Quantum states: how entangled are they, given a bipartition?

A) B

Consider a pure ground state $|\Psi\rangle = \sum_{i,\mu} C_{i\mu} |i\rangle_A \otimes |\mu\rangle_B$

where $\mathcal{H}_A = \{|i\rangle_A, i=1,\dots d_A\}$, $\mathcal{H}_B = \{|\mu\rangle_B, \mu=1,\dots d_B\}$, $|i\rangle$ & $|\mu\rangle$ are orthonormal basis for $\mathcal{H}_A$ and $\mathcal{H}_B$, $C_{i\mu} = d_A \times d_B$ complex matrix with $\mathbb{C}$ entries

* There are 2 different types of states, depending on $C_{i\mu}$:
  1) Separable states;
  2) Entangled states.

Recall, a density matrix is a list of probabilities for the occurrence of various states in $\mathcal{H}_A$, $\mathcal{H}_B$!

① Separable states: $C_{i\mu}$ factorizes into $C_i^A \cdot C_\mu^B$, and so our state (also known as product states) becomes $|\Psi\rangle = |\Psi_A\rangle \otimes |\Psi_B\rangle$, with

$$|\Psi_A\rangle = \sum_i C_i^A |i\rangle_A, \quad |\Psi_B\rangle = \sum_\mu C_\mu^B |\mu\rangle_B.$$

In this case, the reduced density matrix also becomes pure, $\rho_A = |\Psi_A\rangle\langle\Psi_A|$. When we calculate the entanglement entropy, we'll see $S_A = 0$. This means that there are no quantum correlations across A and B, and not knowing about B doesn't affect our knowledge of A.

② Entangled states: $C_{i\mu} \neq C_i^A C_\mu^B$

We can simplify $|\Psi\rangle = \sum_{i,\mu} C_{i\mu} |i\rangle_A \otimes |\mu\rangle_B$ by changing bases into a Schmidt decomposition form, which makes correlations between A and B manifest,

$$|\Psi\rangle = \sum_{k=1}^{r} \sqrt{p_k}\, |\Psi_k\rangle_A \otimes |\Psi_k\rangle_B, \quad \text{w/ } \sum_k p_k = 1 \text{ probabilities!}$$

$r = \min(d_A, d_B)$

$|\Psi_k\rangle_A$ and $|\Psi_k\rangle_B$ are new orthonormal bases for subsystems A and B, and it works for any $d_A \times d_B$ rectangular matrix $C_{i\mu}$. The decomposition is nice because it gives us a mixed state density matrix w/ probability distribution $\{p_k\}$

$$\rho_A = \sum_{i=1}^{d_B} {}_B\langle\Psi_i|\Psi\rangle\langle\Psi|\Psi_i\rangle_B = \sum_{k=1}^{r} p_k |\Psi_k\rangle\langle\Psi_k|_A.$$

From here it's also clear that $S_A = S_B$ since $\sum_{k=1}^{r}$ (→ same) and p_k go into both ρ's

In summary, an entangled state is a superposition of several quantum states. An observer who can access only a subsystem A will find themselves in a mixed state when the pure ground state $|\psi\rangle$ in the total system is entangled.

Recall, given a bipartition $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$, we can construct an operator that acts only on one side of the system, by tracing out the other, for example, the reduced density ρ.

$$\sum_i {}_A\langle i | \rho_{TOT} | i \rangle_A = \rho_B = Tr_A \overbrace{|\psi\rangle\langle\psi|}^{\rho_{TOT}}$$

In short:

$|\psi\rangle$ separable → ρ_A pure state

$|\psi\rangle$ entangled → ρ_A mixed state

⊛ Here we started w/ a pure state and looked at the structure of its coefficients. Is that always possible? What if we start w/ a mixed state?

ex. 2-spin system. $\mathcal{H}_{A,B} = \{ |0\rangle_{A,B}, |1\rangle_{A,B} \}$, $|ij\rangle = |i\rangle_A \otimes |j\rangle_B$

$|\psi\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$, $\rho_A = tr_B |\psi\rangle\langle\psi| = \frac{1}{2}\begin{pmatrix} 1 & \\ & 1 \end{pmatrix}$

$\rightarrow S_A = -tr \rho \ln \rho = -\sum_{i=0}^{1} \lambda_i \ln \lambda_i = \ln 2.$

⇒ Entanglement purification

Once the complete information ρ_A of a system A is given, is it possible to find a pure ρ in an enlarged Hilbert space of $\mathcal{H}_A$ whose partial trace recovers ρ_A? → Yes, one can always construct such an enlarged Hilbert space!

($\mathcal{H}$-space)

The most general form of ρ is $\rho_A = \sum_i p_i |i\rangle\langle i|_A$, where $\{|i\rangle_A\}$ is the orthonormal basis of $\mathcal{H}_A$, and $\sum_i p_i = 1$. We can then copy $\mathcal{H}_A$ into another $\mathcal{H}$-space $\mathcal{H}_{\bar{A}}$ w/ the basis given by $\{|i\rangle_{\bar{A}}\}$ and define a pure ρ by

$$\rho = |\chi\rangle\langle\chi| , \quad |\chi\rangle = \sum_i \sqrt{p_i}\, |i\rangle_A \otimes |i\rangle_{\bar{A}}$$

in the enlarged $\mathcal{H}$-space $\mathcal{H}_A \otimes \mathcal{H}_{\bar{A}}$

So, entanglement purification is a construction of a pure, ~~unentangled~~ state from a mixed entangled state by enlarging the Hilbert space. But basically, a trick to think about mixed states.

This is a generic construction of a pure state w/ some p_i. There is a special case which is important in QG known as the thermofield double state, $|TFD\rangle$,

$$|TFD\rangle = \frac{1}{\sqrt{Z}} \sum_n e^{-\beta E_n/2} |n_A\rangle \otimes |n\rangle_B \, ,$$

where we normalized the state w/ a partition function $Z = \sum_n e^{-\beta E_n}$. Why is this state special: taking a partial trace over A or B, we obtain a Gibbs state (has equilibrium probability distribution),

$$\rho_A = \frac{1}{Z} \sum_n e^{-\beta E_n} |n\rangle\langle n|_A = \frac{1}{Z} e^{-\beta H_A} \, .$$

where H_A is the modular Hamiltonian w/ $H_A |n\rangle_A = E_n |n\rangle_A$ (these are energy eigenstates).

In other words, $|TFD\rangle$ is the entanglement purification of a thermal state w/ Boltzmann weight $p_i = \frac{1}{Z} e^{-\beta E_i}$. And the von Neumann entropy measures the thermal entropy of subsystem A,

$$S_A = -\mathrm{tr}\, \rho_A \left(-\beta H_A - \log Z \right) = \beta \left(\langle H_A \rangle - F \right),$$

w/ $\beta F = -\log Z$, and we used $\rho_A = \frac{1}{Z} e^{-\beta H_A}$ and the fact that $\mathrm{tr}(\rho_A H_A) = \langle H_A \rangle$. An important state in black hole physics!

States in QFT

Having understood a spatial bipartition of a discrete lattice system, can get to the continuum by sending $\epsilon \to 0$ (lattice spacing).

Start w/ some d-dimensional relativistic QFT on some Lorentzian spacetime B which is globally hyperbolic (which means we can pick a well-defined time-slice). For instance, we can pick Minkowski, $B^d_{MINK} = \mathbb{R}^{1,d-1}$, and a Cauchy slice Σ_{d-1}, which is an achronal spacelike slice. We define our state on this time slice.

Can also talk about bipartitions: $A \cup B = \Sigma_{d-1}$, and the boundary of region A is an "entangling surface" ∂A, codim-2.

* All notions of QM hold still here: it makes sense to talk about states, density matrices, entropies etc. However, we also encounter 2 problems when we go to the QFT: UV divergences & non factorization due to gauge fields.

UV divergences

Why are they here? Because the vacuum already has a non-trivial entanglement structure. Recall the 2ptf of a free massive scalar at spatially separated points x and y,

some Bessel function

$$\langle \Omega | \phi(0,x) \phi(0,y) | \Omega \rangle = \frac{1}{4\pi^2} \frac{m}{|x-y|} K_1(m|x-y|)$$

Limits w.r.t. the correlation length m:

$$\begin{cases} |x-y| \ll m^{-1} \Rightarrow 2\text{ptf} \sim \frac{1}{|x-y|^2} \longrightarrow \text{diverges @ short distances!} \\ |x-y| \gg m^{-1} \Rightarrow 2\text{ptf} \sim e^{-m|x-y|} \end{cases}$$

This is an example (free massive scalar), but this power-law divergence and long-distance decay is expected to be true in any relativistic QFT (this condition is known as the Hadamard state condition).

Note that this is non-trivial correlation in the <u>ground state</u> of a QFT: non-interacting, non-relativistic gas of massive particles will <u>not</u> have this property! Loosely speaking, the D $\sim$ # of EPR pairs across ∂A.

One way to interpret this correlation is as an illustration of <u>entanglement</u> of different regions in the <u>vacuum</u> of a relativistic QFT.

Another way to see the entangled nature of the vacuum: <u>Reeh-Schlieder theorem</u>: for any region, by acting on $|\Omega\rangle$ w/ operators located in that region, we can produce a set of states which is dense in the full Hilbert space of the QFT.

In other words, by acting on the vacuum of the QFT w/ some operators located in this classroom, we can create the Moon! This is possible only because of the <u>highly entangled nature of the vacuum</u>: were the field theory degrees of freedom (dofs) in the vicinity of the Moon in a product state w/ the field dofs here, then no operators we act w/ here could do anything there.

Note: it is crucial to allow for <u>non-unitary operators</u> $\rightarrow$ those are allowed to change the amount of entanglement (see w/ $\frac{dS_{vN}}{dt}$ and causal structure later).

Hilbert space factorization

One would think that we could factorize $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_{A^c}$, and for theories with no gauge symmetries, this is indeed sufficient. But if we have gauge fields, we need to tensor factor the $\mathcal{H}$-space in a gauge-invariant way. However, this is problematic: we now have "link" variables as basic operators instead of "site" operators, so when we cut a link, we have to decide where the broken link belongs $\rightarrow$ ambiguity! So one has to make a choice.

Resolution of both issues: algebraic QFT

Both of these issues are resolved if we don't talk about Hilbert spaces, but about algebras of operators in a given subregion.

Algebras of local observables do factorize, and they also provide an understanding of UV divergences at short distances. Namely, we now know that the reason why there are UV divergences is because we have a QFT on a fixed background, and such theories, when considering a subregion, imply an algebra of local operators of a particular type III_1, so the UV divergence is now understood to be a property of the algebra, and not of the state. (for instance, in QM, we have a type I algebra).

We will mostly ignore this, and deal w/ a physicist's way of using "split properties" and "particular choices" for gauge fields, but we should have this in mind.

Quantifying entanglement → von Neumann entropy

Or in other words, it quantifies the amount of ignorance.

DEF: $S_A = -\mathrm{tr}_A \rho_A \log \rho_A$, $\rho_A = \mathrm{tr}_B\, \rho_{TOTAL}$, $\rho_{TOTAL} = |\psi\rangle\langle\psi|$ some pure state

For a finite system, we can explicitly diagonalize ρ_A and obtain eigenvalues λ_i → this is known as obtaining the entanglement spectrum: $S_A = -\lambda_i \sum_i \ln \lambda_i$

iff statement!

For a pure state: $S_A = 0$ → 0 ignorance about our system!

For a mixed state, $S_A \neq 0$ → we have a list of probabilities! (mixed state)

Entanglement entropy (or von Neumann, or fine-grained entropy) measures how much a given state differs from a product (separable) state. It reaches the maximum value when a given state is a superposition of all possible quantum states w/ an equal weight.

Example: Bell pairs: 4 independent maximally entangled states in the 2-qubit system.

$|B_1\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ $\qquad \mathcal{H}_{TOT} = \{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$

$|B_2\rangle = \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle)$ $\qquad |ij\rangle = |i\rangle_A \otimes |j\rangle_B$

$|B_3\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$ $\qquad \langle i|j\rangle = \delta_{ij}\ , \quad i,j = 0,1.$

$|B_4\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$

Let's pick one, say $|B_4\rangle$. What is the entanglement entropy associated w/ a subsystem A? Recall, formula is $S_A = -\operatorname{tr} \rho_A \ln \rho_A$.

First, need to find a reduced density matrix $\rho_A = \operatorname{tr}_B \rho$.

$$\rho_A = \operatorname{tr}_B |\Psi\rangle\langle\Psi| = \sum_{i=0}^{1} {}_B\langle i| \tfrac{1}{2}(|01\rangle - |10\rangle)(\langle 01| - \langle 10|)|i\rangle_B =$$

$$|0\rangle_A \otimes |1\rangle_B$$

$$= \tfrac{1}{2}\, {}_B\langle 0|(|01\rangle - |10\rangle)(\langle 01| - \langle 10|)|0\rangle_B + \tfrac{1}{2}\, {}_B\langle 1|(|01\rangle - |10\rangle)(\langle 01| - \langle 10|)|1\rangle_B =$$

$$= \tfrac{1}{2}\, {}_B\langle 0|10\rangle\langle 10|0\rangle_B + \tfrac{1}{2}\, {}_B\langle 1|01\rangle\langle 01|1\rangle_B =$$

$$= \tfrac{1}{2}|1\rangle\langle 1| + \tfrac{1}{2}|0\rangle\langle 0| = \tfrac{1}{2}\begin{pmatrix}1 & 0\\ 0 & 1\end{pmatrix}.$$

The von Neumann entropy is then: $S_A = -\operatorname{tr}\rho_A \log \rho_A = -\sum_{k=0}^{1} \lambda_k \ln \lambda_k =$

$$= -\tfrac{1}{2}\ln\tfrac{1}{2} - \tfrac{1}{2}\ln\tfrac{1}{2} = \ln 2.$$

$\Rightarrow$ Maximally entangled state since $\ln 2$ saturates bound $d_A = d_B$, $d = \ln \dim \mathcal{H}$

Entanglement entropy inequalities

S_{VN} is most interesting for bipartitions (or multi-partitions),

$\mathcal{H} = \otimes_i \mathcal{H}_i$.

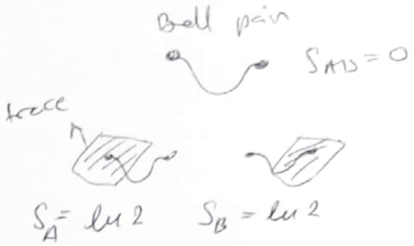

* Subadditivity *

$\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_{A^c} \Rightarrow \boxed{S_A + S_{A^c} \geq S_{AA^c}}$

Saturation occurs if we have unentangled (product) states: then $S_{AA^c} = S_A + S_{A^c}$. But if states are entangled, then RHS goes down. E.g. when we have $S_A = S_{A^c}$ (a pure total state), then $S_{AA^c} = 0$, since all modes purify each other.

In QFT, is S finite? No! Recall, the UV divergences emerge exactly around the entangling surface. So, how does the entropy scale? If whole state pure, then $S_A = S_B$, and so the only information the two regions A and B share is the area of the entangling surface. Therefore, $S \sim \frac{\text{Area}}{\epsilon^2}$ in 4d $\to$ diverges as $\epsilon \to 0$!

* Mutual information *

$$\boxed{I(A{:}B) = S_A + S_B - S_{AB}}$$

$\hookrightarrow$ measures correlations between two subsystems, classical and quantum.

it's also UV finite (divergences cancel out)

* Strong subadditivity *

Requires multi-partition. Say, for $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$. Then:

$$\boxed{S_{ABC} + S_B \leq S_{AB} + S_{BC}}$$

$\to$ Tells us about "monogamy of entanglement" and how strongly are things entangled w/ each other

$\to$ Important for bh info paradox.

Deep and hard to prove.

To see monogamy easier, say we have a system ABCD in a pure state. Then $S_{AB} = S_{CD}$, and $S_{ABC} = S_D$, so from above, we can rewrite

$$\left.\begin{array}{l} S_{AB} \to S_{CD} \\ S_{ABC} \to S_D \end{array}\right\} \Rightarrow S_D + S_B \leq S_{CD} + S_{BC}$$

$$\Rightarrow (S_{CD} - S_D) + (S_{BC} - S_B) \geq 0$$

$$\Rightarrow \boxed{S(C|D) + S(C|B) \geq 0}$$

: a given qubit in C can be maximally entangled w/ D or B, but not both!

* Relative entropy *

Given 2 density matrices ρ and σ, we can define $S(\rho\|\sigma)$ which provides a measure of distinguishability between them.

$$\boxed{S_{rel} = S(\rho\|\sigma) = \mathrm{tr}(\rho \log \rho) - \mathrm{tr}(\rho \log \sigma)}$$

Very useful concept, mainly because of 2 important properties:

a) Positivity: $S(\rho\|\sigma) \geq 0$, $S(\rho\|\sigma) = 0$ iff $\rho = \sigma$

Intuition: need additional information for them to be the same, and more info means non-zero entropy (we're ignorant!), so $S_{rel} > 0$.

b) Monotonicity: $S(\rho_{AB}\|\sigma_{AB}) \geq S(\rho_A\|\sigma_A)$, $\rho_A = \mathrm{tr}_B\, \rho_{AB}$, $\sigma_A = \mathrm{tr}_B\, \sigma_{AB}$

→ Taking a partial trace can only reduce relative entropy!

* Relative entropy useful: can prove the generalized 2nd law (monotonicity) and the Bekenstein bound (positivity – Casini '08)

Wall '11

Time evolution

Entanglement entropy, unlike thermodynamic entropy, is invariant under a unitary transformation that is (more generally) time evolution,

$$\rho_A \to U \rho_A U^\dagger, \quad U(t) = e^{+iHt}, \quad UU^\dagger = 1, \quad \frac{dU}{dt} = +iHU$$

$\boxed{\text{proof}}$ $S = -\mathrm{tr}\,\rho \ln \rho$

$$\frac{d\rho}{dt} = \frac{d}{dt}\left(U(t)\, \rho\, U^\dagger(t)\right) = +iH \overbrace{U\rho U^\dagger}^{\rho(t)} + i\, \overbrace{U\rho U^\dagger}^{\rho(t)} H = -i[\rho, H]$$

Ehrenfest theorem: $\frac{d\langle A\rangle}{dt} = -i\hbar[A, H] + \langle \frac{\partial A}{\partial t}\rangle$ (but ρ_0 doesn't depend on t)

$$\frac{dS}{dt} = -\mathrm{tr}\left(\frac{d\rho}{dt} \ln \rho + \rho \frac{d \ln \rho}{dt}\right) = -\mathrm{tr}\left(-i[\rho, H] \ln \rho + \rho \rho^{-1}(-i[\rho, H])\right) =$$

$$= i\,\mathrm{tr}\left([\rho, H] \ln \rho\right) \underset{\uparrow}{=} i\,\mathrm{tr}(\rho H \ln \rho - \rho \ln \rho\, H) = i\,\mathrm{tr}(\rho \underbrace{[H, \ln \rho]}_{= -\frac{1}{i}\frac{d\ln\rho}{dt}}) =$$

use cyclicity of the trace

Ehrenfest

$$= -\mathrm{tr}\left(\rho \rho^{-1} \frac{d\rho}{dt}\right) = -\frac{d}{dt}\underbrace{\mathrm{tr}\rho}_{1} = 0$$

$\boxed{9}$

So we see that entanglement entropy is preserved under time evolution. This is not in contradiction with the 2nd law: the entropy didn't increase, but it didn't decrease either: this is because we haven't done any coarse-graining to our system, so knowledge is preserved in the density matrix.

Another way of seeing this: data on initial slice is enough to determine (through equations of motion) the solution in the whole causal diamond → so with unitary operators, we get all the information in the causal diamond.

$S(A) = S(A')$

$D[A]$ - domain of dependence

More precisely, if we unitarily evolve ρ_A by transformations that are supported only on $\mathcal{H}_A$, the eigenvalues of ρ_A remain unaffected.

⇒ The entanglement spectrum of ρ_A depends only on $D[A]$ and not on the particular choice of Σ (ρ_A is a "wedge observable").

Fine-grained vs. coarse-grained entropy

S_{VN} is a fine-grained entropy (doesn't change in time), but how do we obtain a coarse-grained one?

→ We fix some macroscopic data, like the energy or temperature (something we can measure) and we stay maximally ignorant about underlying details.

Then, we only measure a set of simple, coarse-grained observables A_i.

In summary, we obtain a coarse-grained entropy in 2 steps:

1) Consider all density matrices $\bar{\rho}$ such that $tr(\rho A_i) = tr(\bar{\rho} A_i)$;

2) Compute $S(\bar{\rho})$ and maximize over all $\bar{\rho}$ (A_i = E or similar).

Renyi entropy (Nishioka, 1801.10352)

Entanglement entropy is sometimes hard to calculate, so we can use some other (easier) quantities instead. One such quantity is known as the Renyi entropy, defined as

$$S^{(n)} = \frac{1}{1-n} \log \operatorname{tr}\rho^n \rightarrow \text{ nth Renyi entropy}$$

In fact, one can prove that $S_{VN} = \lim\limits_{n \to 1} S^{(n)}$.

$\boxed{\text{proof.}}$

$$S = \lim_{n\to 1} S^{(n)} = \lim_{n \to 1} \frac{1}{1-n} \log \operatorname{tr}\rho^n \overset{\text{L'Hopital since both} \to 0}{=} -\lim_{n\to 1} \partial_n (\log \operatorname{tr}\rho^n) =$$

$$= -\lim_{n\to 1} \partial_n \operatorname{tr}\rho^n = -\lim_{n \to 1} \operatorname{tr} \partial_n \rho^n \underset{\partial_a x^a = x^a \ln x}{=} -\lim_{n\to 1} \operatorname{tr}\rho^n \ln \rho = -\operatorname{tr}\rho \ln\rho = S \quad \blacksquare$$

Example : Bell pair $|B_4\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$, recall $\rho_A = \frac{1}{2}\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \rightarrow \lambda_k = 1/2$.

$$\operatorname{tr}\rho^n = \sum_{k=0}^{1} \lambda_k^n = \left(\tfrac{1}{2}\right)^n + \left(\tfrac{1}{2}\right)^n = 2^{1-n}$$

$$\Rightarrow S_A = -\lim_{n\to 1} \partial_n \log 2^{1-n} = -\lim_{n \to 1} \partial_n \log 2^{-n} = \log 2 \quad \blacksquare \checkmark$$

So this was easy (and somewhat unnecessary) in QM, but what about in QFT? Here we will use the Euclidean path integral to calculate.

Recall, to prepare a state: we care about the vacuum since everything else $\forall |\Psi\rangle$ can be obtained from it:

$\langle \phi_0(\vec{x}) | \Psi \rangle =$

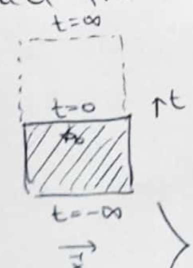

"ket"

$\Big\}$ the path integral performed over the shaded region

$\langle \Psi | \phi_0'(\vec{x}) \rangle =$

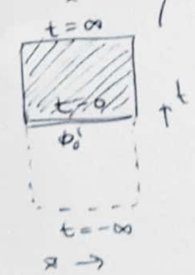

"bra"

In formulas:

$$\Psi[\phi_0(\vec{x})] = \langle \phi_0(\vec{x}) | \Psi \rangle = \int_{t=-\infty,\, \phi_0(\vec{x})}^{t=0,\ \phi(t=0,\vec{x}) = \phi_0(\vec{x})} [D\phi(t,\vec{x})]\, e^{-I_E[\phi]}$$

$$\Psi^*[\phi_0'(\vec{x})] = \langle \Psi | \phi_0'(\vec{x}) \rangle = \int_{t=0,\ \phi_0}^{t=\infty} [D\phi(t,\vec{x})]\, e^{-I_E[\phi]}$$

The wave function of the ground state $|\Psi\rangle$ is given by the wave functional $\Psi[\phi_0(\vec{x})] = \langle \phi_0(\vec{x}) | \Psi \rangle$ whose EPI representation is given above.

To obtain the partition function, we integrate over the whole Euclidean space

$$Z = \int [D\phi_0(t=0, \vec{x})] \underbrace{\langle \psi | \phi_0 \rangle \langle \phi_0 | \psi \rangle}_{\text{just combine previous 2.}}$$

sum over all possible boundary conditions at point $\phi_0(t=0, \vec{x})$

$= \langle \psi | \psi \rangle$

tr_B $\overline{B \,|\, A \,|\, B}$

* How to obtain partial trace: we glue the edges of the 2 sheets $\langle \psi | \phi_0 \rangle$ and $\langle \phi_0 | \psi \rangle$ along B: that is, integrate over every state $\phi^B(t=0, \vec{x})$ w/ support only on $\vec{x} \in B$:

$$\rho_A = \frac{1}{Z} \int [D\phi^B(t=0, \vec{x} \in B) \langle \phi^B | \psi \rangle \langle \psi | \phi^B \rangle$$

$$= \frac{1}{Z} \mathrm{tr}_B | \psi \rangle \langle \psi |$$

This reduced density matrix has two indices: $[\rho_A]_{ab} = \langle \phi_a^A | \rho_A | \phi_b^A \rangle$, where $\phi_a^A(\vec{x} \in A)$ and $\phi_b^A(\vec{x} \in A)$ specify b.c. on A at $t = 0^+$ and $t = 0^-$, respectively.

In other words, $\underset{\langle \phi_a | \rho_A | \phi_b \rangle}{[\rho_A]_{ab}} = \frac{1}{Z} \int [D\phi^B(t=0, \vec{x} \in B)] \{ \langle \phi_a^A | \langle \phi^B | \psi \rangle \langle \psi | \phi_b^A \rangle | \phi^B \rangle \}$

$$= \frac{1}{Z} \int [D\phi^B(t=0, \vec{x} \in B)]$$

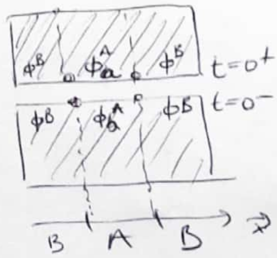

$=$

after integrating ϕ^B

$$= \frac{1}{Z}$$

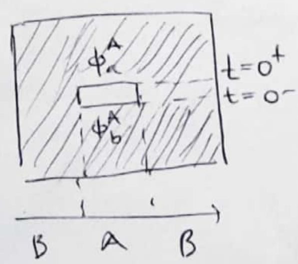

$=$

$$= \frac{1}{Z} \int [D\phi(t, \vec{x})] e^{-I_E[\phi]} \prod_{\vec{x} \in A} \delta(\phi(0^+, \vec{x}) - \phi_a^A(\vec{x})) \delta(\phi(0^-, \vec{x}) - \phi_b^A(\vec{x}))$$

So $\rho \sim$  $=$  → sometimes we use this notation

Can define this manifold as $\mathcal{M}$.

We derived the pictorial representation for ρ. Recall, we want to use this to calculate the Renyi entropy, which involves $\mathrm{tr}\rho^n$.

Pictorially, this means:

$$\mathrm{tr}\rho_A^n = \frac{1}{(Z)^n}$$

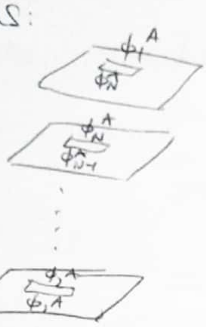

→ we're gluing along the cuts since that's what the trace instructs us to do

→ call ~~view~~ define this weird manifold as $\mathcal{M}_n$,

and write $\mathrm{tr}\,\rho_A^n = \frac{1}{(Z)^n} \cdot Z_n$,

where Z_n is the partition function on this weird manifold $\mathcal{M}_n$, which is an n-fold cover of the original spacetime $\mathcal{M}$

n-fold cover = manifold constructed by gluing n-copies of the single $\mathcal{M}$ sheet along region A.

We can then express the formula for our Renyi entropy as

$$S^{(n)} = \frac{1}{1-n}\log \mathrm{tr}\rho^n = \frac{1}{1-n}\log\left\{\frac{Z_n}{Z^n}\right\}$$

$$\Rightarrow \boxed{S_{vN} = \lim_{n\to 1} S^{(n)} = -\lim_{n\to 1}\partial_n\left(\log Z_n - n\log Z\right)}$$

→ See the benefit of this definition in gravity!

* Note that this n-fold cover has a conical singularity along the codim-2 entangling surface $\Sigma = \partial A$ w/ a deficit angle $2\pi(1-n)$

* Note also that the cyclicity of the trace translates into a cyclic permutation symmetry amongst all of the copies of the functional integral → EPI has a cyclic $\mathbb{Z}_n$ symmetry acting on its components ⇒ this is known as the replica symmetry

* Caveat: ∃ n copies of a manifold, clearly a discrete number. How can we then send $n\to 1$?

→ We need to ensure that ∃ an analytic continuation away from integers (discrete copies).

Not every function allows for an analytic continuation away from integers. A simple example is $\sin(\pi z)$ which is 0 for z integer, but not zero otherwise. There is a theorem, known as the Carlson theorem, which tells us under which conditions we're allowed to do this: an important condition is that the function cannot grow rapidly (a.k.a. exponentially) at $\pm i\infty$ (that's why $\sin(\pi z)$ doesn't have a good anal. continuation). If Carlson's conditions are satisfied, we get a unique analytic continuation. We will assume this.

All that's left now is to calculate the fancy partition function Z_n. One can do this via the heat kernel method (see Nishioka '18), or if we're in a 2D CFT, one can relate the partition function to the correlation f-on of certain "twist operators" (see Adi's lecture).

Or, we can use AdS/CFT and find the gravitational dual to the weird $\mathcal{M}_n$ manifold theory (see Lewkowycz - Maldacena '13).