$ and grows exponentially for a finite amount of time. (We will make no reference to a multiverse in this example.) Every person is born with a watch showing 1 o'clock at the time of their birth, when they first look at it. One hour later, when they look again, the watch shows 2 o'clock; immediately thereafter the person dies. After the time $t_* \gg 1$ hour, no more people are born, but whoever was born before t_* gets to live out their life and observe 2 o'clock on their watch before they die. In this example, there is a well-defined, finite ensemble, consisting of all observers throughout history and their observations of 1 and 2. The ensemble contains an equal number of 1's and 2's. Every observer in the ensemble sees both a 1 and a 2, each with 100% probability. No observer meets a premature demise before seeing 2.

Now suppose that we do not know the full ensemble described above. Instead, we are allowed access only to a finite sample drawn from it, namely everything that happened by the time t , with 1 hour $\ll t < t_*$. This sample contains many observers who died before t ; each of them will have seen both 1 and 2. We refer to these as “histories” of type 12. It also contains observers (those who were born after $t - 1$ hour) who are still alive. Each of them has seen 1 but not yet 2, by the time t , which we refer to as a history of type 1-. What do we say about these observers? Should we declare that there is a nonzero probability that an observer who sees 1 does not live to see 2? In fact, a finite sample of a larger, unknown ensemble allows us to draw no conclusion of this kind, because we have no guarantee that our sampling of the ensemble is fair. The true set of outcomes, and their relative frequency, is determined only by the full ensemble, not by our (possibly unfaithful) sample of it. Similarly, if we had a considered a more complicated system, such as observers with watches of different colors, etc., the relative frequency of outcomes in any subset need not be the same as the relative frequency in the full ensemble, unless we make further assumptions.

If we examined the full ensemble, we could easily verify that every observer who sees 1 also lives to see 2. Thus we would learn that 1- was an artifact of our sampling: imposing a cutoff at fixed time produced a new class of events that does not exist (or more precisely, whose relative frequency drops to zero) once we take the cutoff away. Armed with this knowledge, we could then invent an improved sampling method, in which the 1- cases

are either excluded, or treated as 12 events.

As our second example, let us consider a civilization much like the previous one, except that it perishes not by a sudden lack of births, but by a comet that kills everyone at the time t_* . This, too, gives rise to a finite, well-defined ensemble of observations. But unlike in the previous example, there is a larger number of 1's than 2's: not every observer who sees a 1 lives to see a 2. Thus, the probabilities for the histories 12 and 1- satisfy $p_{1-} > 0$, $p_{12} < 1$. Indeed, if we choose parameters so the population grows exponentially on a timescale much faster than 1 hour, most people in history who see 1 end up being killed by the comet rather than expiring naturally right after seeing 2; that is, $p_{12} = 1 - p_{1-} \ll 1$ in this limit.

Again, we can contemplate sampling this ensemble, i.e., selecting a subset, by considering everything that happened prior to the time $t < t_*$. Note that this sample will look identical to the finite-time sample we were given in the previous example. Again, we find that there are apparently events of type 1-, corresponding to observers who have seen 1 but not 2 by the time t . But in this example, it so happens that (i) events of type 1- actually do exist in the full ensemble, i.e., have nonzero true relative frequency; and (ii) assuming exact exponential growth, our sample is faithful: the relative frequency of 1- vs. 12 in the sample (observers prior to t) is the same as in the full ensemble (observers in all history, up to t_*).²

We learn from the above two examples that a subset of an ensemble need not yield reliable quantitative information about the relative frequencies of different outcomes, or even qualitative information about what the allowed outcomes are. All of this information is controlled only by the full ensemble. In both examples, the set of events that occurred before the time $t < t_*$ contain events of type 1-. But in the first example, these events are a sampling artifact and their true probability is actually 0. In the second example, 1- corresponds to a real possibility with non-zero probability.

The cutoff in eternal inflation defines the ensemble Now let us return to eternal inflation. In order to regulate its divergences, we define a cutoff that picks out a finite spacetime region, for example the region prior to some constant light-cone time t . Naively, this seems rather similar to the examples above, where we sampled a large ensemble by considering only the events that happened prior to a time $t < t_*$. But we learned that such samples cannot answer the question of whether the histories of type 1- are real or artifacts. To answer this question, we had to look at the full ensemble. We found in the first example that 1- was real, and in the second that 1- was an artifact, even though the sample looked exactly the same in both cases. In eternal inflation, therefore, we would apparently need to “go beyond the cutoff” and consider the “entire ensemble” of outcomes, in order to decide whether 1- is something that can really happen.

But this is impossible: the whole point of the cutoff was to *define* an ensemble. An infinite set is not a well-defined ensemble, so the set we obtained by imposing a cutoff *is* the most fundamental definition of an ensemble available to us. We can argue about which cutoff is correct: light-cone time, scale factor time, the causal patch, etc. But whatever

²Actually it is faithful only in the limit as t is much greater than the characteristic growth timescale of the civilization, because of the absence of any observers prior to $t = 0$.

the correct cutoff is, its role is to define the ensemble. It cannot be said to select a sample from a larger ensemble, namely from the whole multiverse, because this larger ensemble is infinite, so relative abundances of events are not well-defined. If they were, we would have had no need for a cutoff in the first place.

10.4 The Guth-Vanchurin paradox

Another way to see that the end of time is a real possibility is by verifying that it resolves a paradox exhibited by Guth and Vanchurin [80]. Suppose that before you go to sleep someone flips a fair coin and, depending on the result, sets an alarm clock to awaken you after either a short time, $\Delta t \ll 1$, or a long time $\Delta t \gg 1$. Local physics dictates that there is a 50% probability to sleep for a short time since the coin is fair. Now suppose you have just woken up and have no information about how long you slept. It is natural to consider yourself a typical person waking up. But if we look at everyone who wakes up before the cutoff, we find that there are far more people who wake up after a short nap than a long one. Therefore, upon waking, it seems that there is no longer a 50% probability to have slept for a short time.

How can the probabilities have changed? If you accept that the end of time is a real event that could happen to you, the change in odds is not surprising: although the coin is fair, some people who are put to sleep for a long time never wake up because they run into the end of time first. So upon waking up and discovering that the world has not ended, it is more likely that you have slept for a short time. You have obtained additional information upon waking—the information that time has not stopped—and that changes the probabilities.

However, if you refuse to believe that time can end, there is a contradiction. The odds cannot change unless you obtain additional information. But if all sleepers wake, then the fact that you woke up does not supply you with new information.

Another way to say it is that there are two reference classes one could consider. When going to sleep we could consider all people falling asleep; 50% of these people have alarm clocks set to wake them up after a short time. Upon waking we could consider the class of all people waking up; most of these people slept for a short time. These reference classes can only be inequivalent if some members of one class are not part of the other. This is the case if one admits that some people who fall asleep never wake up, but not if one insists that time cannot end.

10.5 Discussion

Mathematically, the end of time is the statement that our spacetime manifold is extendible, i.e., that it is isometric to a proper subset of another spacetime. Usually, it is assumed that spacetime is inextendable [180]. But the cutoffs we considered regulate eternal inflation by restricting to a subset of the full spacetime. Probabilities are fundamentally defined in terms of the relative abundance of events and histories in the subset. Then the fact that spacetime is extendible is itself a physical feature that can become part of an observer's history. Time can end.

10.5.1 Assumptions

We do not know whether our conclusion is empirically correct. What we have shown is that it follows logically from a certain set of assumptions. If we reject the conclusion, then we must reject at least one of the following propositions:

Probabilities in a finite universe are given by relative frequencies of events or histories This proposition is sometimes called the assumption of typicality. It forces us to assign a nonzero probability to encountering the end of time if a nonzero fraction of observers encounter it.

Even in a finite universe one needs a rule for assigning relative probabilities to observations. This need is obvious if we wish to make predictions for cosmological observations. But a laboratory experiment is a kind of observation, too, albeit one in which the observer controls the boundary conditions. A comprehensive rule for assigning probabilities cannot help but make predictions for laboratory experiments in particular. However, we already have a rule for assigning probabilities in this case, namely quantum mechanics and the Born rule, applied to the local initial conditions prepared in the laboratory. This must be reproduced as a special case by any rule that assigns probabilities to all observations [33]. A simple way to achieve this is by defining probabilities as ratios of the expected number of instances of each outcome in the universe, as we have done in Eq. (10.1).

Probabilities in an infinite universe are defined by a geometric cutoff This proposition states that the infinite spacetime of eternal inflation must be rendered finite so that the above frequency prescription can be used to define probabilities. Moreover, it states that a finite spacetime should be obtained by restricting attention to a finite subset of the infinite multiverse.³ It is possible that the correct measure cannot be expressed in a geometric form. Imagine, for instance, a measure that makes “exceptions” for matter systems that come into existence before the cutoff, allowing all events in their world volume to be counted. A purely geometric prescription would have chopped part of the history off, but in this measure, the cutoff surface would be deformed to contain the entire history of the system. Such a cutoff would depend not only on the geometry, but also on the matter content of spacetime.⁴ A more radical possibility is that the measure may not involve any kind of restriction to a finite portion of spacetime. For example, Noorbala and Vanchurin [135], who exhibit a paradox similar to that described in Sec. 10.4, but do not allow for the possibility that time can end, advocate a nongeometric type of measure. If such a prescription could be made well-defined and consistent with observation (which

³We have considered measures in which the cutoff is completely sharp, i.e., described by a hypersurface that divides the spacetime into a region we keep and a region we discard. In fact this is not essential. One could smear out the cutoff by assigning to each spacetime event a weight that varies smoothly from 1 to 0 over some region near the cutoff surface. There would still be a finite probability for time to end.

⁴We have not attempted to prove this statement, so it should be considered an additional assumption. Because the metric has information about the matter content, we cannot rule out that a geometric measure could be formulated whose cutoff surfaces never intersect with matter. It seems unlikely to us that such a cutoff could select a finite subset of the multiverse. A related possibility would be to define a global time cutoff such that typical observers live farther and farther from the cutoff in the limit as $t \rightarrow \infty$. This would invalidate our analysis in Sec. 10.3.1, which assumed exponential growth in t .

seems unlikely to us), then one might escape the conclusion that time can end. Similarly, Winitzki [182, 183, 184] defined a measure where only finite spacetimes are considered, and in this measure there is no novel catastrophe like the end of time.

The universe is eternally inflating To prove this proposition wrong would be a dramatic result, since it would seem to require a kind of fundamental principle dictating that Nature abhors eternal inflation. After all, eternal inflation is a straightforward consequence of general relativity, assuming there exists at least one sufficiently long-lived field theory vacuum with positive vacuum energy (a de Sitter vacuum). This assumption, in turn, seems innocuous and is well-motivated by observation: (1) The recent discovery of accelerated expansion [151, 143], combined with the conspicuous lack of evidence that dark energy is not a cosmological constant [115], suggests that our own vacuum is de Sitter. If this is the case, the universe must be eternally inflating. (2) Slow-roll inflation solves the horizon and flatness problems. Its generic predictions agree well with the observed CMB power spectrum. But slow-roll inflation requires a sufficiently flat scalar field potential. Eternal inflation requires only a local minimum and so is less fine-tuned. How could we consider slow-roll inflation, but exclude eternal inflation? — There are also theoretical motivations for considering the existence of de Sitter vacua: (3) In effective field theory, there is nothing special about potentials with a positive local minimum, so it would be surprising if they could not occur in Nature. (4) String theory predicts a very large number of long-lived de Sitter vacua [41, 109, 55], allowing for a solution of the cosmological constant problem and other fine-tuning problems.

10.5.2 Observation

If we accept that time can end, what observable implications does this have? Should we expect to see clocks or other objects suddenly disappear? In measures such as scale factor time or light-cone time, the expected lifetime of stable systems is of order 5 billion years right now, so it would be very unlikely for the end of time to occur in, say, the next thousand years. And even if it did occur, it would not be observable. Any observer who would see another system running into the end of time is by definition located to the causal future of that system. If the cutoff surface is everywhere spacelike or null, as is the case for the light-cone time cutoff and the causal patch cutoff, then the observer will necessarily run into the cutoff before observing the demise of any other system.

Though the end of time would not be observable, the fact that time has *not* ended certainly is observable. If a theory assigns extremely small probability to some event, then the observation of this event rules out the theory at a corresponding level of confidence. This applies, in particular, to the case where the event in question is time-not-having-ended. For example, Eq. (10.7) shows that the proper time measure is thus falsified.

An observation which indirectly probes the end of time is the value of the cosmological constant. For definiteness consider the causal patch measure, which predicts a coincidence between the time when the observers live and the time when the cosmological constant begins to dominate the expansion of the universe, $t_\Lambda \sim t_{\text{obs}}$. This represents the most important phenomenological success of the measure, and we will now argue that it is tied intimately to the end of time.

The most likely value of the cosmological constant is the one which leads to the most observers inside the causal patch. We will assume that there are a constant number of observers per unit mass, and will imagine scanning the possible values of $t_\Lambda \sim 1/\sqrt{\Lambda}$ with t_{obs} held fixed. It's most useful to think of the distribution of values of $\log t_\Lambda$, where the preferred value is largely determined by two competing pressures. First, since the prior probability is flat in Λ , there is an exponential pressure in $\log t_\Lambda$ toward lesser values. Second, if $t_\Lambda < t_{\text{obs}}$ there is an exponential pressure in t_Λ (superexponential in $\log t_\Lambda$) toward greater values. This is a simple consequence of the fact that all matter is expelled from the causal patch at an exponential rate after vacuum domination begins. These two pressures lead to $t_{\text{obs}} \sim t_\Lambda$.

The end of time is implicitly present in this argument. Suppose there are two generations of observers, one living at t_Λ and another at $10t_\Lambda$. Even if local physics says that there are the same number of observers per unit mass in each generation, the second generation must be atypical, and hence have fewer members, if the prediction for the cosmological constant is to remain valid. Where are the missing members of the second generation? The answer is that time has ended for them. They are not counted for the purposes of any calculation, and so they do not exist. Clearly, the setup is identical to the observers who see 1 o'clock and 2 o'clock discussed above.

In this chapter, we have considered sharp geometric cutoffs. However, intuition from AdS/CFT [73, 26, 28] suggests that the cutoff should not be a completely sharp surface, but should be smeared out over a time of order t_Λ . If the cutoff is smeared, there could be observable consequences of approaching the end of time; the details would depend on the precise prescription for smearing the cutoff.

10.5.3 Interpretation

The notion that time can come to an end is not completely new. Space and time break down at singularities, which are guaranteed to arise in gravitational collapse [95]. But our conclusion is more radical: the world can come to an end in any spacetime region, including regions with low density and curvature, because spacetime is incomplete.

One might speculate that semiclassical gravity breaks down on very large time scales, say t_Λ^3 , the evaporation time for a large black hole in de Sitter space, or $\exp(\pi t_\Lambda^2)$, the recurrence time in de Sitter space. But in the most popular measures, we are likely to encounter the end of time on the much shorter timescale t_Λ . Perhaps one could invent a new cutoff that would push the end of time further into the future. But there is no well-motivated candidate we are aware of, and, as we have discussed, one would be likely to lose some of the phenomenological success of the measures in solving, e.g., the cosmological constant problem.

How can we make sense of our conclusion? Is there a way of thinking about it that would make us feel more comfortable about the end of time? Does it fit in with something we already know, or is this a completely new phenomenon? The answer to this question turns out to depend somewhat on which cutoff is used.

All measures One way to interpret the end of time is to imagine a computer just powerful enough to simulate the cutoff portion of the eternally inflating spacetime. The simulation simply stops at the cutoff. If the measure involves taking a late time limit, then one can

imagine building larger and larger computers that can simulate the spacetime until a later cutoff. These computers can be thought of as the definition of the cutoff theory, much in the same way that lattice gauge theory is used. There is no physical significance to any time after the cutoff.⁵ This is an interesting rephrasing of the statement of the end of time, but it does not seem to mitigate its radicality.

Causal patch only Our result appears to admit an intriguing interpretation if the causal patch measure is used. The original motivation for the causal patch came from black hole complementarity [170]. Consider the formation and evaporation of a black hole in asymptotically flat space. If this process is unitary, then the quantum state describing the collapsing star inside the black hole is identical to the state of the Hawking radiation cloud. Since these two states are spacelike separated, two copies of the quantum state exist at the same time. But before the star collapsed, there was only one copy. This amounts to “quantum xeroxing”, which is easily seen to conflict with quantum mechanics.

A way around this paradox is to note there is no spacetime point whose past light-cone contains both copies. This means that no experiment consistent with causality can actually verify that xeroxing has taken place. Thus, the paradox can be regarded as an artifact of a global viewpoint that has no operational basis. A theory should be capable of describing all observations, but it need not describe more than that. Geometrically, this means that it need not describe any system that cannot fit within a causal patch. What the xeroxing paradox teaches us is that we *must* not describe anything larger than the causal patch if we wish to avoid inconsistencies in spacetimes with black holes.

But once we reject the global description of spacetime, we must reject it whether or not black holes are present. In many cosmological solutions, including eternal inflation, the global spacetime is not accessible to any single experiment. This motivated the use of the causal patch as a cutoff to regulate the infinities of eternal inflation [32, 24].

Let us return to the black hole spacetime and consider the causal patch of an outside observer. This patch includes all of the spacetime except for the interior of the black hole. As Susskind has emphasized, to an outside observer, the causal patch is a consistent representation of the entire world. The patch has a boundary, the stretched horizon of the black hole. This boundary behaves like a physical membrane, endowed with properties such as temperature, tension, and conductivity. When another observer falls into the black hole, the outside observer would say that he has been thermalized at the horizon and absorbed into the membrane degrees of freedom. Later the membrane evaporates and shrinks away, leaving behind a cloud of radiation.

It is very important to understand that this really is the unique and complete description of the process from the outside point of view; the black hole interior does not come into it. The process is no different, in principle, from throwing the second observer into a fire and watching the smoke come out. Any object is destroyed upon reaching the horizon. Yet, assuming that the black hole is large, the infalling observer would not notice anything special when crossing the horizon. There is no contradiction between these two

⁵Ken Olum has pointed out for some time that one way to interpret a geometric cutoff is that “we are being simulated by an advanced civilization with a large but finite amount of resources, and at some point the simulation will stop.” The above interpretation adopts this viewpoint (minus the advanced civilization).

descriptions, since they agree as long as the two observers remain in causal contact. Once they differ, it is too late for either observer to send a signal to the other and tell a conflicting story.

The end of time in the causal patch is an effect that fits well with the outside observer's description. When the infalling observer enters the black hole, he is leaving the causal patch of the outside observer. In the language of the present chapter, the outside observer defines a particular causal patch, and the inside observer encounters the end of time when he hits the boundary of this patch. We now see that there is a different, more satisfying interpretation: the inside observer is thermalized at the horizon. This interpretation invokes a relatively conventional physical process to explain why the inside observer ceases to exist. Time does not stop, but rather, the observer is thermalized. His degrees of freedom are merged with those already existing at the boundary of the causal patch, the horizon.

If this interpretation is correct, it can be applied to black holes that form in the eternally inflating universe, where it modifies the theory of the infalling observer. It is no longer certain that an infalling observer will actually make it to the horizon, and into the black hole, to perish only once he hits the future singularity. Instead, time might end before he enters the black hole. How is this possible?

In the traditional discussion of black hole complementarity, one picks an observer and constructs the associated causal patch. It is impossible, by construction, for an observer to leave his own patch. In other words, time cannot end if we live in a causal patch centered on our own worldline. In eternal inflation, however, one first picks a causal patch; then one looks for observers in it. Some of these observers will be closer to the boundary and leave the patch sooner than others, who happen to stay in the patch longer. Equivalently, suppose we do want to begin by considering observers of a given type, such as an observer falling towards a black hole. To compute probabilities, we must average over all causal patches that contain such an observer. In some patches the observer will be initially far from the boundary, in others he will hit the boundary very soon. This yields a probability distribution for the rate at which time ends.

Suppose, for example, that we attempted to jump into a black hole of mass m in our own galaxy (and neglect effects of gravitational tidal forces, matter near the black hole, etc.). Using the ensemble of causal patches defined in Ref. [44], one finds that time would probably end before we reach the horizon, with probability $1 - O(m/t_\Lambda)$. This probability is overwhelming if the black hole is much smaller than the cosmological horizon.

Chapter 11

Conclusion

We live in a universe with a positive cosmological constant. In terms of the full effective potential, we are sitting at a local minimum. There is little reason to believe that the effective potential contains only one local minimum. Indeed, string theory tells us that the effective potential contains an enormous number of minima. Therefore, there is strong reason to believe that our vacuum is metastable: we are in a local minimum, but not a global one. From the perspective of some geodesic, our vacuum will eventually decay, nucleating new bubble universes. The nucleated universes will themselves decay, and the process will continue repeating until a universe with negative cosmological constant is created. This particular geodesic will terminate at the big crunch singularity. However there will be other geodesics that followed a different route and that are still in regions that are expanding. In this way, inflation is eternal and gives rise to an infinite multiverse. In this thesis we have attempted to take some steps towards finding a theory describing the multiverse.

As we have discussed, there is strong reason to believe that the complete theory is not a local one. Rather, the theory lives on a codimension 1 surface. In the context of the multiverse, the natural place for the theory to live is on a spacelike equal time slice on future infinity. In the absence of knowledge of what this theory is, we have tried to make progress by trying to ascertain properties of this theory from contexts in which there exists a complete understanding of quantum gravity.

The case in which an in-principle full understanding of string theory is known is for a stable spacetime with a negative cosmological constant (AdS space). For AdS space, the boundary lives on a timelike hypersurface. The complete nonperturbative description of the bulk is given by a conformal field theory living on the boundary. This AdS/CFT correspondence realizes the holographic principle: the idea that spacetime is not fundamental, but rather emergent. In the context of AdS space, we know what the fundamental degrees of freedom are: they are the degrees of freedom of the CFT. In the context of the multiverse we lack this knowledge.

The first task needed for doing computations in the multiverse is to introduce a late time cutoff (a measure). This cutoff is analogous to the large radius IR cutoff in AdS space. In AdS/CFT, the UV/IR relation is generally regarded as providing a relation between the bulk IR cutoff and the CFT cutoff. In chapter 8 we implemented a similar kind of cutoff in the context of the multiverse so as to define a measure which we called new

light-cone time. In chapter 9 we computed some of the predictions this measure gives. The set of quantities that we can at this time compute is limited, as there are few quantities for which we know the prior probability distribution. Nevertheless, the cosmological constant is a notable exception since regardless of the precise form of the effective potential, near any point Taylor expansion should be valid and doing this around values close to 0 gives a flat prior for the cosmological constant. Our results are encouraging, with the measure itself providing an explanation of the coincidence problem (the coincidence between the era when we are living and the onset of vacuum energy domination). It is notable that in the framework of chapter 9, the value of the cosmological constant is not set by the requirement that it be sufficiently small so as to not disrupt galaxy formation [181]. Rather, the effect is due entirely to the measure. At the same time, in chapter 10 we found that all current measures, including ours, face a serious difficulty: the late time cutoff does not decouple. Thus, our successful explanation of the cosmological constant comes at the expense of some other less palatable and more cataclysmic predictions about our future fate.

These difficulties were in part the motivation to reexamine the UV/IR relation. In chapter 2 we found that in fact a UV cutoff CFT would not appear to be sufficient to describe the portion of the bulk AdS out to some radius cutoff. Indeed, any attempt of this kind to decompose the CFT into different tensor factors does not lend itself to a geometric decomposition of the bulk. Perhaps then, this goes some way towards explaining why the analog of UV/IR that we implemented in the multiverse encountered difficulties in chapter 10. For future work, it would therefore seem appropriate to first gain a better handle on constructing a regulated version of AdS/CFT before trying to apply it to the multiverse. One possibility, that we explored to some extent in chapter 2, is to replace UV regularization of the CFT with something else so as to achieve full description of the desired bulk region. Another possibility is to try to understand what it is in the bulk that a UV regulated CFT can describe. In the context of the multiverse, this would translate into using a nongeometric measure.

In the other chapters in Part 1 we tried to make progress in addressing how the bulk is encoded in the CFT. To this effect, in chapters 3 and 4 we studied if it is possible that there is a duality between a subregion of the boundary and a subregion of the bulk. In chapters 5 and 6, we went on to consider how to reconstruct the bulk from boundary data in the form of local operators. It would be interesting to apply these ideas to the multiverse and ask the analogous question of how would one use the data on future infinity to reconstruct the spacetime.

Bibliography

- [1] Ahmed Almheiri and Joseph Polchinski. Models of AdS2 Backreaction and Holography. arXiv:1402.6334, 2014.
- [2] M. Anderson. On uniqueness and differentiability in the space of Yamabe metrics. *Comm. Contemp. Math*, 7:299–310, 2005.
- [3] T. Aubin. Equations differentielles non lineaires et probleme de Yamabe concernant la courbure scalaire. *J. Math. Pures Appl.*, 55:269–296, 1976.
- [4] Vijay Balasubramanian, Borun D. Chowdhury, Bartłomiej Czech, Jan de Boer, and Michal P. Heller. A holographic spacetime. arXiv:1310.4204, 2013.
- [5] Vijay Balasubramanian and Per Kraus. Spacetime and the holographic renormalization group. *Phys. Rev. Lett.*, 83:3605–3608, 1999.
- [6] Tom Banks, Michael R. Douglas, Gary T. Horowitz, and Emil J. Martinec. AdS dynamics from conformal field theory. hep-th/9808016, 1998.
- [7] J.L.F. Barbon and E. Rabinovici. AdS Crunches, CFT Falls And Cosmological Complementarity. *JHEP*, 1104:044, 2011.
- [8] J. M. Bardeen, B. Carter, and S. W. Hawking. The four laws of black hole mechanics. *Commun. Math. Phys.*, 31:161, 1973.
- [9] Claude Bardos, Gilles Lebeau, and Jeffrey Rauch. Sharp Sufficient Conditions for the Observation, Control, and Stabilization of Waves from the Boundary. *SIAM J. Control Optim.*, 30(5):1024–1065, 1992.
- [10] J. D. Bekenstein. Black holes and the second law. *Nuovo Cim. Lett.*, 4:737–740, 1972.
- [11] Jacob D. Bekenstein. Black holes and entropy. *Phys. Rev. D*, 7:2333, 1973.
- [12] Jacob D. Bekenstein. Generalized second law of thermodynamics in black hole physics. *Phys. Rev. D*, 9:3292, 1974.
- [13] Jacob D. Bekenstein. Statistical black-hole thermodynamics. *Phys. Rev. D*, 12:3077, 1975.
- [14] Jacob D. Bekenstein. A universal upper bound on the entropy to energy ratio for bounded systems. *Phys. Rev. D*, 23:287, 1981.

- [15] Jacob D. Bekenstein. Do we understand black hole entropy? 1994. Presented at the 7th Marcel Grossmann Meeting on General Relativity.
- [16] Iosif Bena. On the construction of local fields in the bulk of AdS(5) and other spaces. *Phys.Rev.*, D62:066007, 2000.
- [17] Eugenio Bianchi and Robert C. Myers. On the Architecture of Spacetime Geometry. 2012. arXiv:1212.5183.
- [18] J.J Bisognano and E.H. Wichmann. On the Duality Condition for a Hermitian Scalar Field. *J.Math.Phys.*, 16:985–1007, 1975.
- [19] W. Boucher and G. W. Gibbons. Cosmic Baldness. *Presented at 1982 Nufeld Workshop on the Very Early Universe, Cambridge, England, Jun 21 - Jul 9, 1982.*
- [20] Raphael Bousso. A covariant entropy conjecture. *JHEP*, 07:004, 1999.
- [21] Raphael Bousso. A covariant entropy conjecture. *JHEP*, 07:004, 1999.
- [22] Raphael Bousso. Holography in general space-times. *JHEP*, 06:028, 1999.
- [23] Raphael Bousso. The holographic principle. *Rev. Mod. Phys.*, 74:825, 2002.
- [24] Raphael Bousso. Holographic probabilities in eternal inflation. *Phys. Rev. Lett.*, 97:191302, 2006.
- [25] Raphael Bousso. TASI Lectures on the Cosmological Constant. *Gen. Rel. Grav.*, 40:607–637, 2008.
- [26] Raphael Bousso. Complementarity in the Multiverse. *Phys. Rev.*, D79:123524, 2009.
- [27] Raphael Bousso and Ben Freivogel. A paradox in the global description of the multiverse. *JHEP*, 06:018, 2007.
- [28] Raphael Bousso, Ben Freivogel, Stefan Leichenauer, and Vladimir Rosenhaus. Boundary definition of a multiverse measure. *Phys.Rev.*, D82:125032, 2010.
- [29] Raphael Bousso, Ben Freivogel, Stefan Leichenauer, and Vladimir Rosenhaus. A geometric solution to the coincidence problem, and the size of the landscape as the origin of hierarchy. *Phys. Rev. Lett.*, 106:101301, 2011.
- [30] Raphael Bousso, Ben Freivogel, Stefan Leichenauer, and Vladimir Rosenhaus. Eternal inflation predicts that time will end. *Phys. Rev.*, D83:023525, 2011.
- [31] Raphael Bousso, Ben Freivogel, Stefan Leichenauer, Vladimir Rosenhaus, and Claire Zukowski. Null Geodesics, Local CFT Operators and AdS/CFT for Subregions. *Phys.Rev.*, D88:064057, 2013.
- [32] Raphael Bousso, Ben Freivogel, and I-Sheng Yang. Eternal inflation: The inside story. *Phys. Rev. D*, 74:103516, 2006.

- [33] Raphael Bousso, Ben Freivogel, and I-Sheng Yang. Boltzmann babies in the proper time measure. *Phys. Rev.*, D77:103514, 2008.
- [34] Raphael Bousso, Ben Freivogel, and I-Sheng Yang. Properties of the scale factor measure. *Phys.Rev.*, D79:063513, 2009.
- [35] Raphael Bousso, Lawrence J. Hall, and Yasunori Nomura. Multiverse Understanding of Cosmological Coincidences. *Phys. Rev.*, D80:063510, 2009.
- [36] Raphael Bousso and Roni Harnik. The Entropic Landscape. *Phys. Rev.*, D82:123523, 2010.
- [37] Raphael Bousso, Roni Harnik, Graham D. Kribs, and Gilad Perez. Predicting the cosmological constant from the causal entropic principle. *Phys. Rev. D*, 76:043513, 2007.
- [38] Raphael Bousso and Stefan Leichenauer. Star Formation in the Multiverse. *Phys.Rev.*, D79:063506, 2009.
- [39] Raphael Bousso and Stefan Leichenauer. Predictions from Star Formation in the Multiverse. *Phys.Rev.*, D81:063524, 2010.
- [40] Raphael Bousso, Stefan Leichenauer, and Vladimir Rosenhaus. Light-sheets and AdS/CFT. *Phys.Rev.*, D86:046009, 2012.
- [41] Raphael Bousso and Joseph Polchinski. Quantization of four-form fluxes and dynamical neutralization of the cosmological constant. *JHEP*, 06:006, 2000.
- [42] Raphael Bousso and Lisa Randall. Holographic domains of Anti-de Sitter space. *JHEP*, 04:057, 2002.
- [43] Raphael Bousso and I-Sheng Yang. Landscape predictions from cosmological vacuum selection. *Phys. Rev. D*, 75:123520, 2007.
- [44] Raphael Bousso and I-Sheng Yang. Global-Local Duality in Eternal Inflation. *Phys. Rev.*, D80:124024, 2009.
- [45] Brandon Bozek, Andreas Albrecht, and Daniel Phillips. Curvature Constraints from the Causal Entropic Principle. *Phys. Rev.*, D80:023527, 2009.
- [46] Pasquale Calabrese and John L. Cardy. Entanglement entropy and quantum field theory. *J.Stat.Mech.*, 0406:P06002, 2004.
- [47] Jr. Callan, Curtis G. and Frank Wilczek. On geometric entropy. *Phys.Lett.*, B333:55–61, 1994.
- [48] Horacio Casini, Marina Huerta, and Robert C. Myers. Towards a derivation of holographic entanglement entropy. *JHEP*, 05:036, 2011.
- [49] S. Coleman and F. De Luccia. Gravitational effects on and of vacuum decay. *Phys. Rev. D*, 21:3305–3315, 1980.

- [50] Bartłomiej Czech, Joanna L. Karczmarek, Fernando Nogueira, and Mark Van Raamsdonk. Rindler Quantum Gravity. *Class.Quant.Grav.*, 29:235025, 2012.
- [51] Bartłomiej Czech, Joanna L. Karczmarek, Fernando Nogueira, and Mark Van Raamsdonk. The Gravity Dual of a Density Matrix. *Class.Quant.Grav.*, 29:155009, 2012.
- [52] Jan de Boer, Erik P. Verlinde, and Herman L. Verlinde. On the holographic renormalization group. *JHEP*, 08:003, 2000.
- [53] Andrea De Simone, Alan H. Guth, Michael P. Salem, and Alexander Vilenkin. Predicting the cosmological constant with the scale-factor cutoff measure. *Phys.Rev.*, D78:063520, 2008.
- [54] Andrea De Simone and Michael P. Salem. The distribution of Ω_k from the scale-factor cutoff measure. *Phys.Rev.*, D81:083527, 2010.
- [55] Frederik Denef and Michael R. Douglas. Distributions of flux vacua. *JHEP*, 05:072, 2004.
- [56] Freeman Dyson. Time without end: Physics and biology in an open universe. *Rev. Mod. Phys.*, 51:447–460, 1979.
- [57] Roberto Emparan. AdS / CFT duals of topological black holes and the entropy of zero energy states. *JHEP*, 9906:036, 1999.
- [58] Netta Engelhardt and Aron C. Wall. Extremal Surface Barriers. 2013. arXiv:1312.3699.
- [59] Andrew M. Essin and David J. Griffiths. Quantum mechanics of the $1/x^2$ potential. *American Journal of Physics*, 74:109, 2006.
- [60] Thomas Faulkner, Monica Guica, Thomas Hartman, Robert C. Myers, and Mark Van Raamsdonk. Gravitation from Entanglement in Holographic CFTs. *JHEP*, 1403:051, 2014.
- [61] Thomas Faulkner, Hong Liu, and Mukund Rangamani. Integrating out geometry: Holographic Wilsonian RG and the membrane paradigm. *JHEP*, 08:051, 2011.
- [62] Brian Feldstein, Lawrence J. Hall, and Taizan Watari. Density perturbations and the cosmological constant from inflationary landscapes. *Phys. Rev. D*, 72:123506, 2005.
- [63] Brian Feldstein, Lawrence J. Hall, and Taizan Watari. Landscape Predictions for the Higgs Boson and Top Quark Masses. *Phys. Rev.*, D74:095011, 2006.
- [64] A.Liam Fitzpatrick and Jared Kaplan. Scattering States in AdS/CFT. 2011. arXiv:1104.2597.
- [65] Eanna E. Flanagan, Donald Marolf, and Robert M. Wald. Proof of classical versions of the Bousso entropy bound and of the Generalized Second Law. *Phys. Rev. D*, 62:084035, 2000.

- [66] Ben Freivogel, Steven B. Giddings, and Matthew Lippert. Toward a theory of precursors. *Phys.Rev.*, D66:106002, 2002.
- [67] Ben Freivogel, Matthew Kleban, Maria Rodriguez Martinez, and Leonard Susskind. Observational consequences of a landscape. *JHEP*, 03:039, 2006.
- [68] Mitsutoshi Fujita, Tadashi Takayanagi, and Erik Tonni. Aspects of AdS/BCFT. *JHEP*, 11:043, 2011.
- [69] Sijie Gao and Robert M. Wald. Theorems on gravitational time delay and related issues. *Class. Quant. Grav.*, 17:4999–5008, 2000.
- [70] Juan García-Bellido and Andrei Linde. Stationary solutions in Brans-Dicke stochastic inflationary cosmology. *Phys. Rev. D*, 52:6730–6738, 1995.
- [71] Juan García-Bellido, Andrei Linde, and Dmitri Linde. Fluctuations of the gravitational constant in the inflationary Brans-Dicke cosmology. *Phys. Rev. D*, 50:730–750, 1994.
- [72] Juan García-Bellido and Andrei D. Linde. Stationarity of inflation and predictions of quantum cosmology. *Phys. Rev.*, D51:429–443, 1995.
- [73] J. Garriga and A. Vilenkin. Holographic Multiverse. *JCAP*, 0901:021, 2009.
- [74] J. Garriga and A. Vilenkin. Holographic Multiverse. *JCAP*, 0901:021, 2009.
- [75] Jaume Garriga. Holographic description of vacuum bubbles. *Prog. Theor. Phys. Suppl.*, 190:261–273, 2011.
- [76] Jaume Garriga, Delia Schwartz-Perlov, Alexander Vilenkin, and Sergei Winitzki. Probabilities in the inflationary multiverse. *JCAP*, 0601:017, 2006.
- [77] Jaume Garriga and Alexander Vilenkin. Anthropic prediction for Lambda and the Q catastrophe. *Prog. Theor. Phys. Suppl.*, 163:245–257, 2006.
- [78] Steven B. Giddings. The Boundary S matrix and the AdS to CFT dictionary. *Phys.Rev.Lett.*, 83:2707–2710, 1999.
- [79] S. S. Gubser, I. R. Klebanov, and A. M. Polyakov. Gauge theory correlators from noncritical string theory. *Phys. Lett.*, B428:105, 1998.
- [80] Alan Guth and Vitaly Vanchurin. Private communication.
- [81] Alan H. Guth. Inflation and eternal inflation. *Phys. Rep.*, 333:555, 1983.
- [82] Alan H. Guth. Inflationary models and connections to particle physics. 2000. astro-ph/0002188.
- [83] Alan H. Guth. Inflation. 2004. astro-ph/0404546.
- [84] Alan H. Guth. Eternal inflation and its implications. *J. Phys.*, A40:6811–6826, 2007.

- [85] Alan H. Guth and Vitaly Vanchurin. Eternal Inflation, Global Time Cutoff Measures, and a Probability Paradox. 2011. arXiv:1108.0665.
- [86] Alan H. Guth and Erick J. Weinberg. Could the universe have recovered from a slow first-order phase transition? *Nucl. Phys.*, B212:321–364, 1983.
- [87] Lawrence J. Hall and Yasunori Nomura. Evidence for the Multiverse in the Standard Model and Beyond. *Phys.Rev.*, D78:035001, 2008.
- [88] Lawrence J. Hall and Yasunori Nomura. Evidence for the Multiverse in the Standard Model and Beyond. *Phys. Rev.*, D78:035001, 2008.
- [89] Alex Hamilton, Daniel N. Kabat, Gilad Lifschytz, and David A. Lowe. Holographic representation of local bulk operators. *Phys.Rev.*, D74:066009, 2006.
- [90] Alex Hamilton, Daniel N. Kabat, Gilad Lifschytz, and David A. Lowe. Local bulk operators in AdS/CFT: A Boundary view of horizons and locality. *Phys.Rev.*, D73:086003, 2006.
- [91] Daniel Harlow and Douglas Stanford. Operator Dictionaries and Wave Functions in AdS/CFT and dS/CFT. 2011. arXiv:1104.2621.
- [92] S. W. Hawking. Gravitational radiation from colliding black holes. *Phys. Rev. Lett.*, 26:1344–1346, 1971.
- [93] S. W. Hawking. Black hole explosions. *Nature*, 248:30–31, 1974.
- [94] S. W. Hawking. Particle creation by black holes. *Commun. Math. Phys.*, 43:199, 1975.
- [95] S. W. Hawking and G. F. R. Ellis. *The large scale structure of space-time*. Cambridge University Press, Cambridge, England, 1973.
- [96] Idse Heemskerk, Donald Marolf, Joseph Polchinski, and James Sully. Bulk and Transhorizon Measurements in AdS/CFT. *JHEP*, 1210:165, 2012.
- [97] Idse Heemskerk, Donald Marolf, Joseph Polchinski, and James Sully. Bulk and Transhorizon Measurements in AdS/CFT. *JHEP*, 1210:165, 2012.
- [98] Idse Heemskerk and Joseph Polchinski. Holographic and Wilsonian Renormalization Groups. *JHEP*, 06:031, 2011.
- [99] Gary T. Horowitz and N. Itzhaki. Black holes, shock waves, and causality in the AdS/CFT correspondence. *JHEP*, 02:010, 1999.
- [100] Veronika E. Hubeny. Precursors seen inside black holes. *Int. J. Mod. Phys.*, D12:1693–1698, 2003.
- [101] Veronika E. Hubeny. Extremal surfaces as bulk probes in AdS/CFT. 2012.
- [102] Veronika E. Hubeny and Henry Maxfield. Holographic probes of collapsing black holes. 2013. arXiv:1312.6887.

- [103] Veronika E. Hubeny and Mukund Rangamani. Causal Holographic Information. *JHEP*, 1206:114, 2012.
- [104] Veronika E. Hubeny, Mukund Rangamani, and Tadashi Takayanagi. A covariant holographic entanglement entropy proposal. *JHEP*, 07:062, 2007.
- [105] Ling-Yan Hung, Robert C. Myers, and Michael Smolkin. Some Calculable Contributions to Holographic Entanglement Entropy. *JHEP*, 1108:039, 2011.
- [106] Daniel Kabat and Gilad Lifschytz. Decoding the hologram: Scalar fields interacting with gravity. 2013. arXiv:1311.3020.
- [107] Daniel Kabat, Gilad Lifschytz, and David A. Lowe. Constructing local bulk observables in interacting AdS/CFT. *Phys.Rev.*, D83:106009, 2011.
- [108] Daniel N. Kabat and M.J. Strassler. A Comment on entropy and area. *Phys.Lett.*, B329:46–52, 1994.
- [109] Shamit Kachru, Renata Kallosh, Andrei Linde, and Sandip P. Trivedi. De Sitter vacua in string theory. *Phys. Rev. D*, 68:046005, 2003.
- [110] Cynthia Keeler, Gino Knodel, and James T. Liu. What do non-relativistic CFTs tell us about Lifshitz spacetimes? *JHEP*, 1401:062, 2014.
- [111] William R. Kelly and Aron C. Wall. Coarse-grained entropy and causal holographic information in AdS/CFT. 2013. arXiv:1309.3610.
- [112] Esko Keski-Vakkuri. Bulk and boundary dynamics in BTZ black holes. *Phys.Rev.*, D59:104001, 1999.
- [113] Alexei Kitaev and John Preskill. Topological entanglement entropy. *Phys.Rev.Lett.*, 96:110404, 2006.
- [114] Igor R. Klebanov, David Kutasov, and Arvind Murugan. Entanglement as a probe of confinement. *Nucl.Phys.*, B796:274–293, 2008.
- [115] E. Komatsu et al. Five-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation. *Astrophys.J.Suppl.*, 180:330–376, 2009.
- [116] Nima Lashkari, Michael B. McDermott, and Mark Van Raamsdonk. Gravitational Dynamics From Entanglement "Thermodynamics". 2013. arXiv:1308.3716.
- [117] J.M. Lee and T. H. Parker. The Yamabe Problem. *Bull Amer. Math Soc.*, 17:37–81, 1987.
- [118] Stefan Leichenauer and Vladimir Rosenhaus. AdS black holes, the bulk-boundary dictionary, and smearing functions. *Phys.Rev.*, D88(2):026003, 2013.
- [119] Aitor Lewkowycz and Juan Maldacena. Generalized gravitational entropy. *JHEP*, 1308:090, 2013.

- [120] A. Linde. Eternally existing selfreproducing chaotic inflationary universe. *Phys. Lett.*, B175:395–400, 1986.
- [121] Andrei Linde. Sinks in the landscape, Boltzmann Brains, and the cosmological constant problem. *JCAP*, 0701:022, 2007.
- [122] Andrei Linde. Towards a gauge invariant volume-weighted probability measure for eternal inflation. *JCAP*, 0706:017, 2007.
- [123] Andrei Linde, Dmitri Linde, and Arthur Mezhlumian. From the big bang theory to the theory of a stationary universe. *Phys. Rev. D*, 49:1783–1826, 1994.
- [124] Andrei Linde, Dmitri Linde, and Arthur Mezhlumian. Nonperturbative amplifications of inhomogeneities in a self-reproducing universe. *Phys. Rev. D*, 54:2504–2518, 1996.
- [125] Andrei Linde and Arthur Mezhlumian. Stationary universe. *Phys. Lett.*, B307:25–33, 1993.
- [126] Andrei Linde, Vitaly Vanchurin, and Sergei Winitzki. Stationary Measure in the Multiverse. *JCAP*, 0901:031, 2009.
- [127] Juan Maldacena. The large N limit of superconformal field theories and supergravity. *Adv. Theor. Math. Phys.*, 2:231, 1998.
- [128] Juan Maldacena. Vacuum decay into Anti de Sitter space. 2010. arXiv:1012.0274.
- [129] Juan Martin Maldacena. Wilson loops in large N field theories. *Phys.Rev.Lett.*, 80:4859–4862, 1998.
- [130] Juan Martin Maldacena. Eternal black holes in anti-de Sitter. *JHEP*, 0304:021, 2003.
- [131] Juan Martin Maldacena. Non-Gaussian features of primordial fluctuations in single field inflationary models. *JHEP*, 0305:013, 2003.
- [132] Juan Martin Maldacena, Jeremy Michelson, and Andrew Strominger. Anti-de Sitter fragmentation. *JHEP*, 9902:011, 1999.
- [133] John March-Russell and Francesco Riva. Signals of Inflation in a Friendly String Landscape. *JHEP*, 0607:033, 2006.
- [134] Donald Marolf. States and boundary terms: Subtleties of Lorentzian AdS/CFT. *JHEP*, 05:042, 2005.
- [135] Mahdiyar Noorbala and Vitaly Vanchurin. Geocentric cosmology: a new look at the measure problem. 2010. arXiv:1006.4148.
- [136] Lukas Novotny and Bert Hecht. *Principles of Nano-Optics*. Cambridge University Press, Cambridge, 2012.
- [137] Ken D. Olum and Delia Schwartz-Perlov. Anthropic prediction in a large toy landscape. *JCAP*, 0710:010, 2007.

- [138] Ken. Olum. Private Communication.
- [139] H. Osborn and A.C. Petkou. Implications of conformal invariance in field theories for general dimensions. *Annals Phys.*, 231:311–362, 1994.
- [140] Don N. Page. Return of the Boltzmann brains. 2006.
- [141] Don N. Page. Cosmological Measures without Volume Weighting. *JCAP*, 0810:025, 2008.
- [142] Eric Perlmutter. A universal feature of CFT Renyi entropy. 2013. arXiv:1308.1083.
- [143] S. Perlmutter et al. Measurements of Omega and Lambda from 42 high-redshift supernovae. *Astrophys. J.*, 517:565–586, 1999.
- [144] Daniel Phillips and Andreas Albrecht. Effects of Inhomogeneity on the Causal Entropic prediction of Lambda. *Phys.Rev.*, D84:123530, 2011.
- [145] Levon Pogosian and Alexander Vilenkin. Anthropic predictions for vacuum energy and neutrino masses in the light of WMAP-3. *JCAP*, 0701:025, 2007.
- [146] Joseph Polchinski. S matrices from AdS space-time. 1999. hep-th/9901076.
- [147] Joseph Polchinski. The cosmological constant and the string landscape. 2006. hep-th/0603249.
- [148] Joseph Polchinski, Leonard Susskind, and Nicolaos Toumbas. Negative energy, superluminality and holography. *Phys.Rev.*, D60:084006, 1999.
- [149] Soo-Jong Rey and Vladimir Rosenhaus. Scanning Tunneling Macroscopy, Black Holes, and AdS/CFT Bulk Locality. 2014. arXiv:1403.3943.
- [150] Soo-Jong Rey and Jung-Tay Yee. Macroscopic strings as heavy quarks in large N gauge theory and anti-de Sitter supergravity. *Eur.Phys.J.*, C22:379–394, 2001.
- [151] Adam G. Riess et al. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *Astron. J.*, 116:1009–1038, 1998.
- [152] Vladimir Rosenhaus. Holography for a small world. 2013. arXiv:1309.4526.
- [153] Vladimir Rosenhaus and Michael Smolkin. Entanglement Entropy: A Perturbative Calculation. 2014. arXiv:1403.3733.
- [154] Shinsei Ryu and Tadashi Takayanagi. Holographic derivation of entanglement entropy from AdS/CFT. *Phys. Rev. Lett.*, 96:181602, 2006.
- [155] Michael P. Salem. Negative vacuum energy densities and the causal diamond measure. *Phys.Rev.*, D80:023502, 2009.
- [156] R. Schoen. Conformal deformation of a Riemannian metric to constant scalar curvature. *J. Diff. Geom.*, 20:479–495, 1984.

- [157] Delia Schwartz-Perlov. Probabilities in the Arkani-Hamed-Dimopolous-Kachru landscape. *J.Phys.*, A40:7363–7374, 2007.
- [158] Delia Schwartz-Perlov. Anthropic prediction for a large multi-jump landscape. *JCAP*, 0810:009, 2008.
- [159] Delia Schwartz-Perlov and Alexander Vilenkin. Probabilities in the Bousso-Polchinski multiverse. *JCAP*, 0606:010, 2006.
- [160] Delia Schwartz-Perlov and Alexander Vilenkin. Measures for a Transdimensional Multiverse. *JCAP*, 1006:024, 2010.
- [161] Sergey N. Solodukhin. Entanglement entropy, conformal invariance and extrinsic geometry. *Phys.Lett.*, B665:305–309, 2008.
- [162] Dam T. Son and Andrei O. Starinets. Minkowski space correlators in AdS / CFT correspondence: Recipe and applications. *JHEP*, 0209:042, 2002.
- [163] Andrew Strominger. The dS/CFT correspondence. *JHEP*, 10:034, 2001.
- [164] Andrew Strominger. Inflation and the dS/CFT correspondence. *JHEP*, 11:049, 2001.
- [165] L. Susskind and E. Witten. The holographic bound in Anti-de Sitter space. 1998.
- [166] Leonard Susskind. The world as a hologram. *J. Math. Phys.*, 36:6377–6396, 1995.
- [167] Leonard Susskind. Holography in the flat space limit. pages 98–112, 1998. Presented at General relativity and relativistic astrophysics. Proceedings: 8th Canadian Conference, Montreal.
- [168] Leonard Susskind. The anthropic landscape of string theory. 2003. hep-th/0302219.
- [169] Leonard Susskind and Larus Thorlacius. Gedanken experiments involving black holes. *Phys. Rev.*, D49:966–974, 1994.
- [170] Leonard Susskind, Larus Thorlacius, and John Uglum. The stretched horizon and black hole complementarity. *Phys. Rev. D*, 48:3743, 1993.
- [171] Leonard Susskind and Nicolaos Toumbas. Wilson loops as precursors. *Phys.Rev.*, D61:044001, 2000.
- [172] G. 't Hooft. Dimensional reduction in quantum gravity. 1993. gr-qc/9310026.
- [173] Tadashi Takayanagi. Holographic Dual of BCFT. *Phys. Rev. Lett.*, 107:101602, 2011.
- [174] Daniel Tataru. Unique continuation problems for partial differential equations. In *Geometric Methods in Inverse Problems and PDE Control*, volume 137 of *The IMA Volumes in Mathematics and its Applications*. Springer, 2004.
- [175] Max Tegmark. What does inflation really predict? *JCAP*, 0504:001, 2005.

- [176] N. Trudinger. Remarks concerning the conformal deformation of Riemannian structures on compact manifolds. *Ann. Scuola Norm. Sup. Pisa*, 22:265–274, 1968.
- [177] William G. Unruh and Nathan Weiss. Acceleration Radiation in Interacting Field Theories. *Phys.Rev.*, D29:1656, 1984.
- [178] Vitaly Vanchurin. Geodesic measures of the landscape. *Phys. Rev. D*, 75:023524, 2007.
- [179] Vitaly Vanchurin and Alexander Vilenkin. Eternal observers and bubble abundances in the landscape. *Phys.Rev.*, D74:043520, 2006.
- [180] Robert M. Wald. *General Relativity*. The University of Chicago Press, Chicago, 1984.
- [181] Steven Weinberg. Anthropic bound on the cosmological constant. *Phys. Rev. Lett.*, 59:2607, 1987.
- [182] Sergei Winitzki. A volume-weighted measure for eternal inflation. *Phys. Rev.*, D78:043501, 2008.
- [183] Sergei Winitzki. Reheating-volume measure for random-walk inflation. *Phys. Rev.*, D78:063517, 2008.
- [184] Sergei Winitzki. Reheating-volume measure in the landscape. *Phys. Rev.*, D78:123518, 2008.
- [185] Edward Witten. Anti-de Sitter space and holography. *Adv. Theor. Math. Phys.*, 2:253, 1998.
- [186] Edward Witten. Quantum gravity in de Sitter space. 2001. hep-th/0106109.
- [187] H. Yamabe. On a deformation of Riemannian structures on compact manifolds. *Oski Math J.*, 12:21–37, 1960.