

Qualia Theorems

The math of subjectivity

If we use category theory as Jon Zingale proposed, then “intersubjectivity is a colimit of subjectivity and objectivity is a limit of interobjectivity” as Jon puts it. In plain English this means intersubjectivity is achieved by adding distinct personal voices together (the colimit), while objectivity and objective truth is distilled by filtering out personal variation (the limit).

If we use topology and Cobordism, subjective individual perspectives can be modeled as lower-dimensional manifold boundaries (e.g., (n-1)-dimensional spaces S_1 and S_2 , while the shared intersubjective space is the n-dimensional cobordism manifold W connecting them.

If we use gauge theory, we are treating objectivity not as an absolute “view from nowhere” but as the invariant structure underlying all local perspectives. Subjectivity can be described as symmetry breaking and objectivity as gauge invariance. A perspective switch for 1st-person views from person A to person B becomes a gauge transformation. For local gauge invariance to hold across different points in space, a field (like the photon field in QED) must exist to connect them. Tools or places that bring people closer together and align their views act as a similar kind of a gauge field (or “reality distortion field”). Cinemas for example are places where world views converge. All members of the audience share for a while the same world view, they view and hear the same story which influences their world view (what is good, what is bad, what should be avoided, etc.). If the story is convincing the audiences adopts part of it and adjust the own world view.

Viewpoint	Perspective Switch	Subjectivity → Intersubjectivity	Interobjectivity → Objectivity
Category Theory	Natural Isomorphism / Automorphism	intersubjectivity is a colimit of subjectivity. Colimit: distinct subjective objects are glued together along shared boundaries	objectivity is a limit of interobjectivity Limit: Filtering down to the invariant core that holds across all perspectives/ relationships
Topology (Cobordism)	A Homeomorphism acting along the boundary manifolds ∂W	Boundary to Interior: $S_1 \sqcup S_2 = \partial W$. Individual minds (S_1, S_2) are glued into a shared manifold (W).	Topological Invariants: The intrinsic geometric properties of W that remain unchanged regardless of perspective
Gauge Theory	Gauge Transformation $U(x)$	intersubjectivity is a gauge field Gauging a Local Symmetry & Introducing the Connection Field (A_μ) $D_\mu = \partial_\mu - igA_\mu$	objectivity is a gauge-invariance Gauge Orbit Quotient Space ($\mathcal{M}/G$) is the geometric manifold remaining after factoring out all local gauge degrees of freedom

The version of the hard problem that matters pragmatically is translating between two different 1st-person perspectives. If we can solve this problem, we can use it to solve the more general problem: each of us knows what it is like to be a human that experiences pleasure or pain, hunger or disgust, love or fear. These are the basic irreducible dimensions of subjective experience. If we learn what combination or mix the other person has experienced how and why and translate that 1st-person experience from the other person to ourselves, we begin to understand what-it-is-like-to-be that person.

This problem has no exact solution, only a quantified approximate one. Exact equivalence of subjective experience is generically unattainable and almost never happens; but sharing a path gives us a approximation down to a floor that is almost surely above zero.

Scale / Domain	Philosophical Concept	Gauge Theory Model	Mathematical / Physical Counterpart
1st-Person	Subjectivity	Local Gauge Fixing	Selecting a specific gauge condition (e.g., Lorenz gauge, Coulomb gauge) at x^μ .
1st A $\rightarrow$ 1st B	Perspective Switch	Gauge Transformation	Applying a phase rotation matrix $U(x)$ between two frame choices.
2nd-Person	Intersubjectivity	Covariant Coupling & Connection (A_μ)	Introducing the vector potential A_μ to preserve covariance under local transformations $D_\mu = \partial_\mu - igA_\mu$.
Systemic	Interobjectivity	Networked Gauge Field Dynamics	The full web of field interactions and field-strength tensors ($F_{\mu\nu}$) across all nodes.
3rd-Person	Objectivity	Gauge Orbit Quotient Space ($\mathcal{M}/\mathcal{G}$)	Moduli space geometry: the invariant manifold obtained by quotienting out all gauge redundancies.

Using a gauge-theoretic formulation borrowed from physics (where it is popular in high-energy physics and Quantum Field Theories) we can formulate these two theorems.

LNC Impossibility Theorem: exact equivalence almost never happens

Claim. Two independent agents almost never have exactly equivalent worldviews.

Idea. Two worldviews are equivalent if one can be turned into the other by relabeling the emotional axes, which in the model means applying a rotation Q . Picture all possible pairs of worldviews as a big space. The equivalent pairs form only a thin sheet inside it, like a sheet of paper in a room. The sheet has fewer free dimensions than the room, so a randomly chosen pair almost surely misses it.

How the proof works: we use a specific gauge-theoretic formulation from gauge theory, where the emotional valuation matrix (or valuation structure) is modeled as the gauge connection $U: E \rightarrow G$ and its associated holonomy matrices $H_k(U) \in G$. In this mathematical model, the holonomy matrices model a loop (or fundamental cycle) that represents a closed chain of experiences or reasoning that returns to the same starting concept. The holonomy matrix H_k measures the net orientation change of the agent's internal valuation axes after completing that full loop.

The Valuation Matrix (V): Represents the agent's internal frame or coordinate system for converting a situation into a set of emotional valences (e.g., along axes of pleasure/pain or approach/avoidance).

The Gauge Connection ($U_e \in G$): Each edge e in the context graph represents a specific cognitive or emotional transition (e.g., encountering a situation, stimulus, or memory). The matrix assigned to that edge, U_e , acts as an operator in a compact Lie group G (such as $SO(3)$). In human terms, it describes how an individual's internal state rotates or shifts its “valuation axes” when moving along that experience

The Observables / Holonomies ($H_k(U) \in G$): When an agent undergoes a closed loop of experience (a cycle in the context graph, such as reflecting on a recurring life event or concept), the product of the edge matrices along that cycle forms a holonomy matrix $H_k(U) \in G$.

If an agent transitions through a series of situations or choices along a path in their cognitive context graph (Edge 1 $\rightarrow$ Edge 2 $\rightarrow \dots \rightarrow$ Edge n), their valuation frame updates sequentially:

$$V_{\text{final}} = U_{e_n} \cdot U_{e_{n-1}} \cdots U_{e_1} \cdot V_{\text{initial}}$$

When that sequence forms a **closed loop** or cycle k (returning to the same baseline situation or concept) the Holonomy Matrix (H_k) is the product of all frame-shift operators along that loop:

$$H_k = U_{e_n} \cdot U_{e_{n-1}} \cdots U_{e_1}$$

These holonomies represent the agent's net perspective shift across fundamental value loops. They describe the path-dependent change of the emotional valuation matrices. The agent's updated valuation matrix after completing the cycle is:

$$V_{\text{after loop}} = H_k \cdot V_{\text{before loop}}$$

What makes it “Unique” to an individual: an individual's unique worldview is the total set of these holonomy matrices ($H_1, H_2, \dots, H_r$). Two people have the *same* emotional valuation structure if and only if there exists a global change-of-basis matrix $Q \in G$ that aligns all their loop matrices simultaneously ($H_k(Y) = QH_k(X)Q^{-1}$ for all k).

The entire point of the LNC Impossibility Theorem is proving that because these emotional matrices $H_k(U)$ inhabit a continuous Lie group like $SO(3)$ or $SO(N)$, no two independent human minds will ever possess the exact same set of valuation matrices. The proof evaluates whether two agents A and B have equivalent subjective value systems, and asks: can we find a single frame shift Q that aligns Agent A's valuation changes with Agent B's valuation changes across all life loops simultaneously ($H_k^B = QH_k^A Q^{-1}$ for all k)? Because the holonomies H_k capture the path-dependent evolution of each person's unique valuation system over complex life cycles, their unalignable angles guarantee that no single transformation Q can map one person's value calculation system onto another's ($d(A, B) > 0$).

Assumptions

Label	Assumption
G	A compact connected Lie group with $\dim G \geq 1$ (for example $SO(3)$), acting by a faithful unitary representation, so that the Frobenius norm is invariant under multiplication by group elements.
Graph	A finite connected context graph with $r = E - V + 1$ independent cycles. We require $r \geq 1$: if the graph has no cycles, every connection is gauge-trivial and every pair of agents is equivalent.
Worldview	A connection $U: E \rightarrow G$, with U on a reversed edge equal to the inverse. Values are exact (infinite precision).
A1 (density)	The holonomies of the two agents have a joint density with respect to Haar measure (independent, continuously distributed). This is an assumption about the world, not a derived fact.
Bridge	Equivalent experience corresponds to gauge equivalence of valuation structure. This is assumed, not derived.

Classification and Distance

Fix a spanning tree and let $H_k(U) \in G, k = 1, \dots, r$, be the holonomies of the fundamental cycles (tree path plus one non-tree edge) at a base vertex.

Fact 1 (standard). Two connections X and Y are gauge equivalent if and only if there is $Q \in G$ with $H_k(Y) = QH_k(X)Q^{-1}$ for all k .

Define the distance:

$$d(X, Y) = \min_{Q \in G} \max_k \|H_k(Y) - QH_k(X)Q^{-1}\|$$

- The minimum is attained because G is compact.
- $d = 0$ if and only if X and Y are gauge equivalent.
- d satisfies the triangle inequality (compose the rotations and use the invariance of the norm), so it is a metric on the quotient of connections by gauge.

The Theorem

The LNC Impossibility Theorem says two independent agents almost never have exactly equivalent worldviews X and Y :

Assume $r \geq 1$ and A1. Then $P(X \cong Y) = 0$, and therefore $d(X, Y) > 0$ almost surely.

Proof

Let $\Phi: G \times G^r \rightarrow G^r \times G^r$ be $\Phi(Q, h) = (h, QhQ^{-1})$. The set E of equivalent holonomy pairs is the image of Φ . It is compact (a continuous image of a compact set), hence measurable. It suffices to show E has Haar measure zero, since a joint density assigns probability zero to such sets.

Step 1: The kernel of the differential. At a point (Q, h) , the differential of Φ vanishes on a direction $(\xi, 0)$ exactly when ξ commutes with every h_i , that is, ξ lies in the Lie algebra of the centralizer of h . Writing $s(h)$ for its dimension, the rank of $d\Phi$ is $(r + 1) \cdot \dim G - s(h)$.

Step 2: Case $r \geq 2$. The domain has dimension $(r + 1) \cdot \dim G$, which is smaller than the target dimension $2r \cdot \dim G$. The rank is below the target dimension at every point, so every point is critical. By Sard's theorem, the image has measure zero.

Step 3: Case $r = 1$. Domain and target have the same dimension $2 \cdot \dim G$. Every element h of a compact connected group lies in a maximal torus, whose Lie algebra commutes with h and has dimension $\text{rank}(G) \geq 1$. Hence $s(h) \geq 1$, the rank is at most $2 \cdot \dim G - 1$ everywhere, every point is critical, and Sard's theorem again gives measure zero.

Step 4: $d > 0$. Since the minimum in d is attained, $d = 0$ would mean exact equivalence, which has probability zero. ■

Numerical check of the dimension counts

Counting dimensions is the key for the proof: One worldview is modeled as r fundamental cognitive agentic loops in form of holonomies. Since $\text{SO}(3)$, the symmetry group of 3D rotations, has 3 parameters like pitch, roll, yaw, one worldview takes $3r$ numbers to specify. Two independent worldviews X and Y together require $3r + 3r = 6r$ free numbers. This $6r$ -dimensional space of all pairs is our “room”. X and Y are equivalent only if there is a single perspective transformation Q (3 parameters) that lines up all r loops of X with Y simultaneously. An equivalent pair is built by choosing one worldview (r holonomies, 3 free numbers each in $\text{SO}(3)$) and one rotation Q (3 numbers), giving $3r + 3$ free numbers. For $r \geq 2$,

$3r + 3 < 6r$, so the sheet is thinner than the room. We could say the worldview is an own ‘slice of the world’. Two randomly chosen sheets or slices almost surely are not identical. For $r = 1$ the counts tie, but rotations around a holonomy’s own axis leave it unchanged, which makes the sheet effectively one dimension thinner.

Check of Dimension of $E = (r + 1) \cdot \dim G - s(h) = (r + 1) \cdot 3 - s(h)$ where $G = \text{SO}(3)$, so $\dim G = 3$

$s(h)$ is the set of commuting directions. It is defined as the dimension of the Lie algebra of the joint centralizer (or stabilizer) of the holonomies $(h_1, \dots, h_r)$ inside G

$$s(h) = \dim\{\xi \in \mathfrak{g} \mid \xi h_i = h_i \xi \text{ for all } i = 1, \dots, r\}$$

r	Domain dimension 3r+3	Target dimension 6r	Dimension of E (Jacobian rank)
1	6	6	5
2	9	12	9
3	12	18	12
4	15	24	15

The numbers match the formula $(r + 1) \cdot \dim G - s$ with $s(h) = 1, 0, 0, 0$ respectively. For a finite group G , or without **A1**, the statement fails, and the probability of equivalence is positive.

If we model internal emotional or cognitive states using an N-dimensional Lie group like $\text{SO}(N)$ (where $N = 10, 100$, or 1000), then the dimension of the rotation group $\text{SO}(N)$ is given by the formula:

$$\dim \text{SO}(N) = \frac{N(N-1)}{2}$$

As human cognitive or AI valuation spaces grow into hundreds or thousands of dimensions, the space of possible worldviews expands so rapidly that the slice of “equivalent worldviews” becomes infinitesimally thinner.

Dimensions N	Lie Group G	dim G = dim SO (N)	Domain dimension 3 dim G	Target dimension 4 dim G
3	SO (3)	3	9	12
10	SO (10)	45	135	180
100	SO (100)	4,950	14,850	19,800
1,000	SO (1,000)	499,500	1,498,500	1,998,000

Asymptotic Qualia Theorem: asymptotic approximation is possible

Claim. If two agents share a path, the gap between two agents' worldviews converges to a floor, and the floor is almost surely above zero. Sharing a path gives asymptotic approximation down to a floor

The theorems states that while two agents can never reach exact gauge equivalence ($d(A, B) = 0$), they can progressively reduce their unresolvable interpretation gap down to a strictly positive floor ($\epsilon_0 > 0$) by repeatedly sharing contextual paths and experiences.

Idea: Each agent holds one rotation matrix per transition between contexts. Every step of experience nudges some of these matrices toward the agent's long-run target. On transitions both agents share, the targets are identical. On the others, the targets are different random draws. Because each holonomy is a product of a bounded number of edge matrices, small errors on the edges become small errors on the holonomies, and the distance between the agents settles to the distance between the targets.

That final distance cannot be zero except by coincidence: exact alignment would need one rotation Q that leaves the shared part unchanged and also lines up the random part, and such Q exist only for a thin, lower-dimensional set of target pairs, which a random draw almost surely misses.

How the proof works:

When an agent lives through a sequence of experiences it moves along a path $\gamma = (e_1, e_2, \dots, e_M)$ where each edge e_m is a specific step of a trajectory, a shift from one context to another. Edges represent transitions between two contexts/situations. Experiencing an edge e is what triggers the adaptive agent to update their frame/worldview via the discrete transition matrix $U_t(e) \in G$. It is the specific event that nudges the agent's internal frame toward their long-term target $U_\infty(e)$.

We distinguish between shared and unshared experience.

- Unshared experiences I_u are represented by unshared edges (N)
- Shared experiences I_s are represented by shared edges (S) which agents A and B share

The distinction between shared and unshared edges is the mathematical secret to why worldviews get closer without ever becoming identical. Shared Edges (S) are experiences where both agents face the exact same environmental reality (e.g., physical laws, shared culture, living together, or undergoing the same crisis). On these edges, both agents share the exact same target:

$$U_\infty^A(e) = U_\infty^B(e)$$

Unshared Edges (N) are experiences where individual genetics, distinct personal histories, or unique internal biological states dictate the outcome. On these edges, the targets are independent random draws:

$$U_\infty^A(e) \neq U_\infty^B(e)$$

The more shared edges we have, the more the worldviews of the two agents converge. In everyday language, “sharing a path” sounds like two people simply walking down the same street or living in the same house. In this formal mathematical model, “sharing a path” is defined as encountering a sequence of edges where both agents are exposed to the same environmental driving force, meaning they learn toward

the exact same target matrices.

1. **Environmental Forcing:** When two agents undergo a common experience (e.g., growing up in the same culture, same family or same place, going through a shared trauma, or engaging in continuous dialogue), the external environment imposes a common objective target $U_\infty(e)$ for those shared steps.
2. **Convergence Driver:** Because they share the target on those edges, each step along that shared path nudges both Agent A's matrix $U_t^A(e)$ and Agent B's matrix $U_t^B(e)$ toward the same destination.

This shared target dynamic is what drives the exponential convergence of their worldviews down toward the floor ϵ_0 . On unshared steps, where their targets differ, their individual substrates retain unique orientations—guaranteeing that the final floor ϵ_0 remains strictly greater than zero.

The Connection (U_t^a) as a worldview: Assigns a specific rotation matrix $U_t(e) \in G$ to every single edge $e \in E$ and agent $a \in \{A, B\}$. This connection dictates how the agent shifts their valuation frame when transitioning from one context to another. An agent's connection is the collection of all their rotation matrices - one matrix for every edge (transition) in their mental context graph.

- full connection U_t (the collection of all edge matrices) is the entire worldview map
- a single edge matrix $U_t^a(e)$ (Current Connection): The rotation matrix an agent a holds at step t for a specific experience e represents a worldview change
- $U_\infty^a(e)$ (Limiting Connection / Target): The ultimate, long-term matrix target that agent a 's mind is learning toward. It represents the final, fully formed perspective or value the agent will assign to that experience after unlimited learning.

Crucially, each agent has their own limiting connection (U_∞^A and U_∞^B). They are not forced to arrive at the same destination on every topic

Lemma 0 provides a concrete mechanism in form of an update rule for how learning happens step-by-step: at time step t , the agent experiences a specific edge e_t in life. The agent takes their current matrix $U_t(e_t)$ and moves it a fraction λ (e.g., 20%) along the shortest path (geodesic) toward their target $U_\infty(e_t)$. All other edge matrices in their mind remain as they were. Because every life transition is revisited periodically (every m steps), the distance between an agent's current matrix and their target matrix shrinks exponentially by a factor of $(1 - \lambda)$ with each revisit. Mathematically, this guarantees that their current connection U_t converges exponentially toward U_∞ at rate γ .

Setup and Notation

- **Context Graph:** A finite connected graph $\Gamma = (V, E)$ with $r = |E| - |V| + 1 \geq 1$ independent cycles. (With $r = 0$ every pair of agents is equivalent.)
- **Group:** A compact connected Lie group G with $\dim G \geq 1$ (for example $SO(3)$), acting by a faithful unitary representation, so that the Frobenius norm is invariant under multiplication by group elements.
- **State:** At each discrete time $t = 0, 1, 2, \dots$, an agent $a \in \{A, B\}$ holds a connection $U_t^a: E \rightarrow G$, one matrix per edge, with the reverse of an edge carrying the inverse (transpose). The experience steps are what update this connection.

- **Holonomies:** Fix a spanning tree and a base vertex. For each non-tree edge $e_k = (u_k \rightarrow v_k)$, $k = 1, \dots, r$, the holonomy is:

$$H_k^a(t) = T_{v_k}^{-1} \cdot U_t^a(e_k) \cdot T_{u_k}$$

where T_w is the ordered product of the tree-edge matrices from the base vertex to w . Let L be the maximum number of edge matrices in these words.

- **Distance:** if X, Y are the two agents' connections (their worldviews),

$$d(X, Y) = \min_{Q \in G} \max_k \|H_k^B - QH_k^A Q^{-1}\|$$

The minimum is attained (G is compact), $d = 0$ if and only if the agents are gauge equivalent, and d is a metric on the gauge quotient. Evaluating d at the limiting connections X_∞ and Y_∞ means:

$$\epsilon_0 = d(X_\infty, Y_\infty) = \min_{Q \in G} \max_k \|H_k^B(\infty) - QH_k^A(\infty)Q^{-1}\|$$

Assumptions

Label	Assumption
A1 (Exponential Convergence)	Each agent has a limiting connection U_∞^a , and for all edges e , both agents $a \in \{A, B\}$, and all t : $\ U_t^a(e) - U_\infty^a(e)\ \leq \delta_0 e^{-\gamma t}, \quad \text{with } \delta_0, \gamma > 0$ The two agents have separate limits. Lemma 0 gives a concrete update rule that implies A1.
A2 (Shared and Unshared Structure)	The edge set splits into shared edges S and unshared edges N . The limits coincide on shared edges: $U_\infty^A(e) = U_\infty^B(e)$ for $e \in S$. Let I_s be the indices k whose holonomy word uses only shared edges (so $H_k^A(\infty) = H_k^B(\infty) = : h_k$), and I_u the others, with $r_u = \ I_u\ \geq 1$.
A3 (Density)	Conditional on the shared limit holonomies $h_s = (h_k)_{k \in I_s}$, the pairs $(H_k^A(\infty), H_k^B(\infty))$ for $k \in I_u$ have a joint density with respect to Haar measure on $G^{r_u} \times G^{r_u}$.
Exactness	All values are exact (infinite precision).
Bridge	Equivalent experience corresponds to gauge equivalence of the valuation structure. This is assumed, not derived.

Sufficient edge-level condition for A3: Suppose the shared edges contain a spanning tree, and the limits $U_\infty^A(e), U_\infty^B(e)$ for each unshared non-tree edge e are independent with a density with respect to Haar measure, independent of the shared edges. Then each $H_k^a(\infty)$ ($k \in I_u$) equals a fixed (shared) left factor times $U_\infty^a(e_k)$ times a fixed (shared) right factor, and Haar measure is invariant under such translations. Hence $(H_k^A(\infty), H_k^B(\infty))$ has a joint density, and A3 holds.

Lemma 0: A Concrete Learning Rule That Implies A1

Assumption A1 can be derived from a simple update rule, instead of being taken as given. Lemma 0 derives exponential convergence from a concrete contractive update with revisits. Each time the agent experiences edge e , Lemma 0 specifies that $U_t(e)$ moves a fraction λ along the shortest geometric path (geodesic) toward $U_\infty(e)$

Lemma 0. Let d_{geo} be the geodesic distance for a bi-invariant metric on G , with diameter D , and let C be a constant such that $\|\rho(g) - \rho(h)\| \leq C d_{\text{geo}}(g, h)$. At step t , the agent traverses one edge e_t and moves $U_t(e_t)$ along a minimizing geodesic toward its target $U_\infty(e_t)$, covering a fraction $\lambda \in (0,1)$ of the remaining distance; all other edges are unchanged.

Suppose every edge is traversed at least once in every window of m consecutive steps. Then A1 holds with:

$$\gamma = -\frac{\ln(1-\lambda)}{m} \quad \text{and} \quad \delta_0 = \frac{CD}{1-\lambda}$$

Proof. Each traversal of an edge multiplies its geodesic distance to the target by exactly $(1-\lambda)$. By time t , an edge has been traversed at least $\lfloor t/m \rfloor$ times, so:

$$d_{\text{geo}}(U_t(e), U_\infty(e)) \leq D(1-\lambda)^{\lfloor t/m \rfloor} \leq D(1-\lambda)^{(t/m-1)} = \left(\frac{D}{1-\lambda}\right) e^{-\gamma t}$$

Multiplying by C gives A1. ■

This is a modeling choice, a learning-toward-target dynamic, not a derived fact about minds. On shared transitions the two agents share the target; this is what "sharing a path" means in the model.

The Theorem

The **Asymptotic Qualia Theorem** says if two agents share a path, the gap between two agents' worldviews converges to a floor ϵ_0 , and the floor is almost surely above zero. Under assumptions A1–A3, let X_t, Y_t be the agents' connections at time t , and let $\epsilon_0 = d(X_\infty, Y_\infty)$. Then:

1. **Convergence:** For all $t \geq 0$: $|d(X_t, Y_t) - \epsilon_0| \leq 2L\delta_0 e^{-\gamma t}$
2. **Positive Floor:** $\epsilon_0 > 0$ almost surely.
3. **Consequence:** $d(X_t, Y_t) \geq \epsilon_0/2$ for all $t \geq t^*$, where: $t^* = \max\left(0, \frac{1}{\gamma} \ln\left(\frac{4L\delta_0}{\epsilon_0}\right)\right)$

Proof

Part 1: Convergence

Step 1: Edges to Holonomies. For unitary matrices, $\|A_1 \cdots A_n - B_1 \cdots B_n\| \leq \sum_{i=1}^n \|A_i - B_i\|$ (telescoping, using invariance of the norm). Also $\|A^{-1} - B^{-1}\| = \|A^{-1}(B - A)B^{-1}\| = \|A - B\|$, so inverted edges contribute the same error. Each holonomy word has at most L factors, so by A1:

$$\|H_k^a(t) - H_k^a(\infty)\| \leq L\delta_0 e^{-\gamma t}$$

for each agent $a \in \{A, B\}$ and each loop k .

Step 2: Holonomies to Distance. Fix Q and k . Conjugation by Q is an isometry, so $\|H_k^B(t) - QH_k^A(t)Q^{-1}\|$ differs from $\|H_k^B(\infty) - QH_k^A(\infty)Q^{-1}\|$ by at most the sum of the two holonomy errors, which is at most $2L\delta_0 e^{-\gamma t}$. This bound is uniform in Q and k , so it survives taking the maximum over k and the minimum over Q . Hence:

$$|d(X_t, Y_t) - \epsilon_0| \leq 2L\delta_0 e^{-\gamma t}$$

Remark: The earlier discrete version stated a bound with an extra factor $1/(1 - e^{-\gamma})$. It is not needed here, because at time t every edge is already within $\delta_0 e^{-\gamma t}$ of its limit.

Part 2: The Floor is Almost Surely Positive

Step 1: Shared cycles restrict the gauge.

The size (dimension) of the joint centralizer subgroup K determines how much “freedom” the agents have left to align their remaining unshared experiences I_u . More shared experiences I_s mean a small K . Having shared edges pulls the limit ϵ_0 down toward zero, while unshared edges keep it strictly above zero ($\epsilon_0 > 0$)

Let:

$$K = \{Q \in G \mid Qh_kQ^{-1} = h_k \text{ for all } k \in I_s\}$$

be the joint centralizer of the shared limit holonomies. K is a compact subgroup. If X_∞ and Y_∞ were gauge equivalent, some $Q \in K$ would satisfy $QH_k^A(\infty)Q^{-1} = H_k^B(\infty)$ for all $k \in I_u$.

Step 2: The equivalence map. Define $\Phi_K: K \times G^{r_u} \rightarrow G^{r_u} \times G^{r_u}$ by:

$$\Phi_K(Q, h_1, \dots, h_{r_u}) = (h_1, \dots, h_{r_u}, Qh_1Q^{-1}, \dots, Qh_{r_u}Q^{-1})$$

The set of exactly alignable limit pairs is the image of Φ_K , which is compact and so measurable.

Step 3: Rank. The differential of Φ_K at (Q, h) vanishes exactly on directions $(\xi, 0)$ with ξ in the Lie algebra of K commuting with every h_i . Writing $s_K(h)$ for the dimension of the joint centralizer of h inside K , the rank is $\dim K + r_u \dim G - s_K(h)$. The target dimension is $2r_u \dim G$, so the dimension gap is:

$$\Delta = r_u \dim G - \dim K + s_K(h)$$

Step 4: The gap is at least 1.

- **Case A (K a proper subgroup of G):** A closed subgroup of the same dimension as G would contain the identity component of G , which is all of G because G is connected. So $\dim K \leq \dim G - 1$, and:

$$\Delta \geq r_u \dim G - \dim K > \dim G - \dim K \geq 1$$

- **Case B ($K = G$, shared holonomies are central):** If $r_u \geq 2$, then $\Delta \geq (r_u - 1) \dim G \geq 1$. If $r_u = 1$, every element of a compact connected group lies in a maximal torus whose Lie algebra commutes with it, so $s_K(h) \geq \text{rank}(G) \geq 1$ and $\Delta \geq 1$.

Step 5: Sard's Theorem. In every case the rank is below the target dimension at every point of the domain, so every point is critical and, by Sard's theorem, the image of Φ_K has Haar measure zero. By A3, conditional on h_s the pair of unshared limit holonomies lands in this null set with probability zero. Integrating over h_s :

$$P(d(X_\infty, Y_\infty) = 0) = 0$$

Since the minimum defining d is attained, $\epsilon_0 > 0$ almost surely. ■

Part 3: Consequence

From Part 1, $d(X_t, Y_t) \geq \epsilon_0 - 2L\delta_0 e^{-\gamma t}$.

This is at least $\epsilon_0/2$ as soon as $e^{-\gamma t} \leq \epsilon_0/(4L\delta_0)$, that is for: $t \geq \frac{1}{\gamma} \ln\left(\frac{4L\delta_0}{\epsilon_0}\right)$

Numerical Sanity Check ($G = \text{SO}(3)$)

A simulation with 6 vertices and 9 edges (a shared spanning tree, one shared non-tree edge, three unshared non-tree edges; so $r = 4$, $r_u = 3$, $L = 8$) used the update rule of Lemma 0 with $\lambda = 0.2$, edges traversed in round-robin order ($m = 9$), and independent random targets on unshared edges. The distance d was computed by multi-start numerical minimization over Q .

t	$d(X_t, Y_t)$	$\ d - \epsilon_0\ $	Bound $2L\delta_0 e^{-\gamma t}$
0	1.996	0.068	88.86
100	1.861	0.067	7.45
200	2.005	0.076	0.62
300	1.928	0.001	0.05
400	1.929	0.000	0.00

Here $\epsilon_0 = 1.929$. The bound held at every time point checked ($t = 0, 50, \dots, 400$). Over 60 independent draws of the targets, ϵ_0 ranged from 1.34 to 2.71 (median 2.03), always clearly above zero.

Control: When all targets were shared (digital twins), ϵ_0 was zero to numerical precision (at most 9×10^{-13} over 10 draws).

Two honest observations: The bound is loose, since actual convergence is much faster than the guaranteed rate. Also, $d(X_t, Y_t)$ does not shrink from a large value: with random initial states it already starts near the floor, illustrating convergence to a floor rather than monotone contraction. The floor is large here because three of the four independent cycles are fully unshared; with more shared structure it would be smaller.

So what does it tell us

What can we reply to sceptics who still argue that all this does not solve the hard problem and that's not the hard problem they were asking about? There are two aspects of the explanatory gap. One is explaining the difference between 1st and 3rd person views (why does the 1st person view feeling like something at all), and the other one is finding an exact translation between two first-person perspectives.

The second question is the one that matters pragmatically in our lives. If we cannot understand subjective 1st-person experience in an objective way by measuring it, can we at least try to understand it by transforming a 1-person experience from one person to another? The answer to this is - as the theorems show - that exact equivalence of subjective experience is unattainable, and shared experience gets you only to a floor that is almost surely above zero. But even this minimal understanding gives us access to the larger question. As members of the same biological species, we know what it is like to be a human being that experiences pleasure and pain and all the other basic emotions. Each of us knows what it is like to feel good or bad, happy or sad. The longer we see what the other person went through, the more we understand what it must be like to be that person.

The answer to the first question (what it is like to feel something at all and why structure is accompanied by experience in the first place) can therefore be answered by everybody personally. This answer is in my opinion ultimately anchored in our biological nature. Our minds are inextricably linked to a single, unique, biological body. Our genes have created biomolecular control mechanisms - emotions - to control us. They tell us in every moment if something is good or bad, if something should be approached or not, if something should be consumed or not. These are in my opinion the irreducible dimensions of subjective experience.

They are linked to the needs of our biological body - mainly survival and reproduction. Whenever something is happening in nature in the biological world it is either supper or pairing time. The most basic emotions hunger and disgust (consume it-expel it), pleasure and pain (do it-stop it), love and fear (approach it-get away from it) are direct diametrical opposites of primary subjective dimensions.

Dimension	Biological Meaning	Operational Meaning
Pleasure vs. Pain	Tissue Maintenance vs. Tissue Damage	"Engage with this / Yes, do this" vs. "Protect yourself / No, stop this"
Hunger vs. Disgust	Energy Intake vs. Toxicity Avoidance	"Incorporate into Self" vs. "Expel from Self"
Love/Attraction vs. Fear	Reproduction vs. Threat Evasion	"Approach / Merge" vs. "Flee / Run away"

Thus, if we try to find out what it is like to feel something we eventually will hear behind the whole bio-molecular machinery of our emotions the whisper of our primary biological directive that is encoded in our genes. They are the ultimate judges if something should feel good or bad. The message of pleasure which our genes whisper is "This is going well. Continue, engage, consolidate this. Do this. You are in alignment with your primary directive".

Pleasure and pain are the most fundamental irreducible dimensions of subjective experience. They feel like they do because we have the bodies we have, which are built for survival and experience emotional valuations to guide our behavior. If we mean by the 'Hard Problem' explaining why a biological body built for survival experiences valence (pleasure/pain) as a directional force, then the answer is embodied homeostatic necessity — the irreducible dimensions of qualia are the physical orientation vectors of self-preservation.