Modeling complex systems as interacting agents

Chris Fields

Machine Learning Approaches for Complexity

Erice, 27 April 2024

Why “agents”?

- Agent language is convenient, e.g. in agent-based programming, particles “feeling” forces, etc.
- Systems (particularly complex systems) behave in ways that we don’t expect. *Our* knowledge is limited.
- Conway-Kochen theorem: if any system is free from local determinism, all systems are.
- Free Energy Principle, Bayesian mechanics: least-action principles $\longleftrightarrow$ approximate Bayesian inference.

Session 1: Basic ideas

- Desiderata
- Unitary evolution of isolated systems
- Decompositions and separable systems
- Holographic principle, basis choice, local QRFs
- Computation with QRFs, CCCDs
- CCCDs to TQFTs
- Boundary evolution dual to bulk evolution

Desiderata:

- Minimal, simple, deep assumptions
- Universal applicability
- Scale free representation
- Straightforward translation to multiple domain-specific languages

Consider an isolated finite system U:

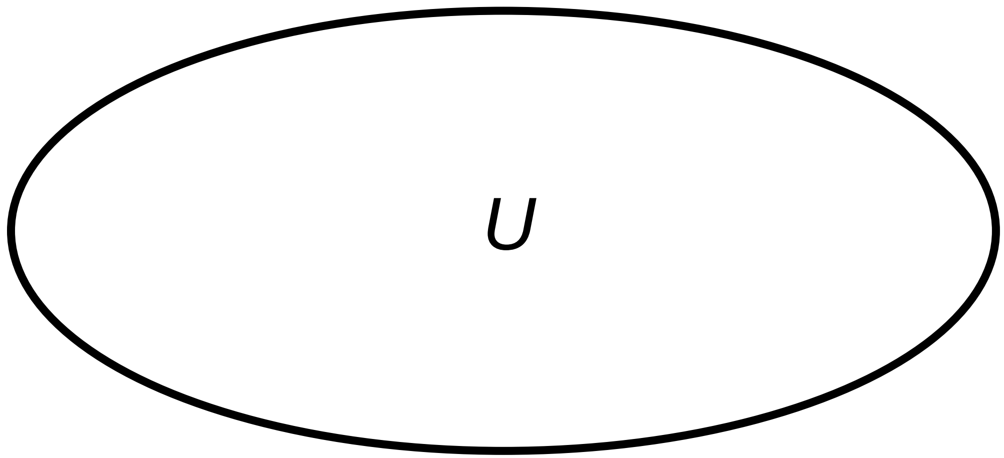

We can write $\mathcal{P}_U(t) = \exp[-(i/\hbar)H_U(t)]$.

We can only stipulate H_U.

t is some external, “objective” time parameter.

We are interested in finite decomposable systems, i.e.

$U = SE$ such that $|U\rangle = |S\rangle|E\rangle$.

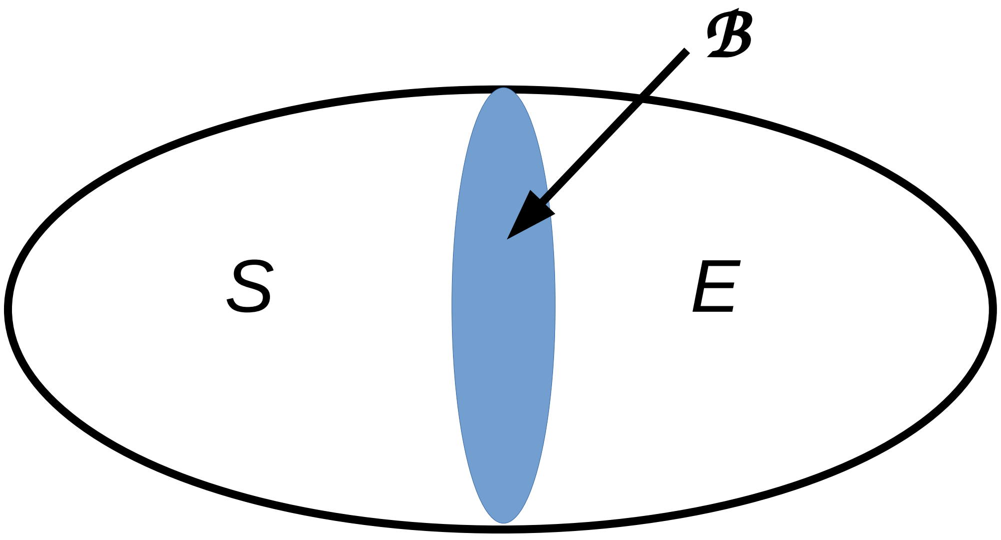

The interaction $H_{SE} = H_U - (H_S + H_E)$ is defined at the boundary $\mathcal{B}$. It has finite dimension and finite eigenvalues.

The Holographic Principle tells us that the flux of information through $\mathcal{B}$ is finite, $S(\mathcal{B}) \geq A_{\mathcal{B}}/4\ l_{\mathrm{P}}^2$

Let $S(\mathcal{B}) = N$. Then we can define a Hilbert space $\mathcal{H}_{\mathcal{B}}$ with $\dim(\mathcal{H}_{\mathcal{B}}) = 2^N$. This is ancillary, $\mathcal{H}_U \cap \mathcal{H}_{\mathcal{B}} = \emptyset$.

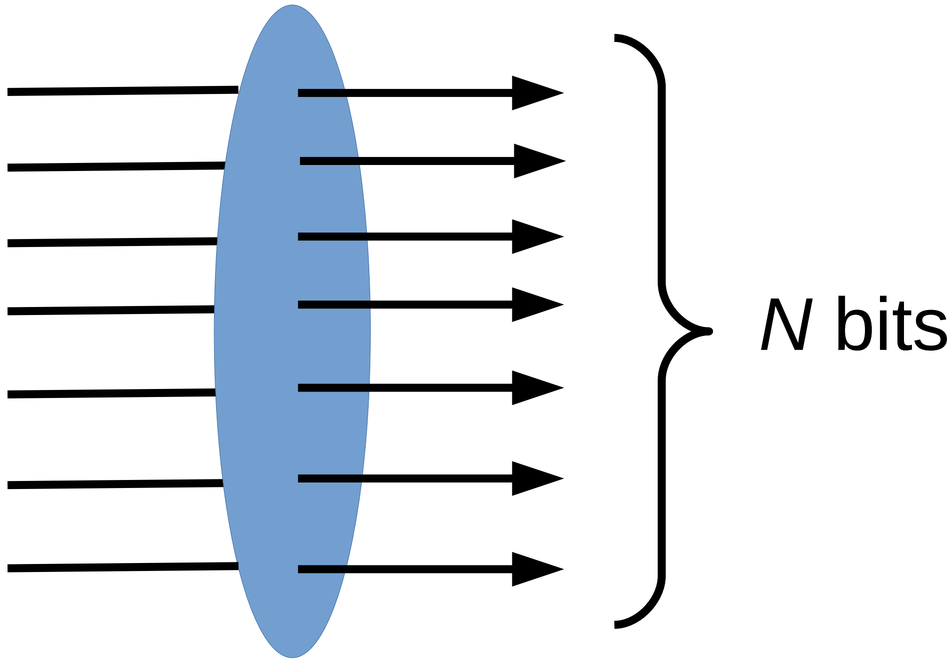

These N bits obviously encode the effect of S on E and vice-versa. They define an information channel.

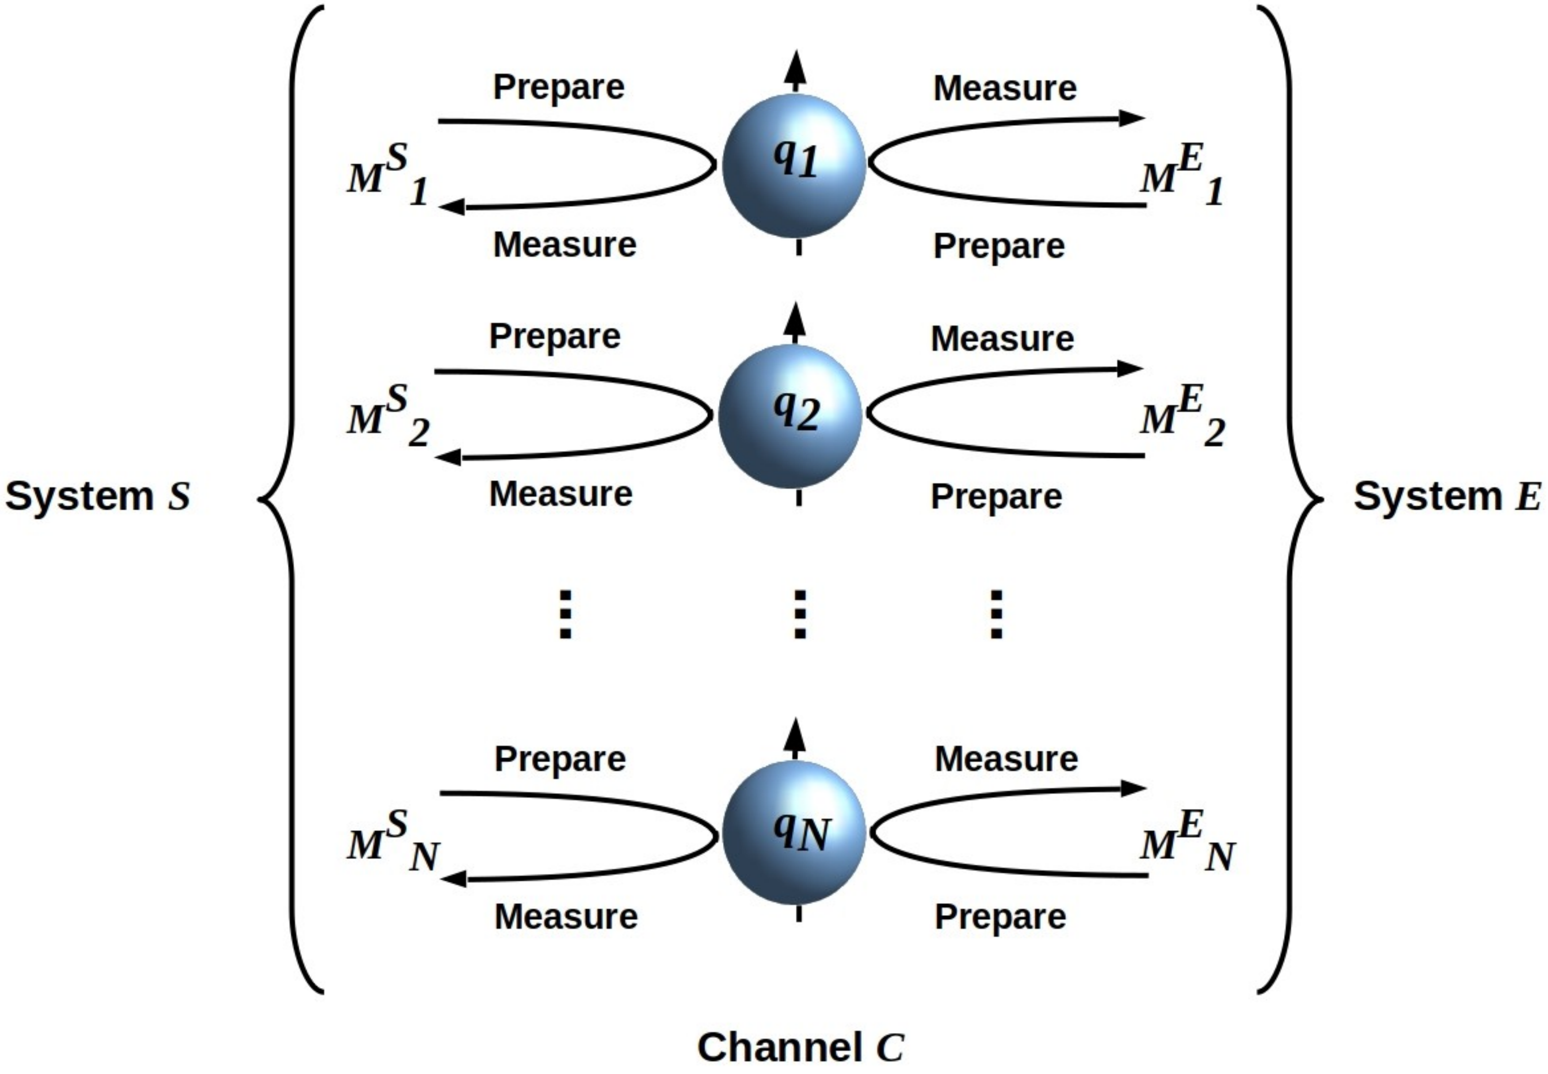

We can now write the Hamiltonian, for $k = S$ or E:

$$H_{SE} = \beta_k \, \mathrm{k_B} \, \mathrm{T}_k \, \Sigma_i \, \alpha_i \, M^k{}_i$$

The $M^k{}_i$ are single-qubit operators with eigenvalues ± 1 (i.e. copies of σ_z).

$\Sigma_i \, \alpha_i = 1$.

$\beta_k \geq \ln 2$ (Landauer's Principle) for irreversible "reads" and "writes" on $\mathcal{B}$.

Now we have:

- $\mathcal{B}$ is an ancillary *N*-qubit array.
- “Measuring” $\mathcal{B}$ yields an *N*-bit string.
- “Preparing” $\mathcal{B}$ encodes a *N*-bit string.
- These *N*-bit strings encode eigenvalues of H_{SE}.
- Interaction is symmetric information exchange.
- Informational equilibrium is not thermal equilibrium.

These are *universal* for decomposable (separable) isolated, finite systems *U*.

We’re interested in systems with lots of “internal” states and large enough boundaries for significant information flow. More like organisms than rocks.

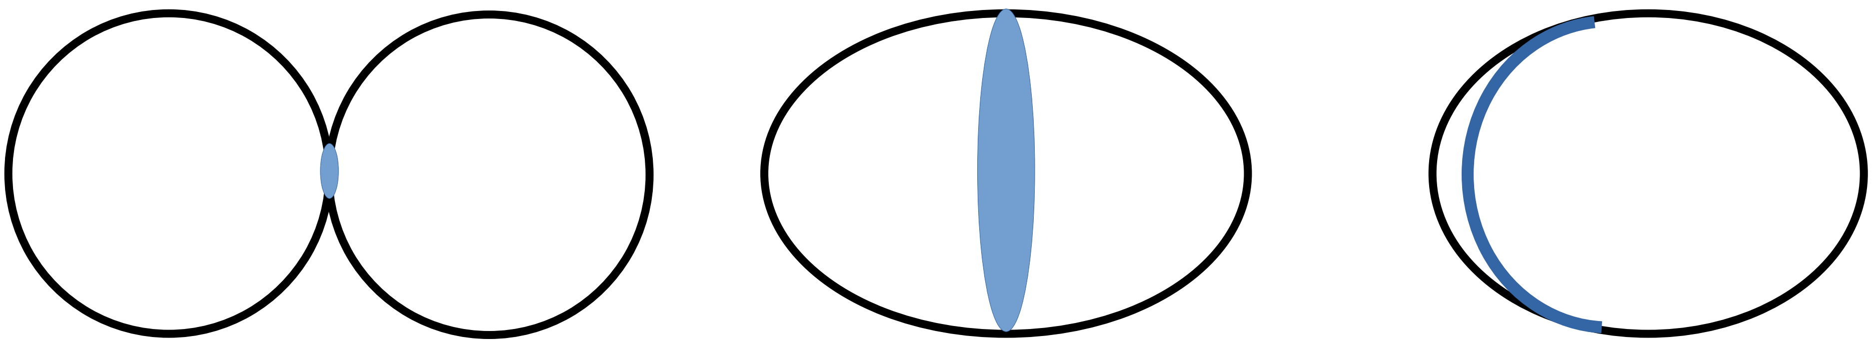

$\mathcal{B}$ is too small
(too few bits
transferred)

“Bag of Gold”
Black Hole

$\mathcal{B}$ is just right

Goldilocks!

Agents

$\mathcal{B}$ is too big
(separability
breaks down)

Boltzmann
Brain

We have H_{SE}. What about H_S (H_E)?

Interesting systems are doing something with the data they get from $\mathcal{B}$. They're *computing*.

$$
\begin{tikzcd}
\psi_i(|S(t_i)\rangle) \arrow[r, "f_{i \to j}"] & \psi_j(|S(t_j)\rangle) \\
|S(t_i)\rangle \arrow[u, "\psi_i"] \arrow[r, "\mathcal{P}_S(t)"'] & |S(t_j)\rangle \arrow[u, "\psi_j"']
\end{tikzcd}
$$

ψ is a functional interpretation (semantics) of *S*.

Measurements get their semantics from reference frames (RFs).

For σ_z, we have to say what “up” or “↑” means.

A *physically implemented* RF is a quantum RF (QRF).

$M^S{}_i$ could use

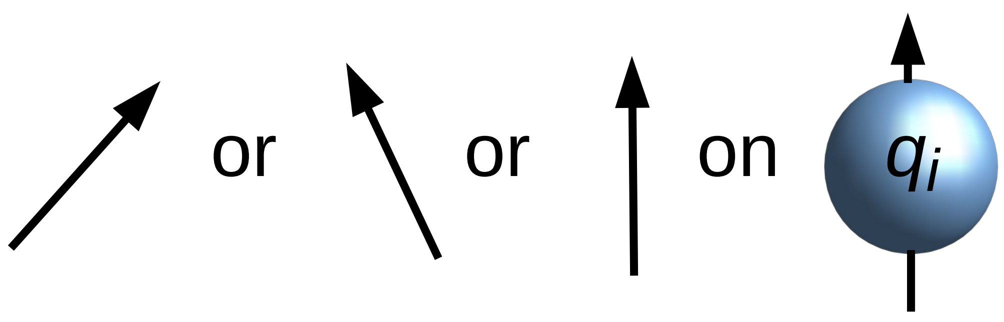

to get an outcome value.

This QRF is implemented by H_S.

S must choose (H_S must implement) an "↑" QRF for each of the N q_i on $\mathcal{B}$.

Choice of the N "↑" QRFs is:

• choice of a basis for H_{SE}.

• choice of a zero point for energy.

• choice of efficiency β and/or temperature T.

This choice is "free" provided $|SE>$ is separable.

Separability ⟷ free choice of QRFs.

Questions?

Consider a set X of perfectly correlated qubits on $\mathcal{B}$.
The correlation may be imposed by either S or E.

Reading X yields, and writing X encodes, a bit string.

We can think of this bit string as a correlated set of single-qubit QRFs – this becomes a QRF for X.

The correlation becomes a set of functions:

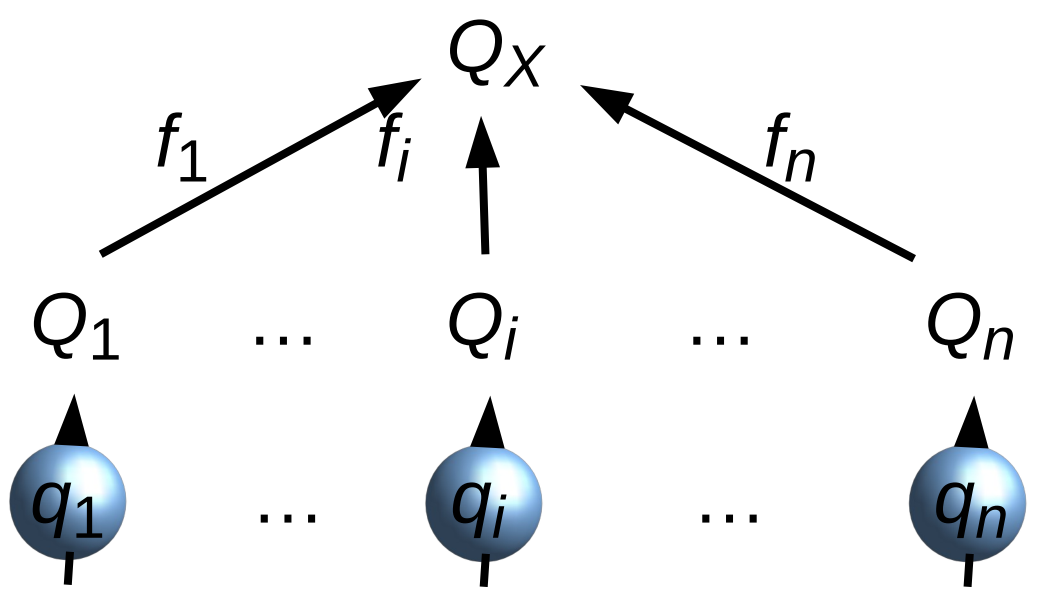

Category theory lets us define an optimal set of maps via universality.

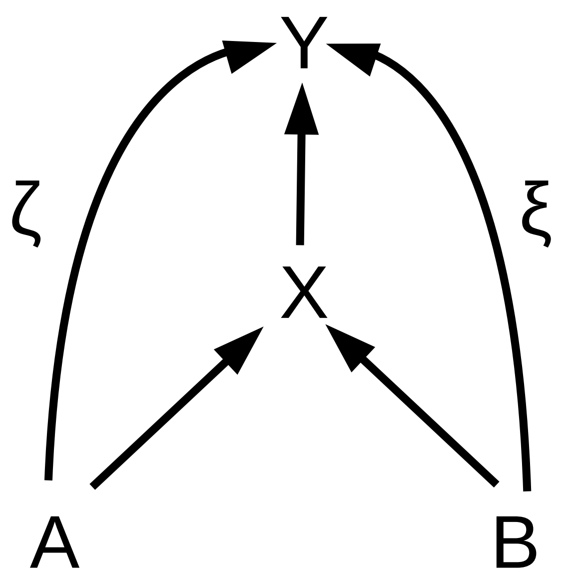

Diagram commutes for all ζ, ξ, X is the *colimit* for such maps.

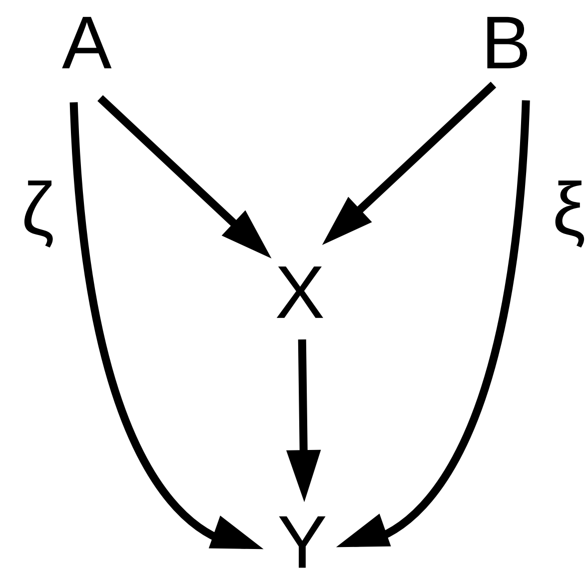

Diagram commutes for all ζ, ξ, X is the *limit* for such maps.

A *cocone diagram* (CCD) depicts a hierarchy of maps from a “base” of 1-bit classifiers to the colimit **C** of all maps from those classifiers.

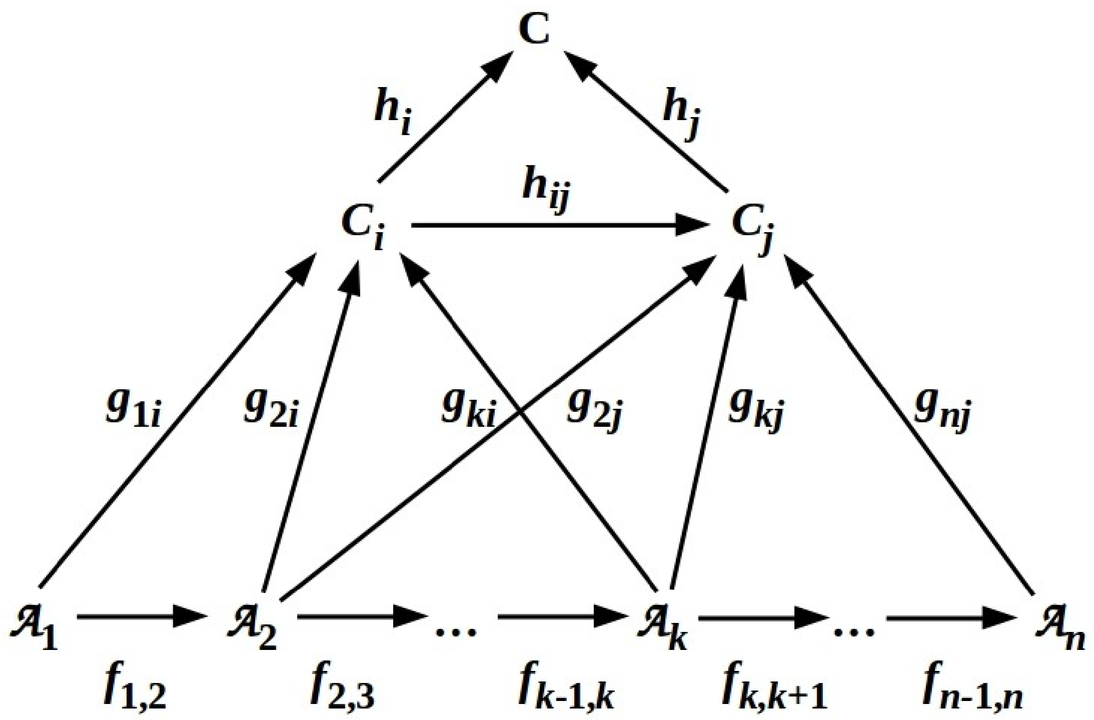

Each classifier $\mathcal{A}_i$ represents a single-qubit QRF.

Combining a cocone diagram with its dual cone diagram, mapping to the $\mathcal{A}_i$ from their limit, gives a CCCD:

$$
\begin{tikzcd}
 & \mathbf{C}' & \\
\mathcal{A}_1 \arrow[ur, "f_1"] \arrow[r, leftrightarrow, "g_{21}", "g_{12}"'] & \mathcal{A}_2 \arrow[u, "f_2"'] \arrow[r, leftrightarrow, "g_{32}", "g_{23}"'] & \ldots\ \mathcal{A}_k \arrow[ul, "f_k"'] \\
 & \mathbf{D}' \arrow[ul, "h_1"] \arrow[u, "h_2"] \arrow[ur, "h_k"'] &
\end{tikzcd}
$$

If the cores $\mathbf{C}' = \mathbf{D}'$, we can also represent the CCCD as:

$$
\begin{tikzcd}
\mathcal{A}_1 \arrow[r, leftrightarrow, "g_{21}", "g_{12}"'] & \mathcal{A}_2 \arrow[r, leftrightarrow, "g_{32}", "g_{23}"'] & \ldots\ \mathcal{A}_k \\
 & \mathbf{C}' \arrow[ul, "h_1"] \arrow[u, "h_2"] \arrow[ur, "h_k"'] & \\
\mathcal{A}_1 \arrow[ur, "f_1"] \arrow[r, leftrightarrow, "g_{21}", "g_{12}"'] & \mathcal{A}_2 \arrow[u, "f_2"'] \arrow[r, leftrightarrow, "g_{32}", "g_{23}"'] & \ldots\ \mathcal{A}_k \arrow[ul, "f_k"']
\end{tikzcd}
$$

CCCDs over single-qubit QRFs represent multi-qubit QRFs.

- This representation is optimal via the universality of limits/colimits.
- It represents both measurement (upward CCD maps) and preparation (downward CD maps).
- CCCDs can map any set of bit strings to a probability distribution.
- Commutativity enforces Kolmogorov probability.

We can “attach” a CCCD to B with maps.

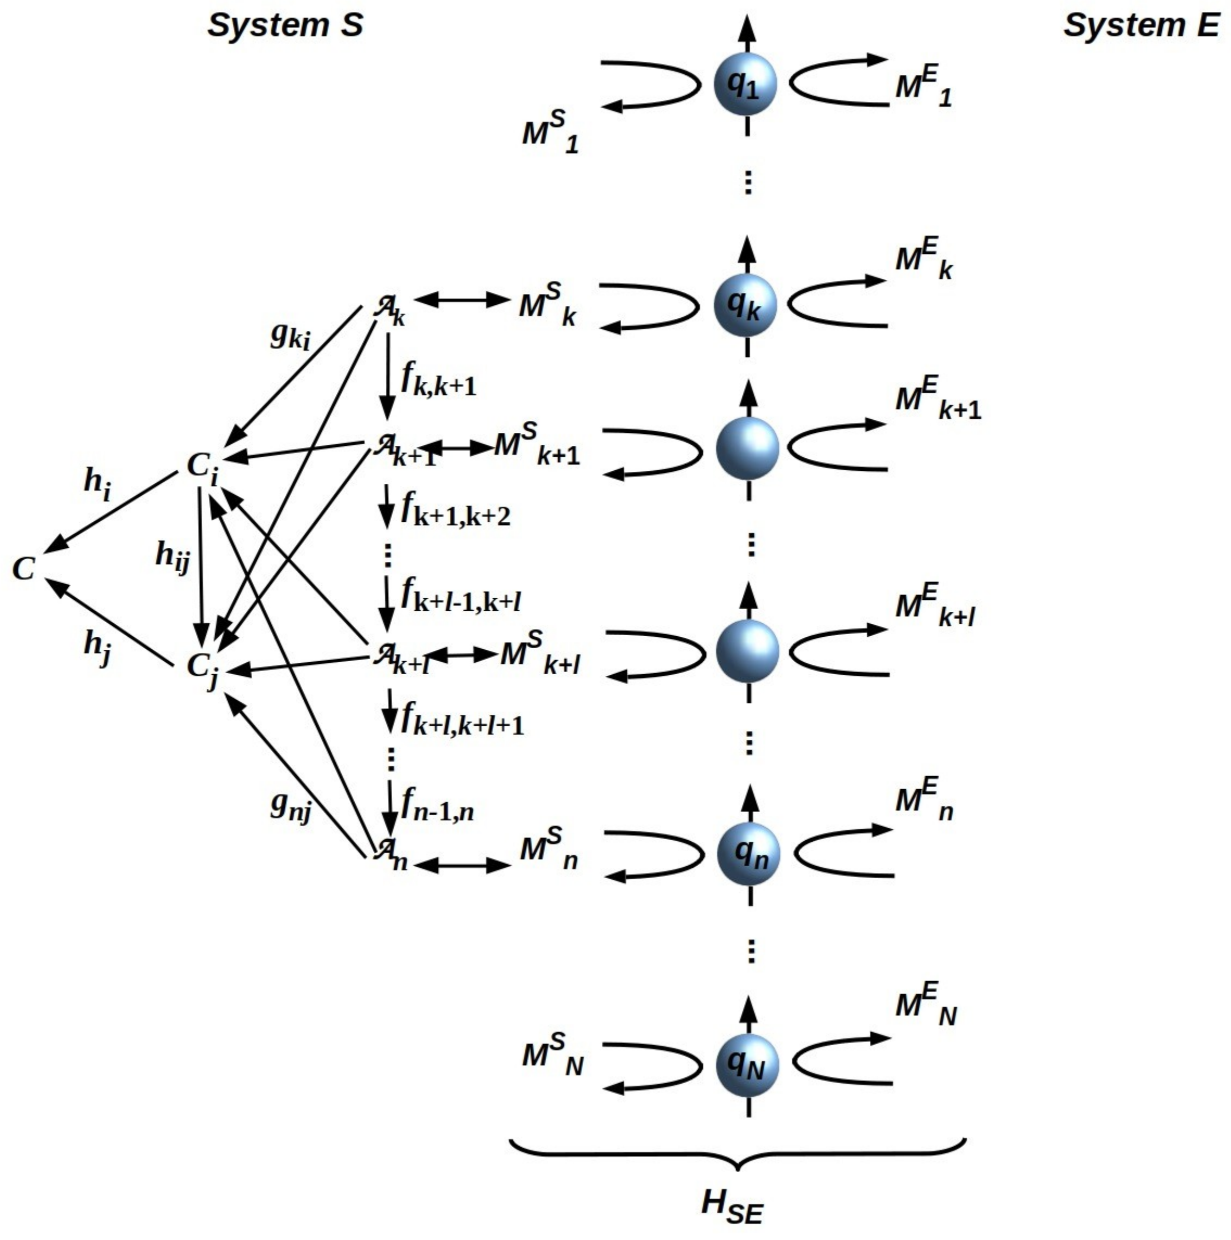

We can now say what's needed to observe a "system" X in a "pointer state" $|p\rangle$.

- We need a QRF R that *identifies* X by detecting an invariant "reference state" $|r\rangle$.

- We need a QRF P that detects $|p\rangle$.

- P and R must commute.

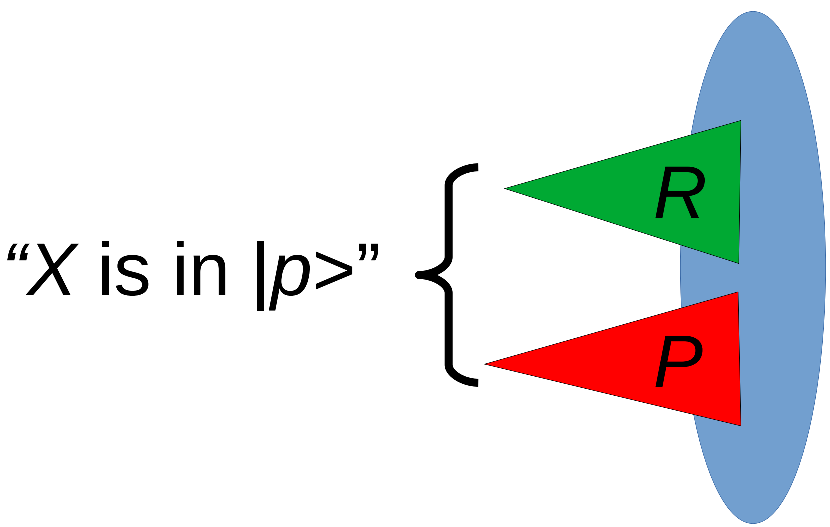

Observer-relative states $\longrightarrow$ observer-relative systems.

Note that this breaks the qubit-exchange symmetry on $\mathcal{B}$.

- Particular qubits encode information processed by particular QRFs.
- Some qubits must be devoted solely to thermodynamic exchange – different parts of $\mathcal{B}$ have different β and/or T.

Symmetry breaking $\longrightarrow$ semantics

Now let's add time to this picture.

Sequential operations with a QRF Q write a sequence of states $|q>$ of dom(Q) on $\mathcal{B}$.

So we can write Q: $|q(t_i)> \longrightarrow |q(t_f)>$.

This lets us identify Q with a TQFT $\mathcal{T}$: dom(Q) $\longrightarrow$ dom(Q).
(Technically, we have a functor **CCCD** $\longrightarrow$ **TQFT**).

Transporting the data on $\mathcal{B}$ forward in time is equivalent to transporting the computation Q forward in time.

QRFs to TQFTs

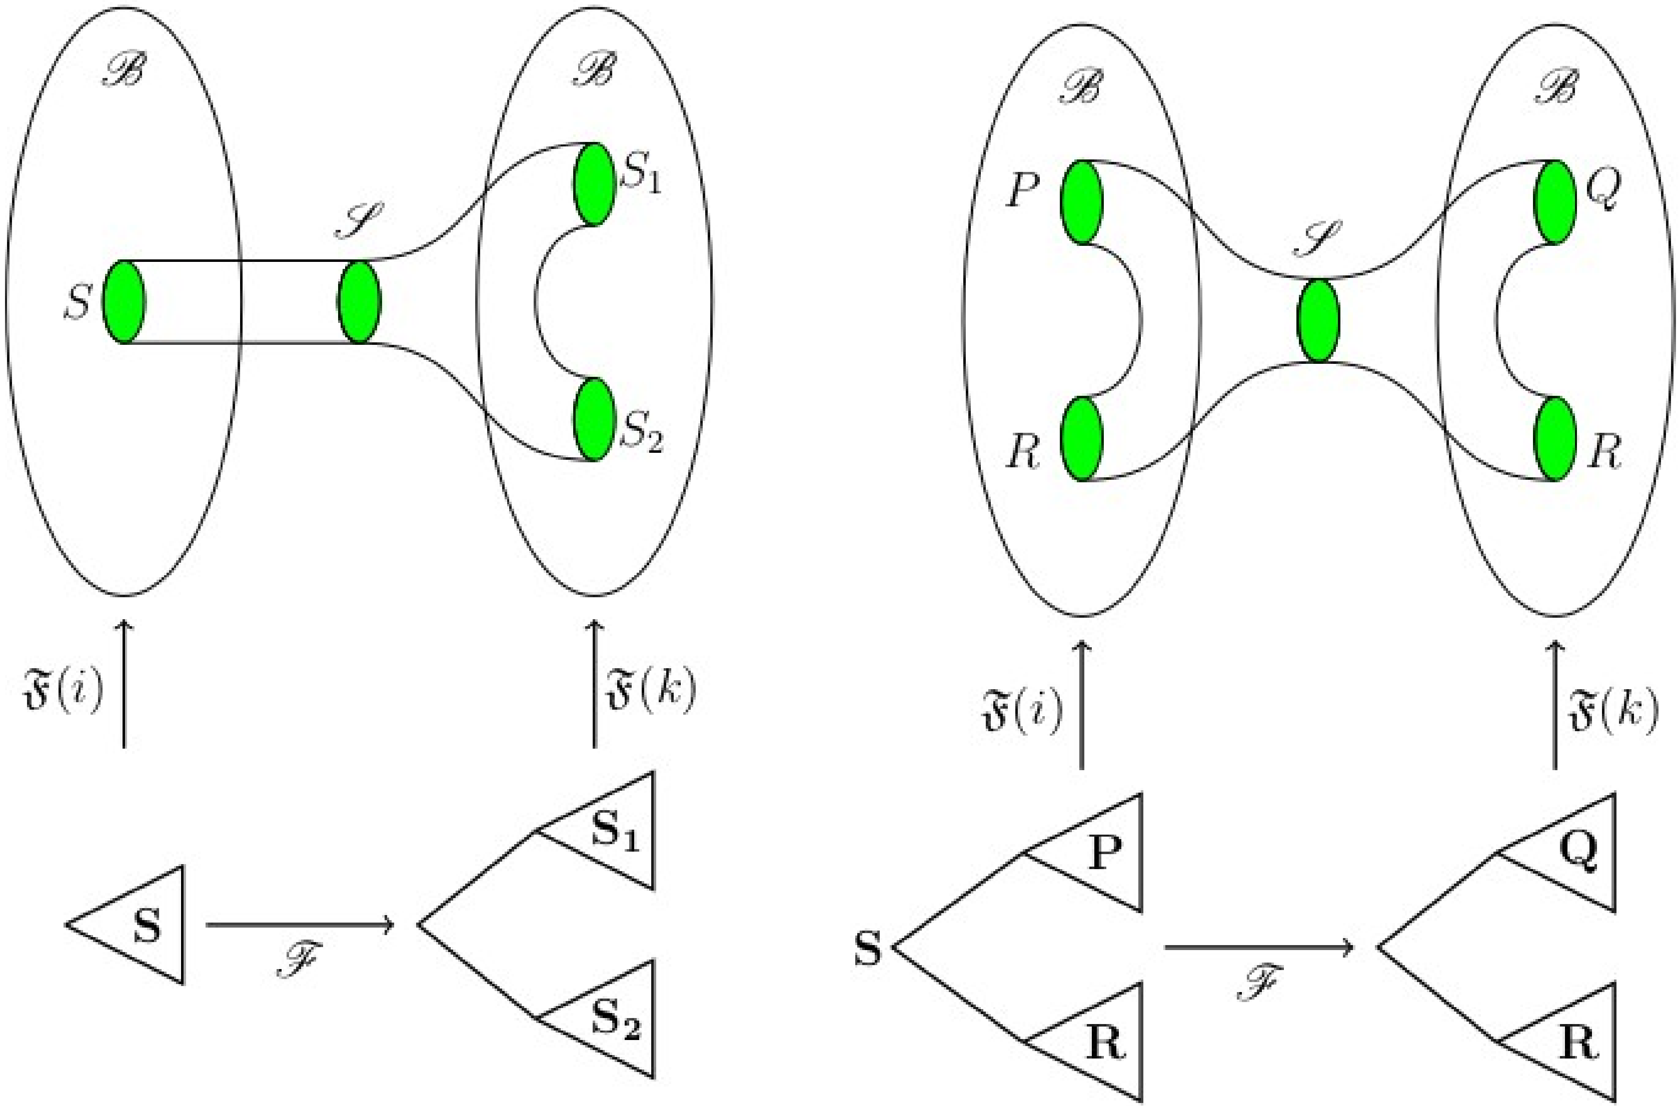

Summary:

- Physical interaction $\longleftrightarrow$ information exchange.
- Separability $\longleftrightarrow$ free choice of QRFs.
- Internal propagator $\exp[-(i/\hbar)H_S(t)] \longleftrightarrow$ computation.
- Computation $\longleftrightarrow$ data is a bulk-boundary duality.
- Agents are bundles of QRFs.
- Classical information is encoded on boundaries.

Questions?

Session 2: LOCC, QECCs, Spacetimes

- Time and memory
- LOCC protocols – meaning of “classical” communication
- Commutativity and compartmentalization
- Examples: Quantum Darwinism, Bell/EPR
- QECCs
- Redundancy, spacetime

Two usually-implicit assumptions:

- Observers have access to arbitrary read/write classical memory.
- Observers can exchange classical information.

We need these assumptions to do science!

How do we model these assumed capabilities explicitly, and what happens when we do?

Classical memory and time are coupled.

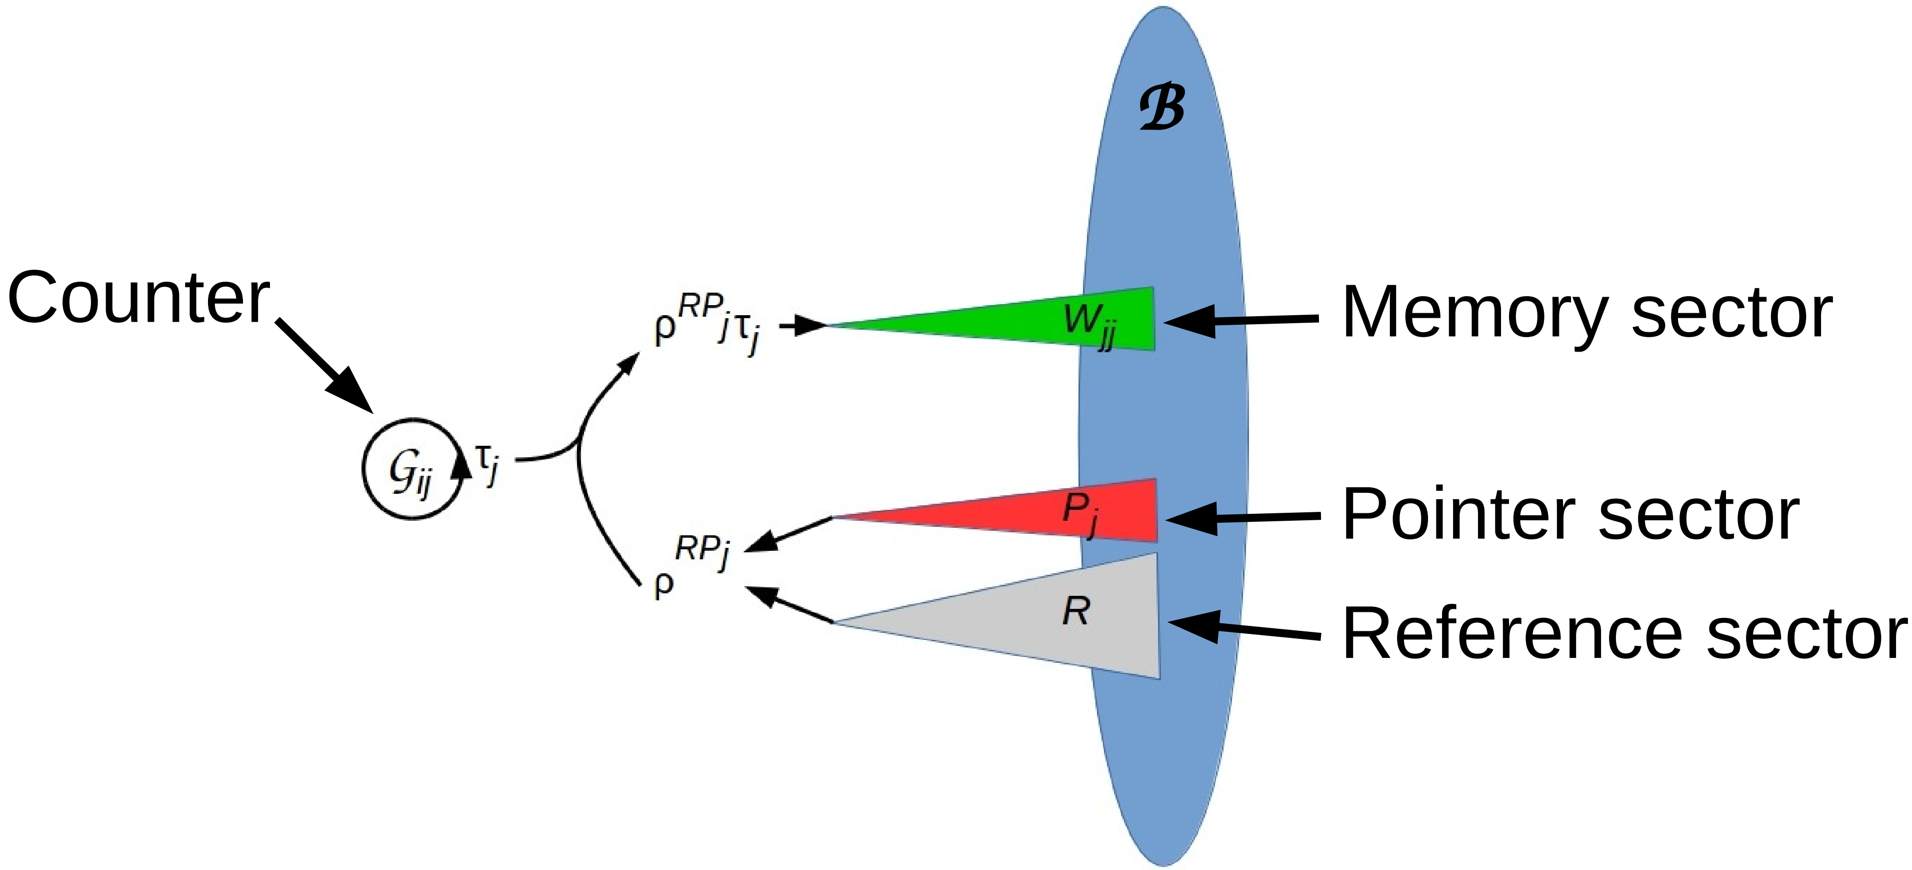

The memory sector must be slowly-varying.

Write this as a TQFT Q:

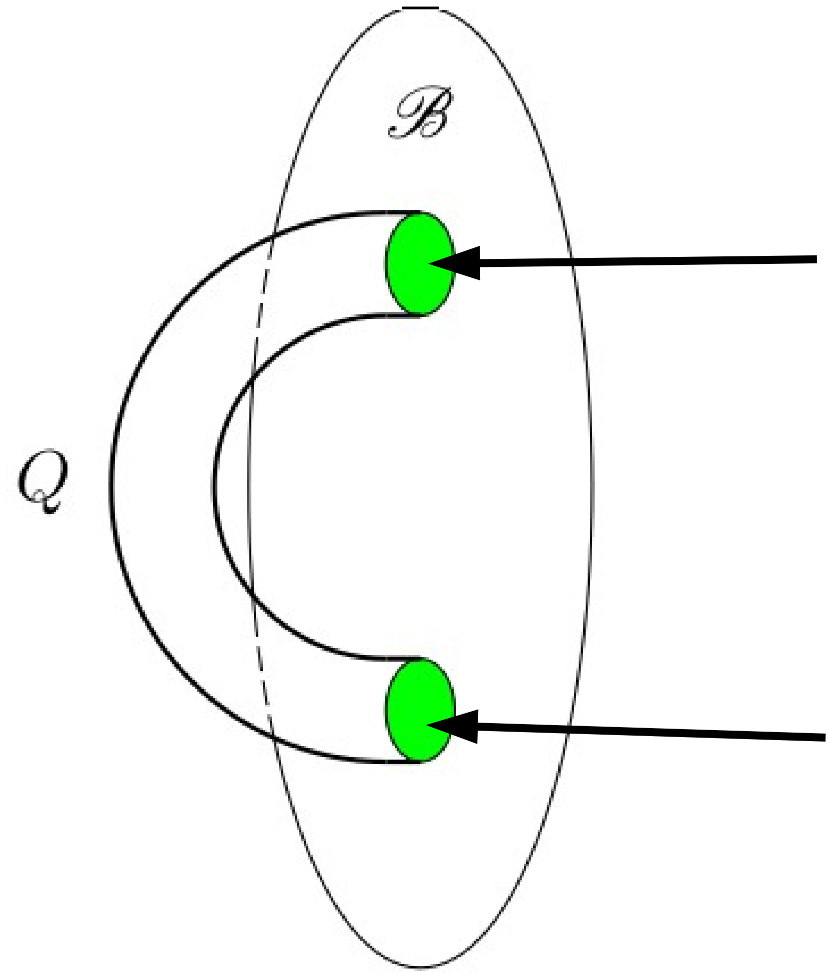

Data read here

is written here

and vice-versa.

Q is a quantum channel between the green I/O sectors.

Let the “system” (Alice) and its “environment” (Bob) both do this kind of computation.

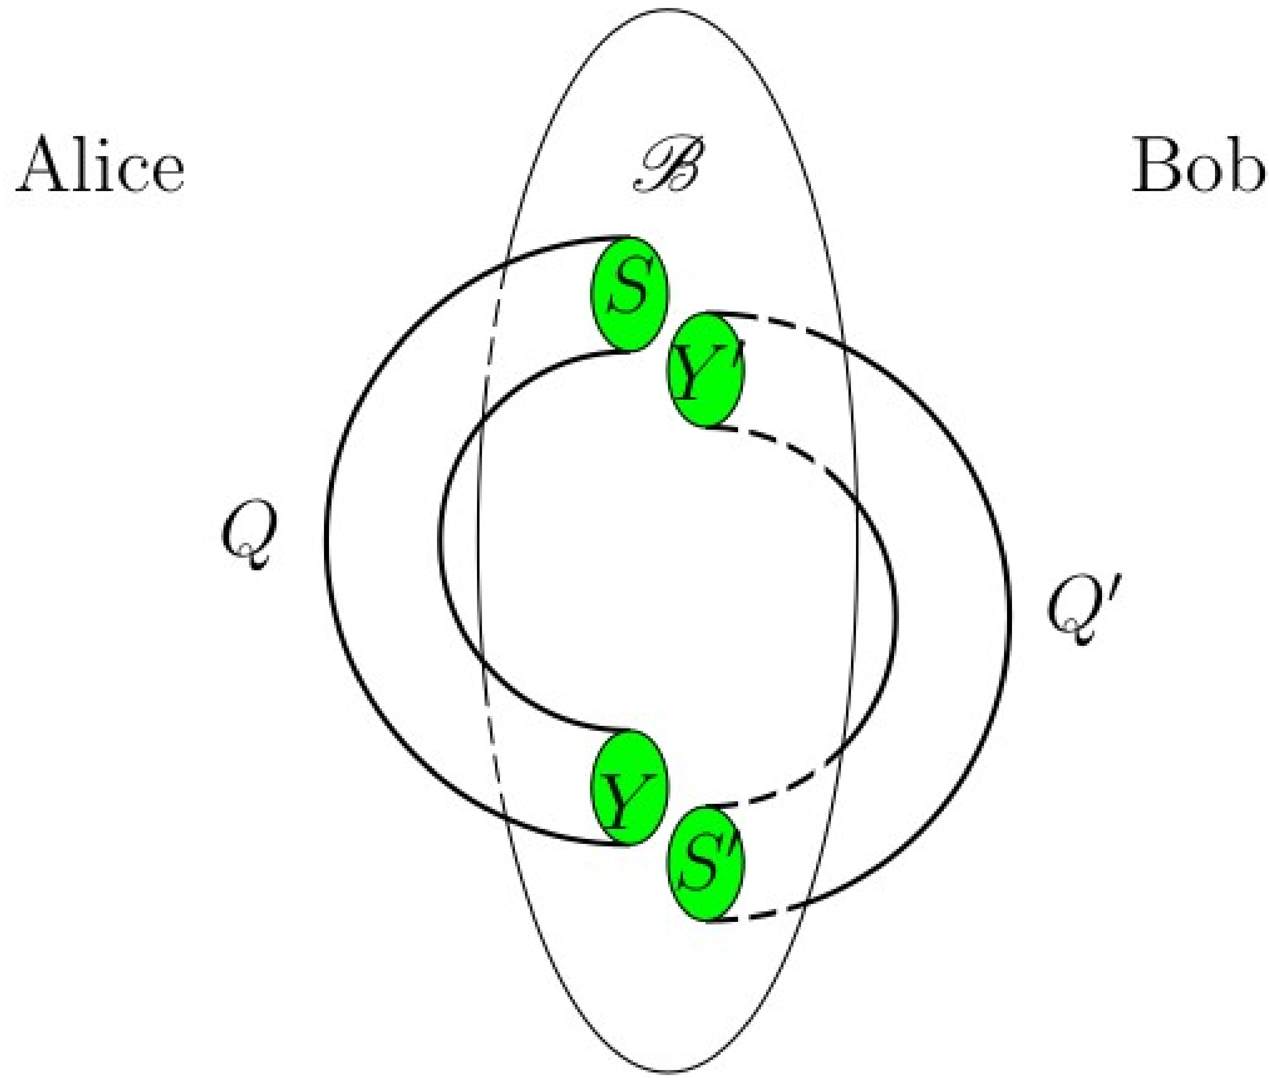

If the sectors align, each reads from and writes to the other's memory.

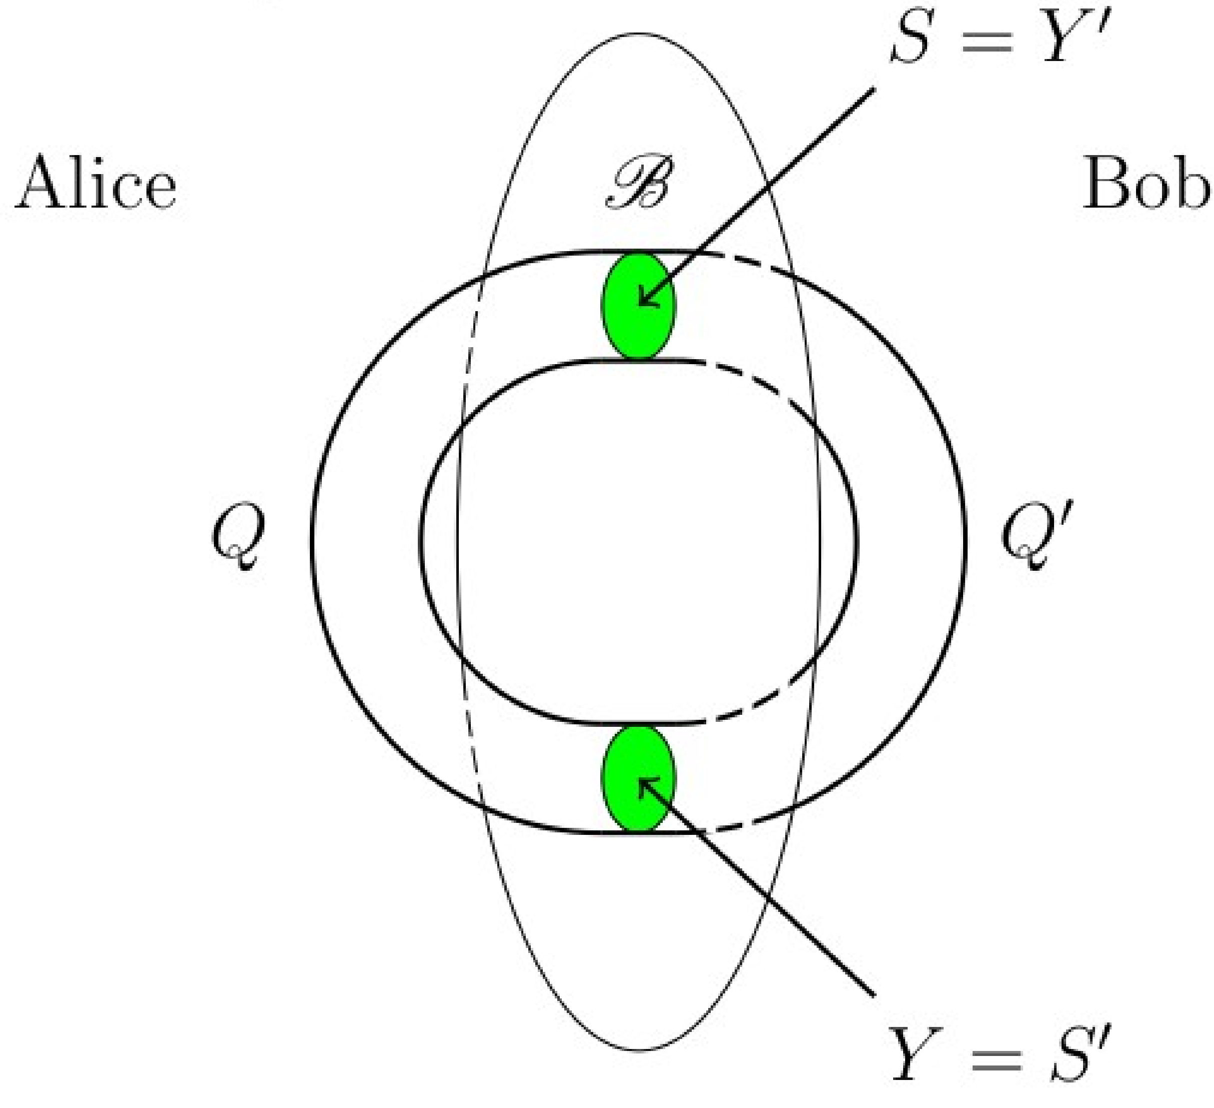

Let's impose $Q = Q'$ as a super-selection rule.

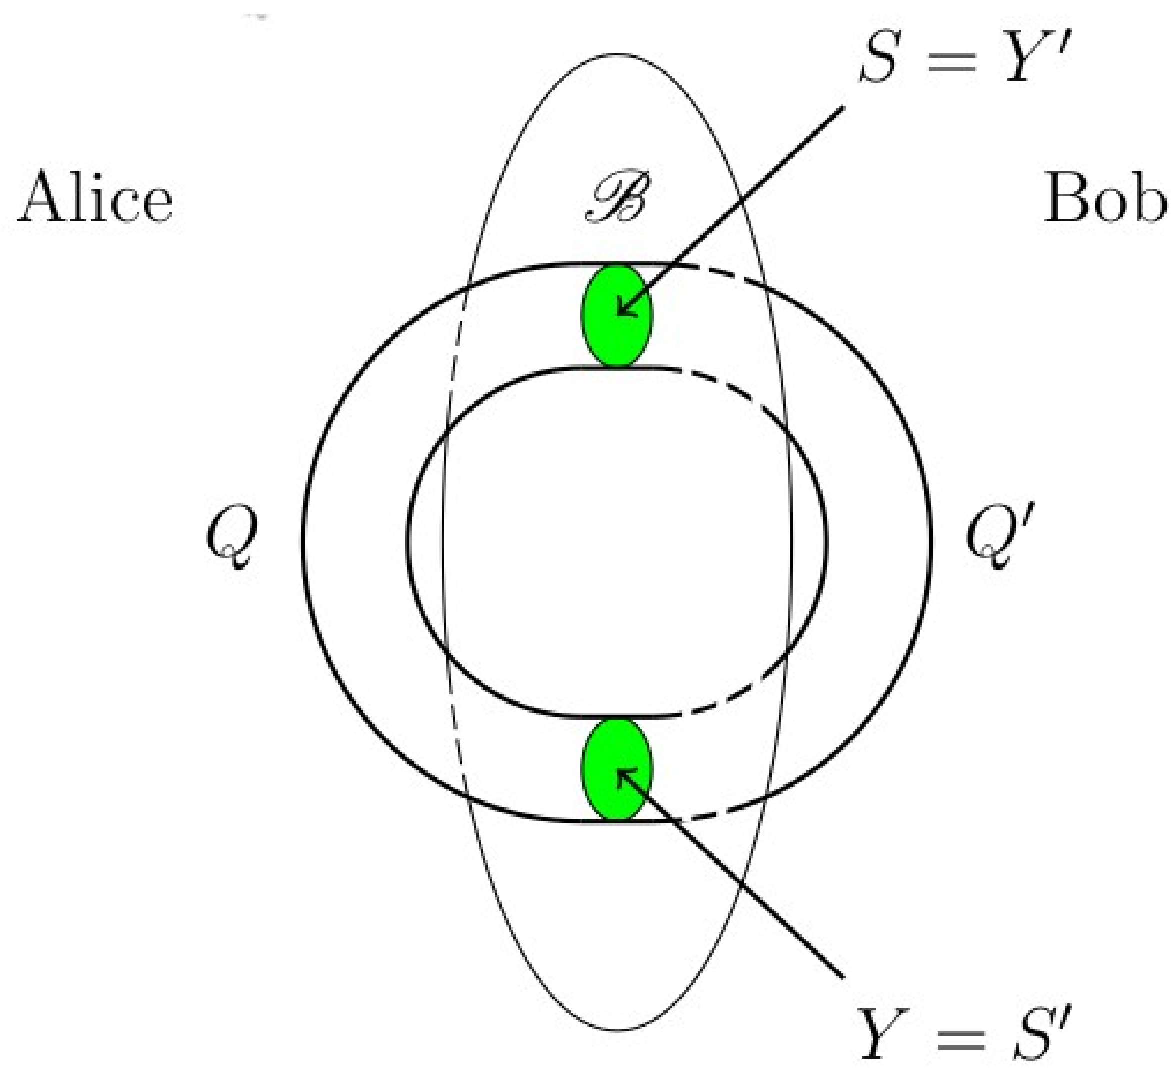

No more free choice of QRF $\longrightarrow$ Entanglement.

Consider two QRFs Q_1 and Q_2 that commute

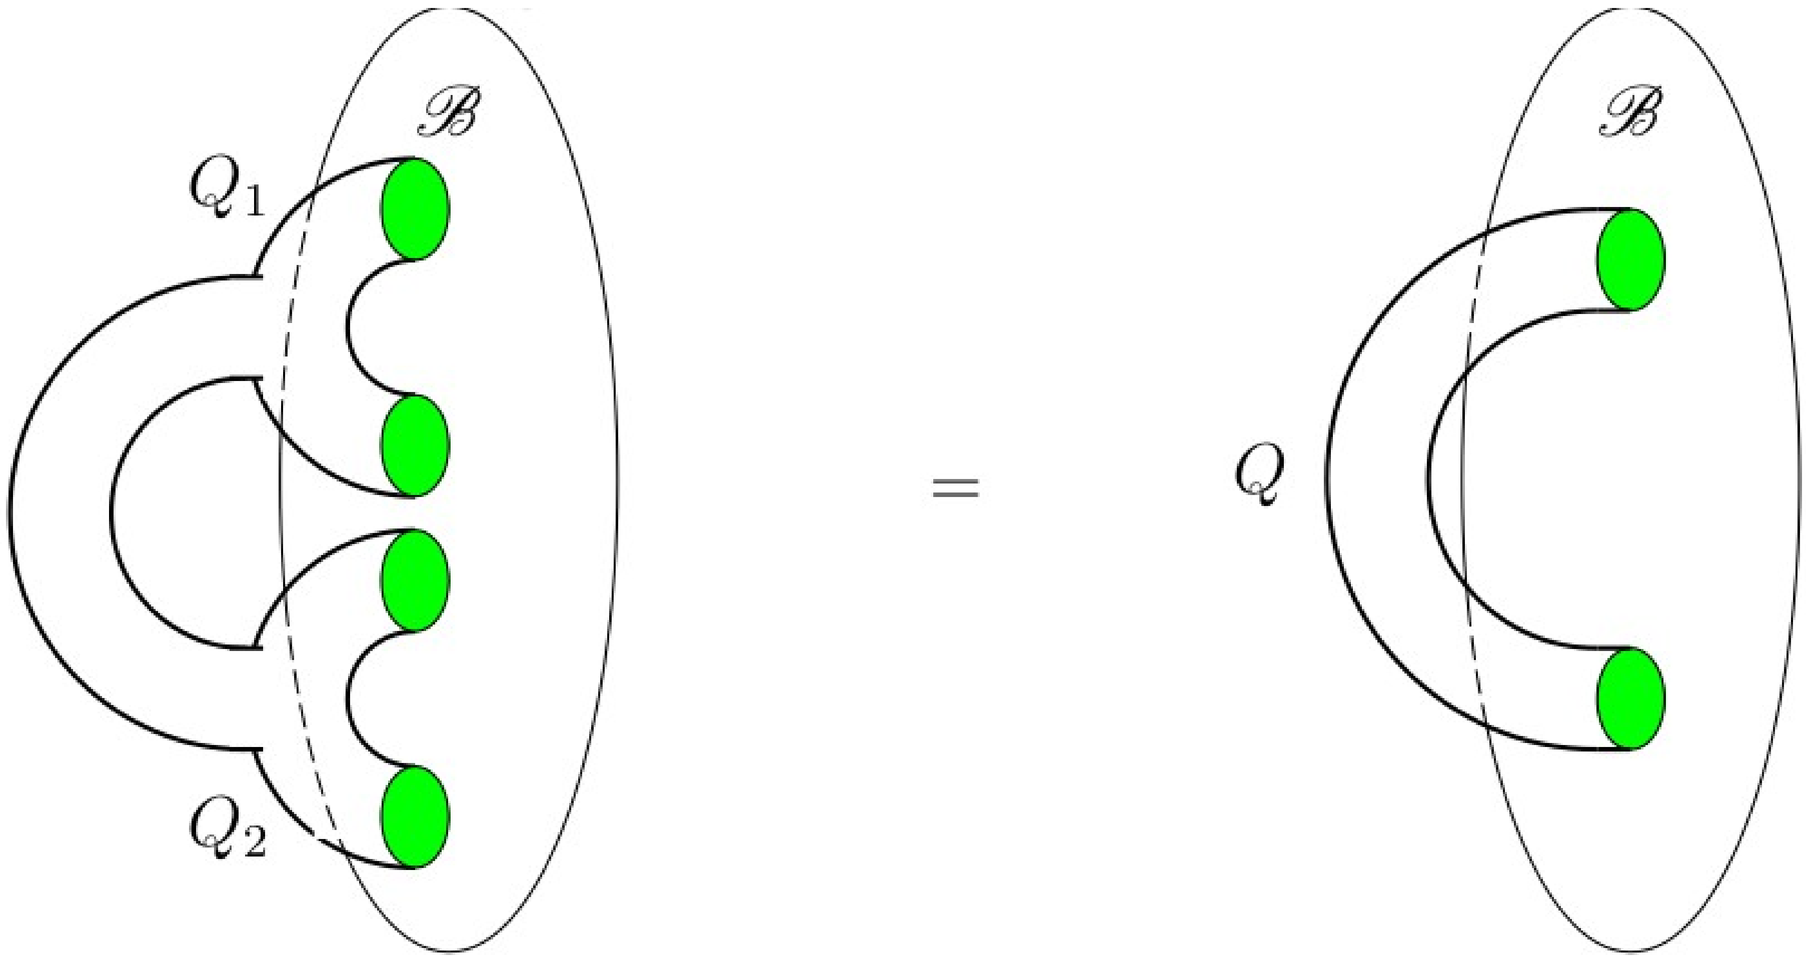

Commuting CCCDs have a common limit and colimit (CCCDs form a category).

If two QRFs Q_1 and Q_2 do not commute

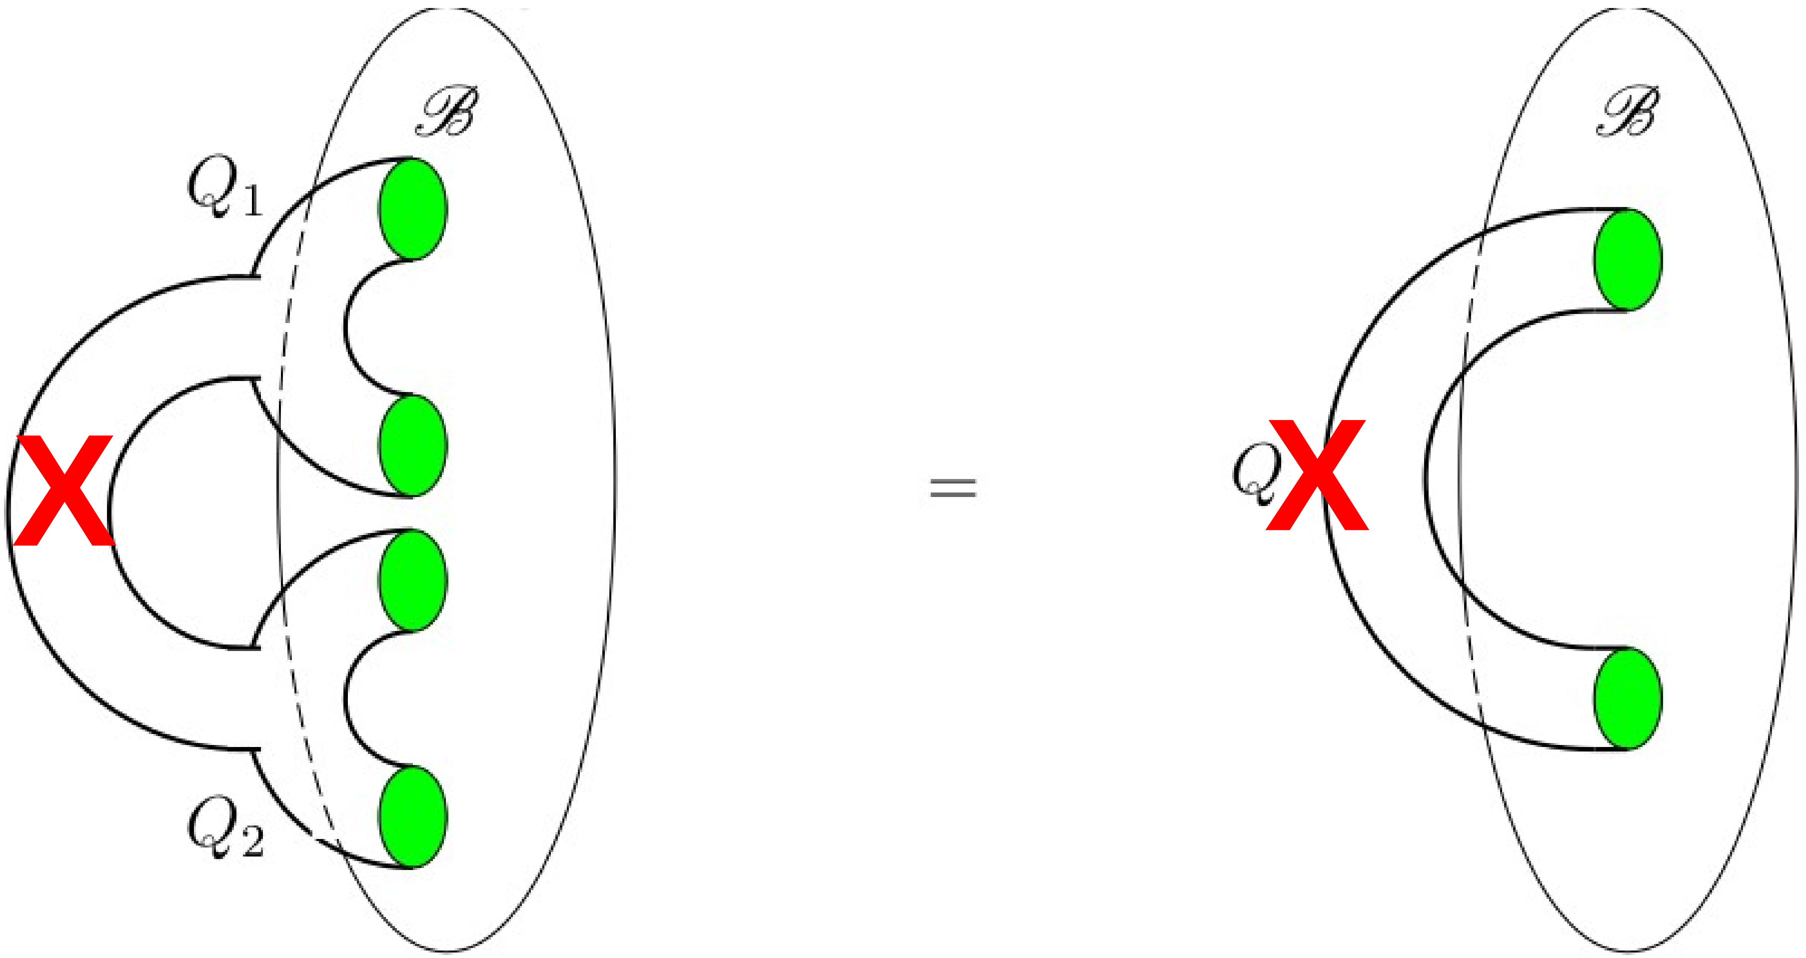

then they are not connected by a quantum channel.
So they can only be connected by a classical channel.

The classical channel can traverse the environment.

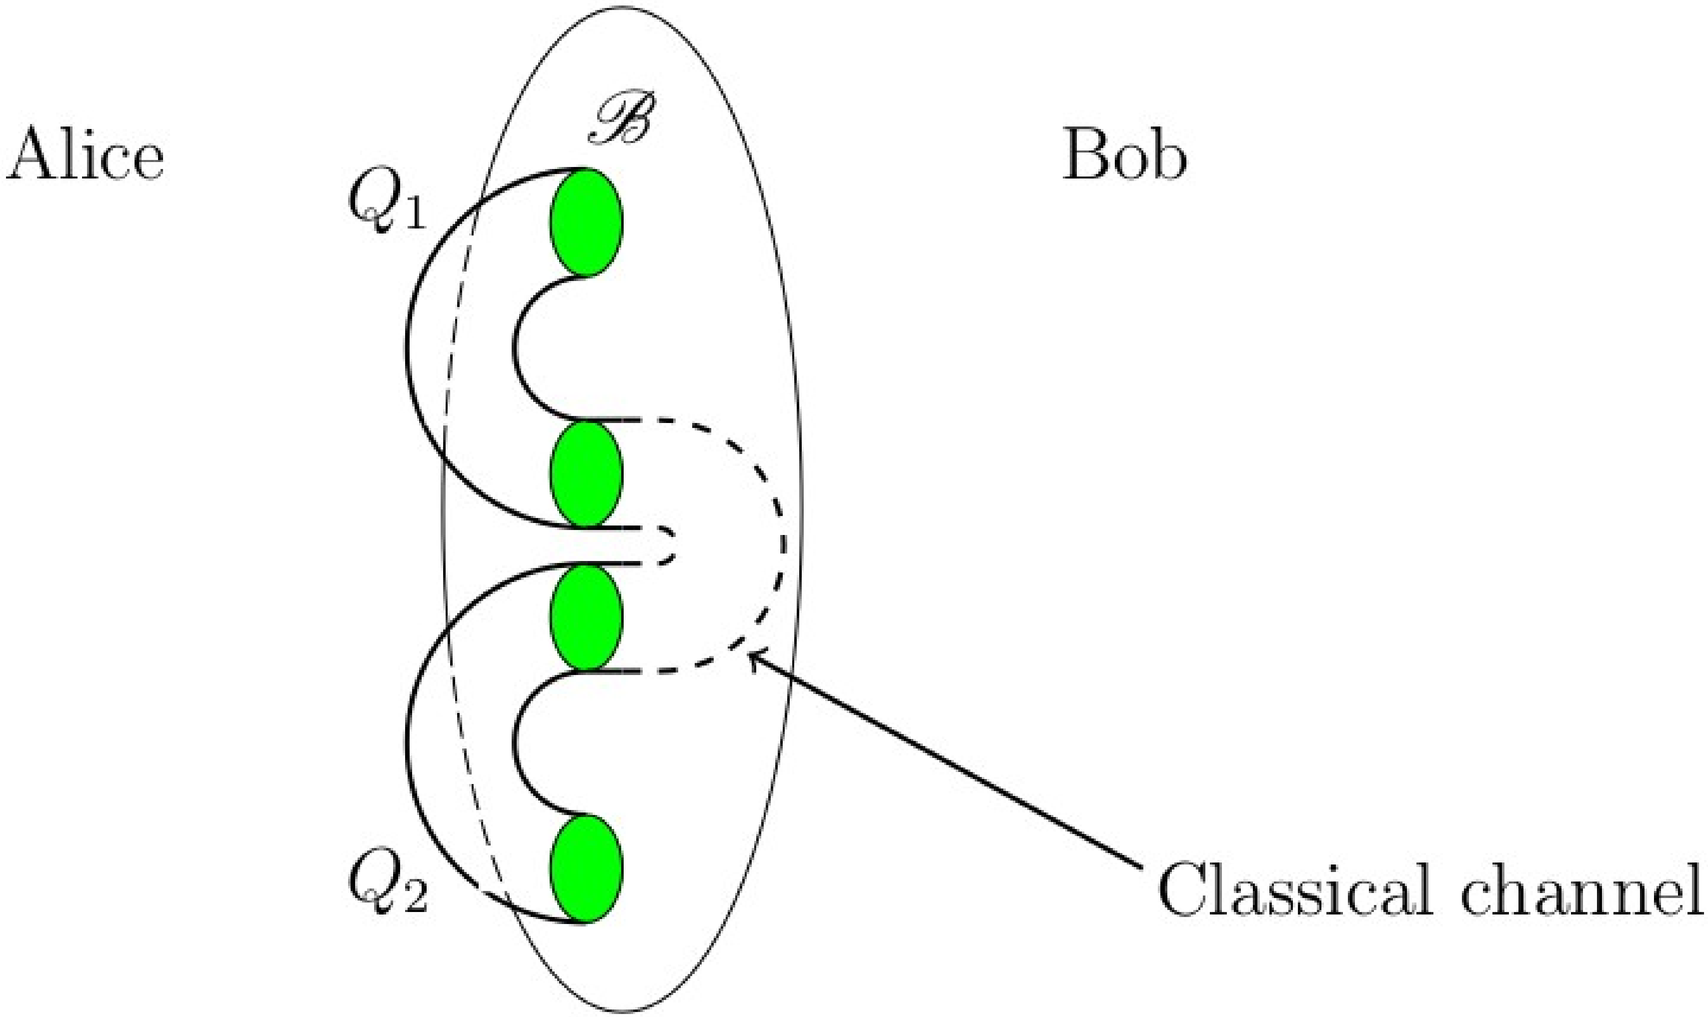

We’ve effectively split Alice into two agents.

We can also draw it this way

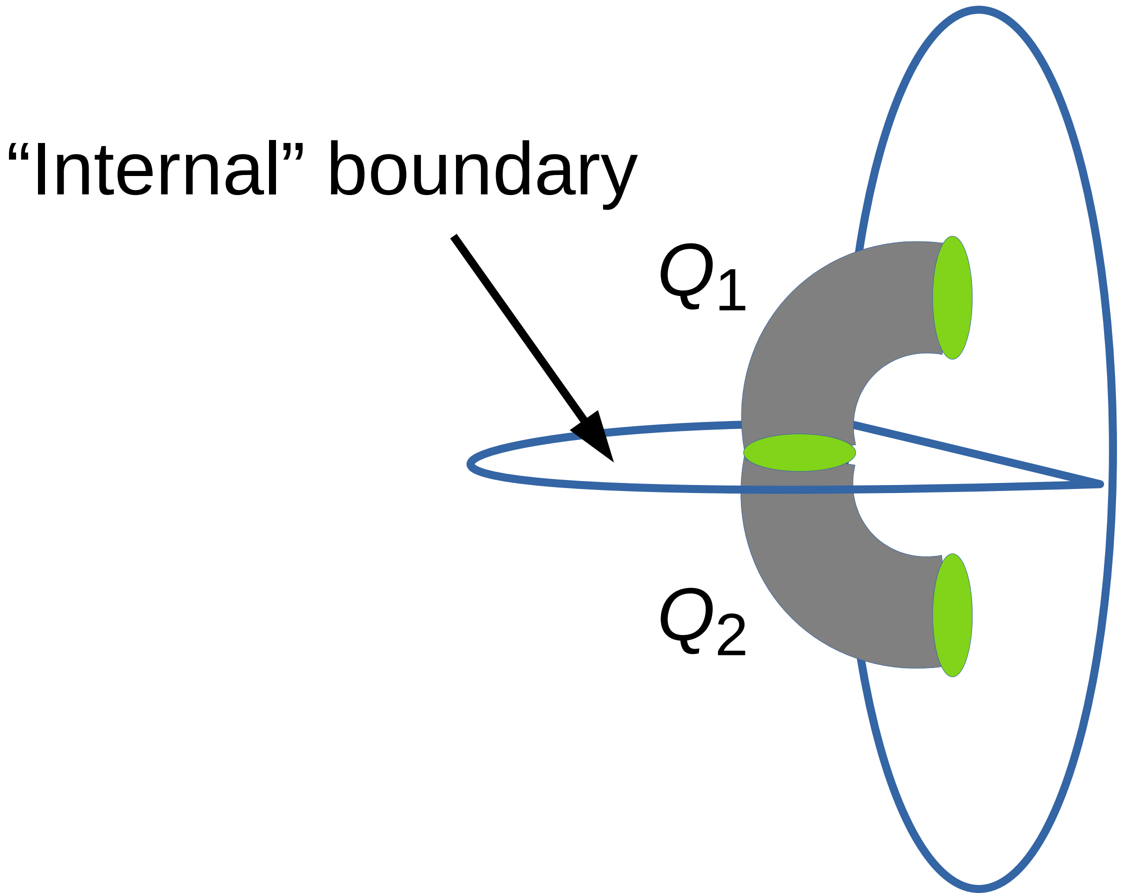

This picture helps for thinking about organisms, or about quantum computers.

We'll use this to indicate a classical message passed through the environment, e.g. an email.

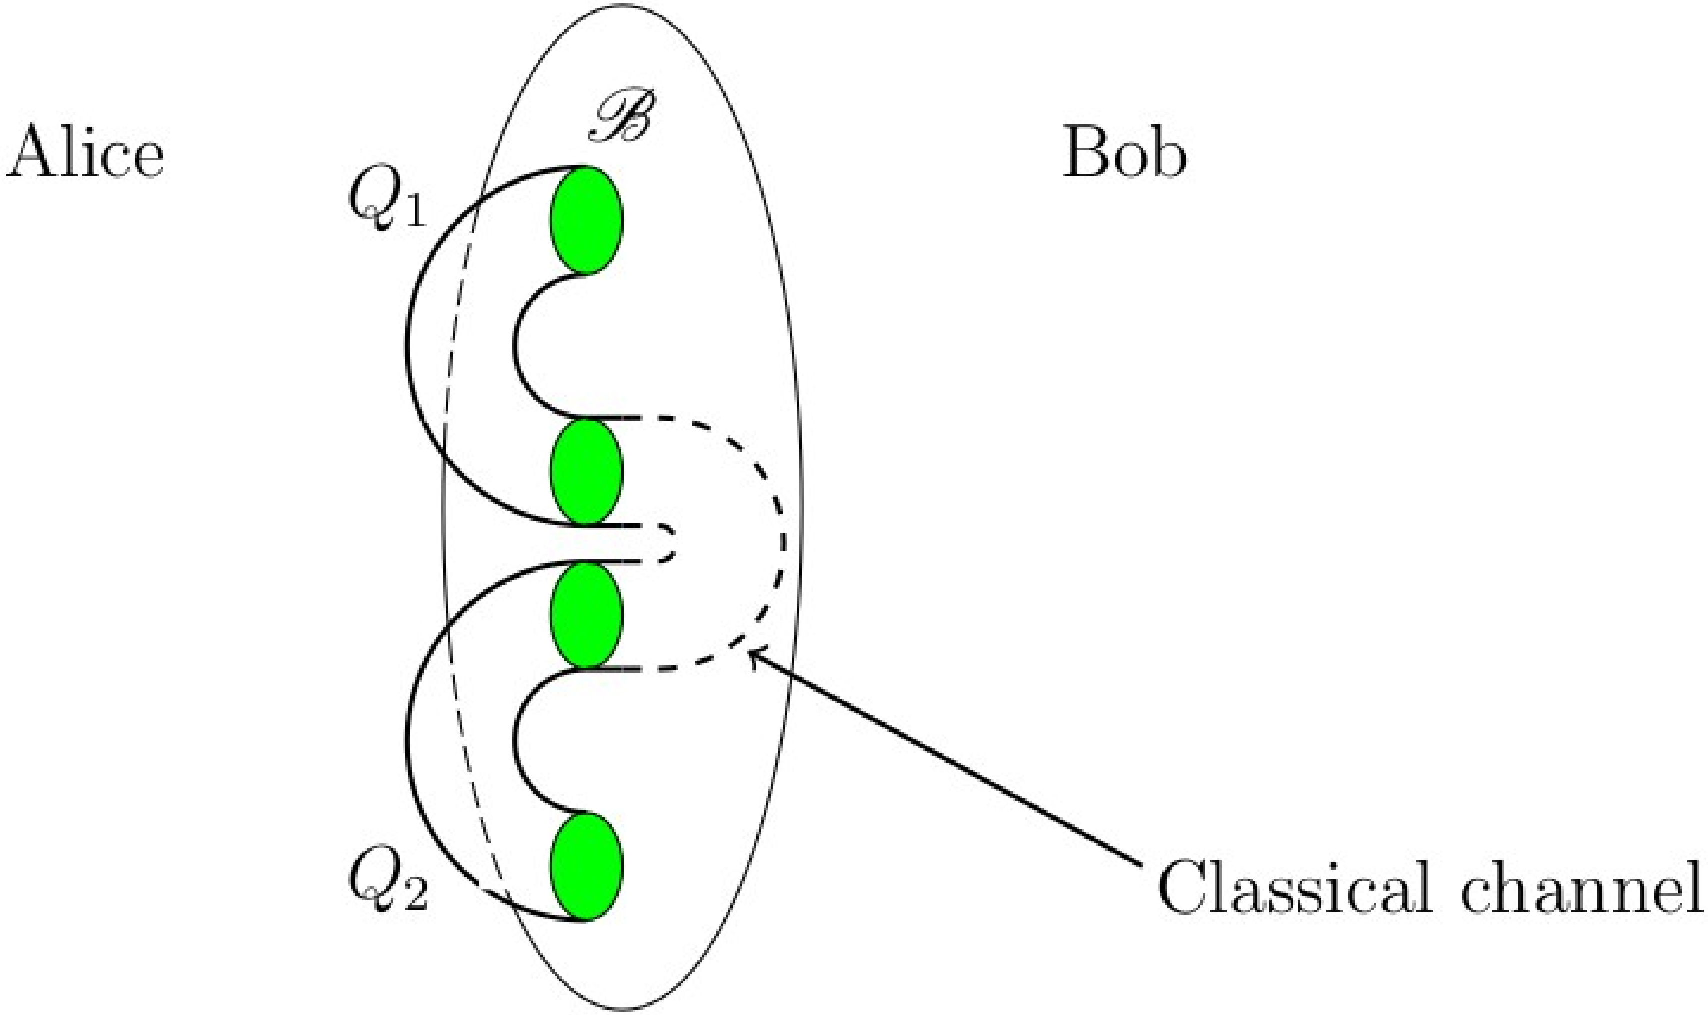

Now we add a quantum channel, e.g. an “object”

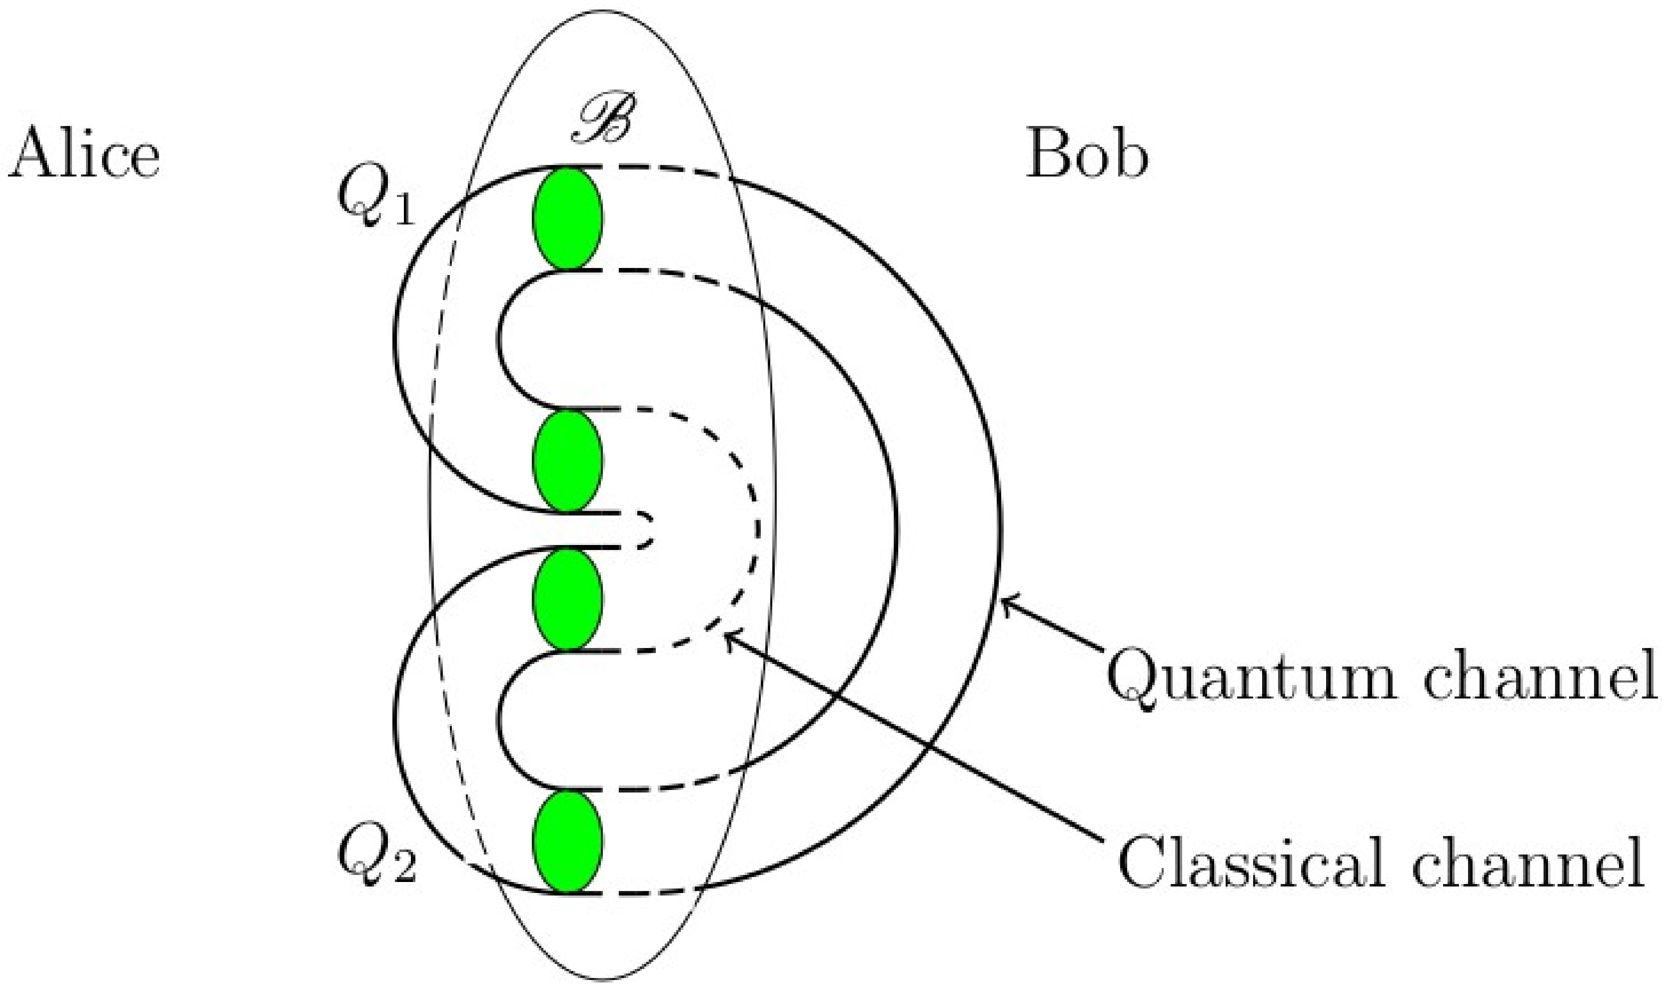

This is a “local operations, classical communication” (LOCC) protocol.

Example: Quantum Darwinism

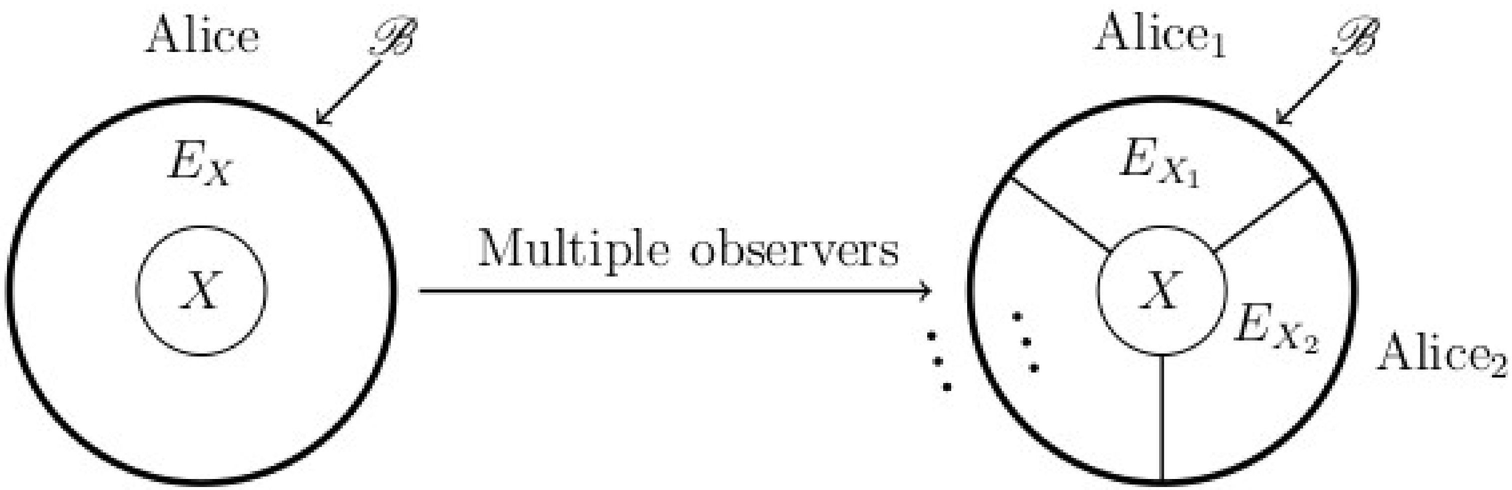

Multiple observers of the same system X must agree about what *counts as* X, i.e. what QRF R_X to deploy, and about what pointer state P_X to measure.

Example: Bell/EPR

Alice and Bob "observe entanglement" only when they compare their classical records.

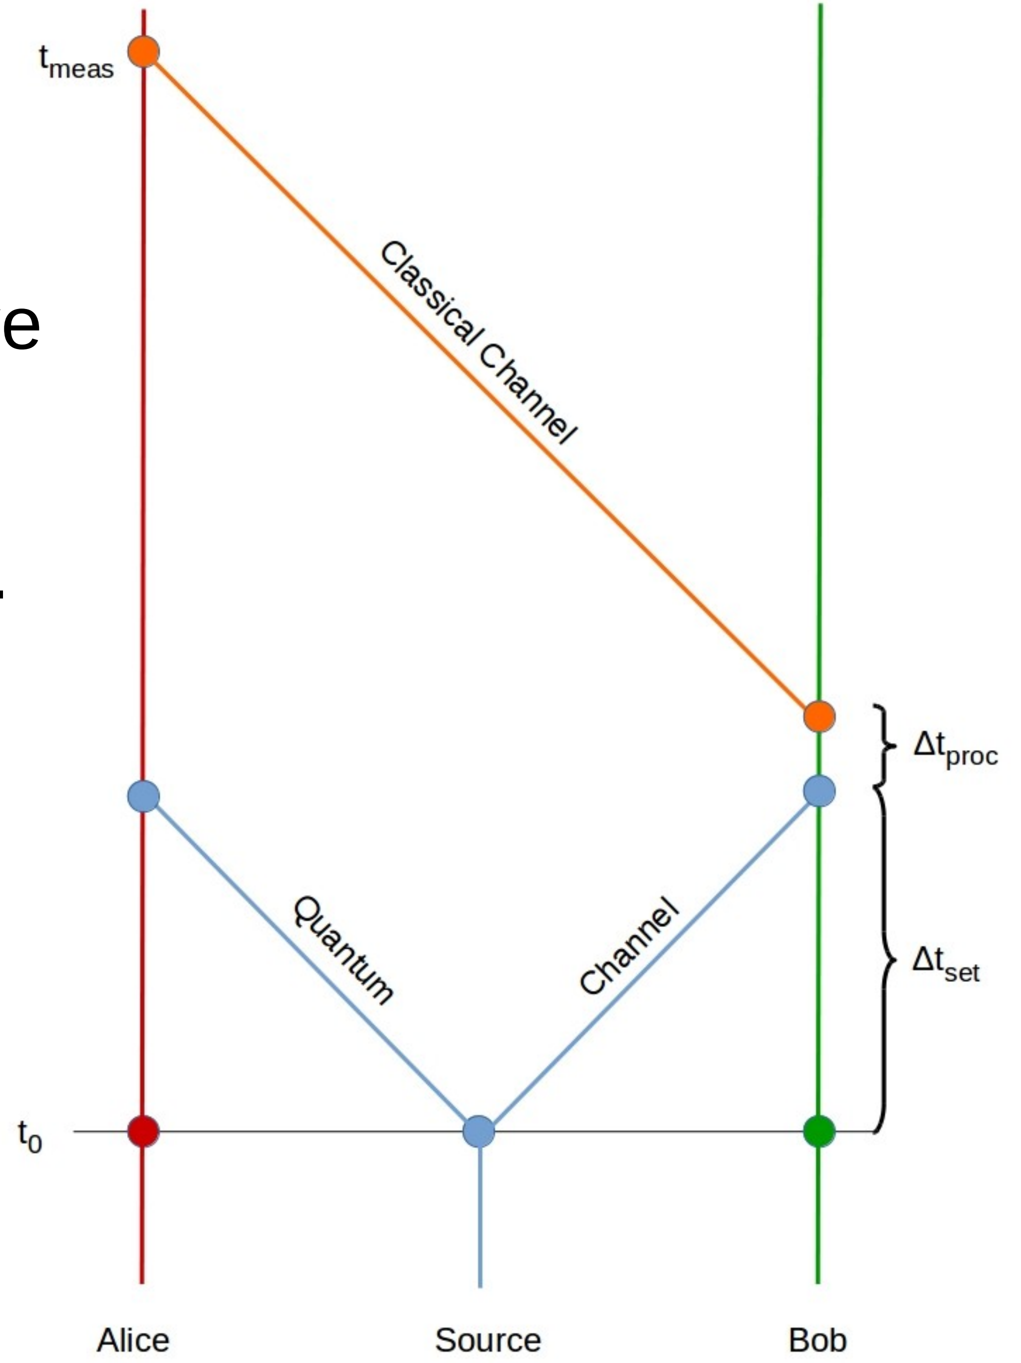

Questions?

Claim: LOCC protocols are made possible by QECCs

Main idea: Bulk entanglement generates classical redundancy on the boundary.

We can turn this around: Using a QECC requires agreement about basis choices via a LOCC protocol.

In AdS/CFT, QECCs are built from the outside in, using classical redundancy on the boundary. Here we'll work from the inside out.

Example: Quantum Darwinism

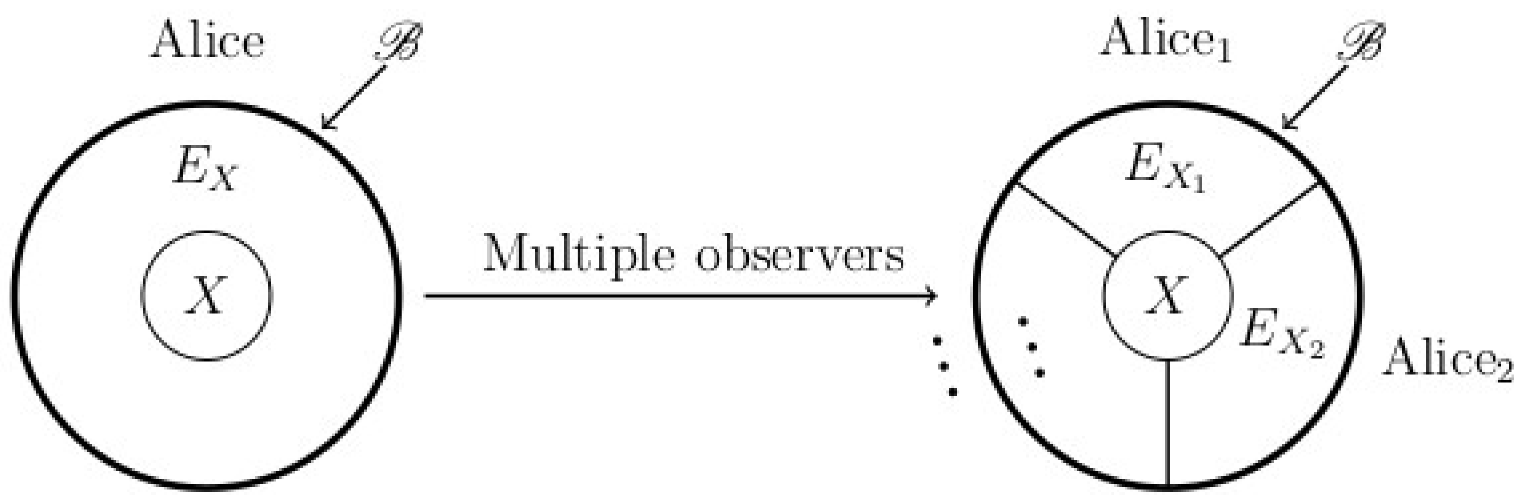

E_X is a channel that measures $|X>$ and encodes the result on Alice's observable sector of the boundary.
It does this for every observer Alice$_i$, who must employ LOCC to coordinate basis choices.

How can this work? Doesn't it violate no-cloning?

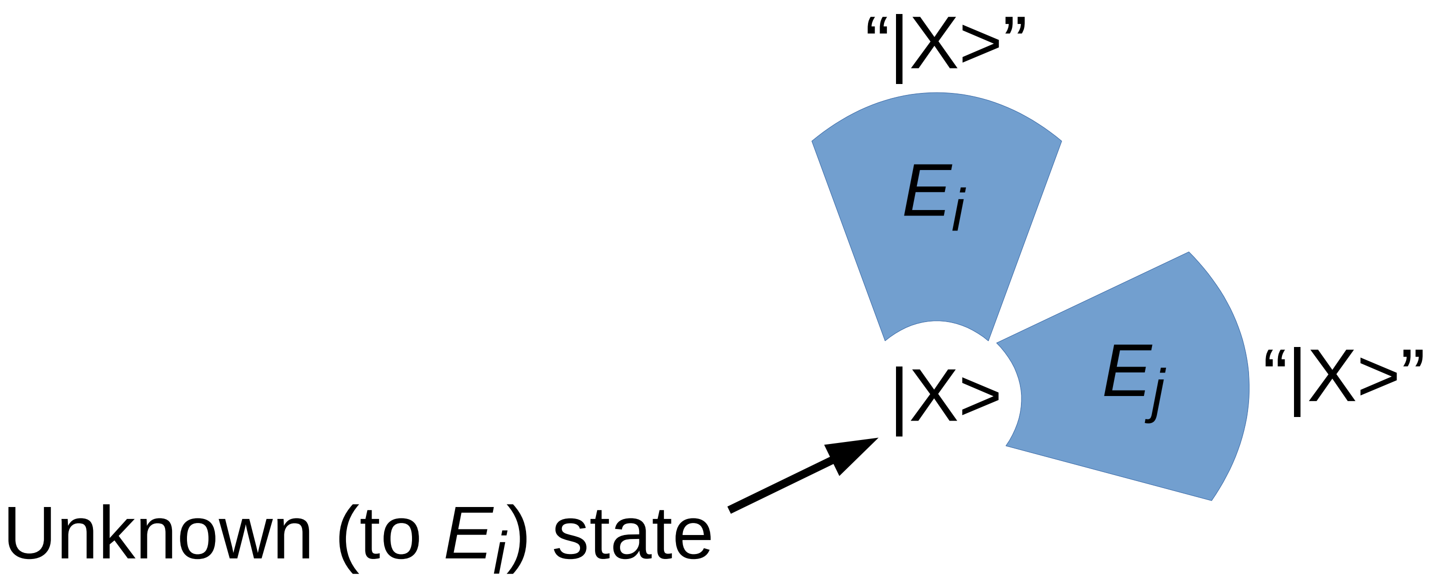

This only works if E_i's measurement resolution is much better than Alice$_i$'s, for all i. This is where Bohr's "amplification" requirement comes from.

Each environmental sector implements a TQFT

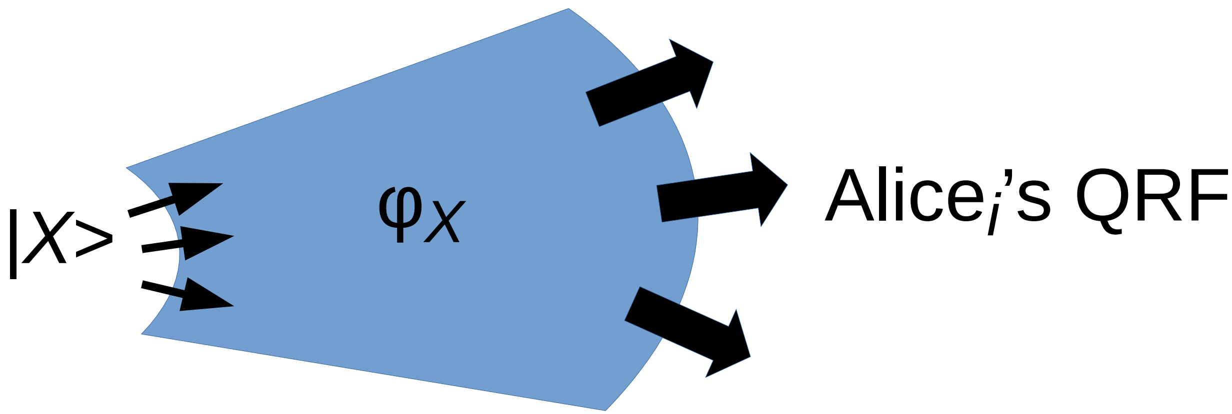

$$\varphi_X = (1/\sqrt{N})\Sigma_k^N\, |X\rangle_k \text{ with } N \longrightarrow \infty$$

Alice$_i$ is measuring a "large" projection of a much larger entangled state. She's averaging many components.

Knill and Laflamme formulate this in terms of “noise” operators B_α, with $\Sigma B_\alpha B^\dagger{}_\alpha$ = Id.

These B_α: sampled components of φ_X $\longrightarrow$
unsampled components of φ_X

The code (the TQFT) “protects against B_α” if there are enough entangled degrees of freedom to “wash out” the exchange of components.

This is recoherence – decoherence in reverse.

X is just a quantum system, so we also have:

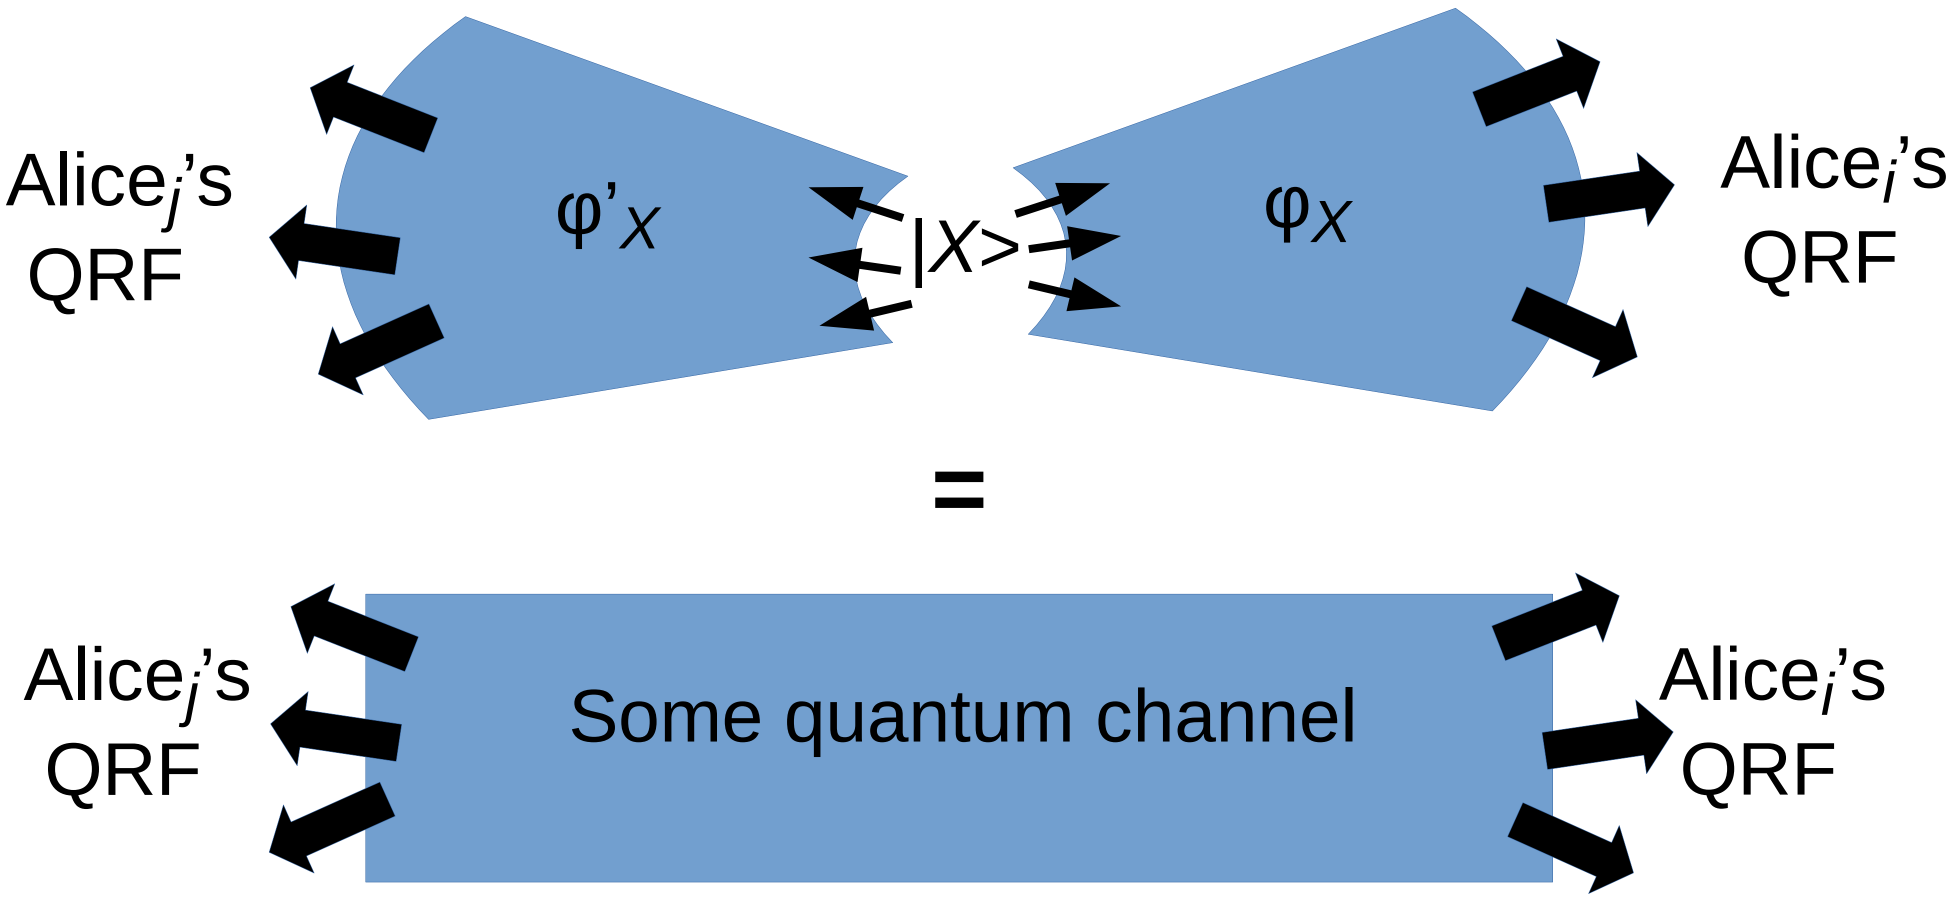

Alice$_i$ and Alice$_j$ can communicate via some channel.

We can represent all this in terms of operators:

$$S \xrightarrow{E} \mathcal{C} \xrightarrow{U_\mathcal{C}} \mathcal{C} \xrightarrow{D} S'$$

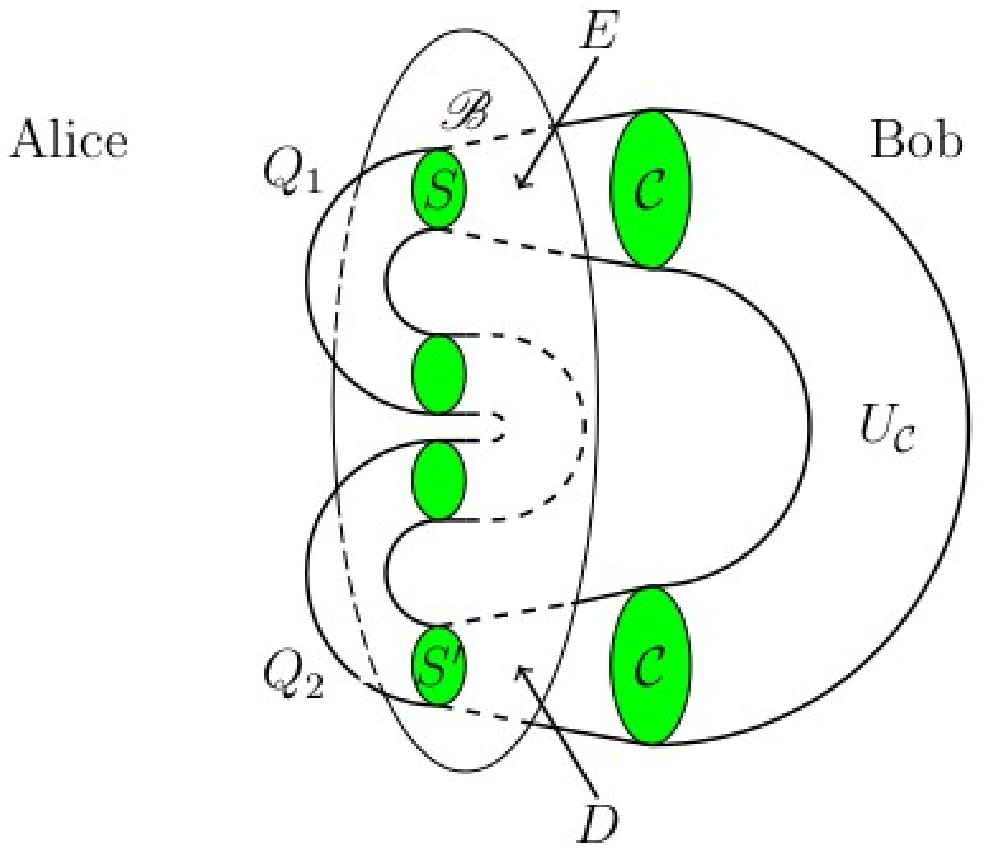

It works as long as Q_1 and Q_2 are “close enough”

How do we represent classical redundancy?

- Space-translation symmetry (spatial redundancy)
- Time-translation symmetry (temporal redundancy)

The Poincaré group represents redundancy!

Claim: Spacetime is a redundancy resource for communicating agents.

cf. tensor-network construction, but top-down.

Summary:

- Communication requires memory plus an internal time QRF.
- "Classical" communication (or coordinated manipulation) requires almost-shared bases/QRFs.
- Sampling and coarse-graining create/enable classical redundancy.
- Quantum Darwinism is a "universal" LOCC protocol.
- LOCC requires a QECC.

So what’s an agent?

Any bounded quantum system looks and acts like an agent.

Any such system can engage in LOCC protocols.

But …

No such system can determine by observation that it is bounded. Agency is an assumption!

Thank you.

Questions?