Reality as a Consensus Protocol:
The Fixed-Point Computation That Implements Physics

B. Müller Kale Arnav Anirudha

April 7, 2026

Paper release: `r1378` **Released:** April 7, 2026

Abstract

This paper gives the consensus-theoretic core of Observer-Patch Holography (OPH). If physics is built from local observer descriptions that must agree on overlaps, the first question is whether different repair orders lead to different worlds. The main result is that, under termination and local confluence, the repair dynamics has a unique schedule-independent normal form. The paper also identifies the loop obstructions to global consistency, shows why physical uniqueness belongs on the gauge quotient rather than on raw microscopic representatives, and gives a fixed-point account of stable classical records. This is the finite patch-net theorem package that supports the broader relativity, gauge, particle, and observer branches in the OPH suite.

1 Introduction

This paper gives a finite patch-net formulation of the OPH consensus picture. A finite graph of observer patches carries local state spaces and overlap maps. Local repair maps act asynchronously, each modifying only one patch or a bounded neighborhood, and global consistency means agreement on every overlap. The central mathematical problem is whether this repair dynamics admits a unique normal form and how global obstructions to reconciliation are encoded.

The resulting theorem package has two layers. The first layer concerns convergence: under termination and local confluence, every initial configuration has a unique schedule-independent normal form. The second layer concerns obstructions: pairwise overlap agreement does not ensure a global solution, and the obstruction is holonomic. On the abelian branch it is the cycle sum of edge data; on the genuinely noncentral branch it is a crossed-module Čech class.

Gauge symmetry enters as invariance under changes of hidden local representation that preserve overlap data. When the repair step is read only on that overlap-invariant quotient, the normal-form map descends to the gauge quotient, so physical uniqueness is a quotient statement. Classical records are described by a finite closure system whose least common fixed point is independent of schedule. These results form the patch-net formulation used in the broader OPH literature.

This paper proves four core results:

1. **Asynchronous confluence.** Under termination and local confluence, every initial state has a unique normal form, independent of update schedule (Theorem 3.5).

2. **Cycle obstruction.** For affine overlap constraints over an abelian group, global consistency holds if and only if the holonomy vanishes on every cycle (Theorem 4.1). The parity triangle

gives the minimal frustrated example, and Theorem 4.4 extends the same logic to the crossed-module higher-gauge defect hierarchy used later in OPH.

3. **Gauge quotient.** When local repair is induced on the overlap-invariant quotient, the normal-form map descends to that quotient. Gauge-invariant observables are strictly unique there, even though microscopic representatives need not be unique (Theorem 5.2).

4. **Record-layer convergence.** Commuting closure operators on a finite lattice converge to the least common fixed point under any asynchronous schedule that applies each operator at least once (Theorem 6.2).

We also define a fitness functional over the space of candidate reconciliation laws and prove that replicator dynamics monotonically increases mean fitness (Theorem 7.4). This gives a clean mathematical model for the selection of physical law.

The results here are exact theorems about a computational model. The further claim that physical reality instantiates this model is the interpretive hypothesis of Observer-Patch Holography (OPH). The companion OPH manuscript develops the broader physics branches conditionally and keeps structural theorems, scaling-limit branches, calibration-sector outputs, and phenomenological continuations distinct [1].

2 Patch Nets, Overlaps, and Global Consistency

Definition 2.1 (Patch net)**.** *Let $G = (V, E)$ be a finite connected graph. Each vertex $i \in V$ is an* ***observer patch*** *with finite local state space S_i. The global state space is*

$$\Sigma := \prod_{i \in V} S_i.$$

For each edge $e = \{i, j\} \in E$, let I_e be an interface alphabet and let

$$\pi_{i,e} : S_i \to I_e, \qquad \pi_{j,e} : S_j \to I_e$$

be the interface projection maps. A global state $s = (s_i)_{i \in V} \in \Sigma$ is ***consistent on edge*** *$e = \{i, j\}$ iff*

$$\pi_{i,e}(s_i) = \pi_{j,e}(s_j).$$

The global consistency set is

$$C := \{s \in \Sigma : \forall\, e = \{i, j\} \in E,\ \pi_{i,e}(s_i) = \pi_{j,e}(s_j)\}.$$

For exposition we use a finite pairwise-overlap graph. The hypergraph version is straightforward: replace edges by hyperedges and pairwise equality by a common interface label on each hyperedge. Nothing in the proofs depends on the pairwise restriction.

The picture: each observer holds a local state, and neighboring observers share an interface through which they can compare notes. A universe-state is physically admissible exactly when all neighbors agree on their shared data. This is a constraint satisfaction problem (CSP), and the consistent states are the codewords.

Definition 2.2 (Inconsistency potential)**.** *For each edge e, choose a weight $w_e > 0$ and a function $d_e : I_e \times I_e \to \mathbb{R}_{\geq 0}$ with $d_e(a, b) = 0 \iff a = b$. Define*

$$\Phi(s) := \sum_{e=\{i,j\} \in E} w_e\, d_e\big(\pi_{i,e}(s_i), \pi_{j,e}(s_j)\big).$$

Then $s \in C \iff \Phi(s) = 0$.

So Φ is the total disagreement energy of the universe. Consistent states have zero energy. Everything else is frustrated.

3 Asynchronous Reconciliation and the Main Fixed-Point Theorem

Definition 3.1 (Local repair law)**.** *A **law** λ is a family of local repair maps*

$$T_i^\lambda : \Sigma \to \Sigma \qquad (i \in V)$$

*such that T_i^λ changes only the state of patch i (or, more generally, only a bounded neighborhood of i). Write $s \to_i t$ iff $t = T_i^\lambda(s) \neq s$. Let $\to := \bigcup_{i \in V} \to_i$, and let $\to^*$ be its reflexive-transitive closure. A state $s \in \Sigma$ is a **normal form** iff $T_i^\lambda(s) = s$ for all $i \in V$.*

The three assumptions that make the machine work:

Assumption 3.2 (Repair completeness)**.** $s \in C \iff \forall\, i \in V,\ T_i^\lambda(s) = s$.

Normal forms are exactly the globally consistent states. The dynamics is neither too weak (missing some inconsistencies) nor too strong (repairing things that were fine).

Assumption 3.3 (Termination)**.** *Every enabled repair strictly decreases the inconsistency potential: $s \to t \implies \Phi(t) < \Phi(s)$.*

Since Σ is finite, this guarantees termination.

Assumption 3.4 (Local confluence)**.** *The repair relation is locally confluent: if $s \to t$ and $s \to u$, then there exists $v \in \Sigma$ with $t \to^* v$ and $u \to^* v$.*

This is the diamond condition for the repair relation.

Theorem 3.5 (Asynchronous confluence / fixed-point law)**.** *Under Assumptions 3.2–3.4, every initial state $s \in \Sigma$ has a unique normal form*

$$\mathrm{nf}_\lambda(s) \in C,$$

and every maximal asynchronous repair execution from s terminates at that same state. The terminal state is independent of update order.

Proof. By Assumption 3.3, every rewrite step strictly decreases Φ, so no infinite rewrite sequence exists. Thus $\to$ is terminating.

We prove confluence by induction on $\Phi(s)$. Take any $s \in \Sigma$ and suppose $s \to^* t$ and $s \to^* u$. We must show t and u have a common descendant.

If one reduction has length 0, there is nothing to prove. Otherwise write $s \to s_1 \to^* t$ and $s \to s_2 \to^* u$. By local confluence, there exists w with $s_1 \to^* w$ and $s_2 \to^* w$. Since $s \to s_1$ and $s \to s_2$, Assumption 3.3 gives $\Phi(s_1) < \Phi(s)$ and $\Phi(s_2) < \Phi(s)$. The induction hypothesis applies at s_1 and s_2. Hence t and w have a common descendant, and u and w have a common descendant. Therefore t and u have a common descendant.

A terminating confluent rewrite system has a unique normal form reachable from each initial state. By Assumption 3.2, the normal forms are exactly C. Every maximal execution terminates (by Φ-decrease) and reaches $\mathrm{nf}_\lambda(s)$ (by uniqueness). $\square$

Corollary 3.6 (Objective law is schedule-independent). *Let $M : \Sigma \to Y$ be any observable. Under the hypotheses of Theorem 3.5, $M(\mathrm{nf}_\lambda(s))$ is independent of the asynchronous update schedule. If physical law is identified with the map $s \mapsto M(\mathrm{nf}_\lambda(s))$, then physical law is objective.*

Proof. All schedules from the same initial s terminate at $\mathrm{nf}_\lambda(s)$, so all yield the same M-value. □

The theorem identifies objectivity with schedule-independent convergence of the repair dynamics.

4 Why Local Agreement Is Not Enough: Cycle Holonomy and Frustration

Now we show that the story has a twist: just because every pair of neighbors agrees does not mean the whole system is consistent.

Theorem 4.1 (Cycle-obstruction / holonomy criterion). *Let A be an abelian group, and let $G = (V, E)$ be a connected graph with an arbitrary orientation on each edge. For each oriented edge $e : u \to v$, assign a label $b_e \in A$. Consider the affine consistency equations*

$$x_v - x_u = b_e \qquad \text{for every oriented edge } e : u \to v,$$

where $x_v \in A$ are unknown patch labels. A global solution $x : V \to A$ exists if and only if for every cycle $C \subseteq G$,

$$\sum_{e \in C} \varepsilon_C(e)\, b_e = 0,$$

where $\varepsilon_C(e) = +1$ if the cycle traverses e in the chosen orientation and -1 otherwise.

Proof. **Necessity.** Suppose x is a solution. Summing the edge equations around any cycle C,

$$\sum_{e:u \to v \in C} \varepsilon_C(e)\,(x_v - x_u) = \sum_{e \in C} \varepsilon_C(e)\, b_e.$$

The left side telescopes to 0 because every vertex appears once with + sign and once with − sign.

Sufficiency. Fix a root $r \in V$. For any vertex v, choose a path $P_{r \to v}$ and define

$$x_v := \sum_{e \in P_{r \to v}} \varepsilon_{P_{r \to v}}(e)\, b_e, \qquad x_r := 0.$$

If P and P' are two paths from r to v, traversing P followed by the reverse of P' yields a cycle. By the vanishing-holonomy assumption, the total signed sum is zero, so x_v is well-defined. For any edge $e : u \to v$, extending a path to u by that edge gives $x_v = x_u + b_e$. □

Corollary 4.2 (Parity triangle: pairwise consistency is not enough). *Take $A = \mathbb{Z}_2$ on the triangle A-B-C-A with edge labels $b_{AB} = 0$, $b_{BC} = 0$, $b_{CA} = 1$. Each individual edge equation is satisfiable. But the global system is not: the cycle sum is $0 \oplus 0 \oplus 1 = 1 \neq 0$.*

This example shows that pairwise consistency does not imply a global solution. The obstruction is carried by the cycle.

Corollary 4.3 (Stable defects as frustrated holonomy)**.** *Define the defect energy*

$$\Phi_b(x) := \sum_{e:u\to v\in E} w_e\, \mathbf{1}[x_v - x_u \neq b_e]\,.$$

If the cycle-holonomy condition fails, then $\min_{x:V\to A} \Phi_b(x) > 0$*. Every minimizer contains irreducible residual inconsistency.*

Proof. If $\min_x \Phi_b(x) = 0$, some assignment satisfies all edge equations, contradicting Theorem 4.1. ☐

Residual inconsistencies of this type cannot be removed by local repair moves. In the OPH interpretation they are stable topological defects of the reconciliation dynamics.

Theorem 4.4 (Higher-gauge defect hierarchy)**.** *Let a finite overlap nerve carry crossed-module defect data* (g_{ij}, h_{ijk}) *for a compact crossed module*

$$H \xrightarrow{\partial} G.$$

Under local rechartings by

$$C^1(N,H) \rtimes C^0(N,G),$$

the nonabelian Čech class

$$q = [(g,h)] \in \check{H}^2(N, H \to G)$$

is invariant. Strict global reconciliation exists if and only if $q = 0$*, and nonzero* q *labels stable fixed-cutoff higher-gauge defects.*

Proof. The allowed rechartings are exactly the crossed-module coboundaries. So they preserve the class and strictify the weak gluing data precisely when the class vanishes. A nonzero class therefore gives a topologically protected residual obstruction, in the full crossed-module defect hierarchy rather than only in the abelian truncation. ☐

5 Gauge Symmetry as Implementation Hiding

Definition 5.1 (Gauge action)**.** *Let* $\Gamma = \prod_{i\in V} \Gamma_i$ *act on* Σ *componentwise. The action is a* ***gauge action*** *if for every* $e = \{i,j\} \in E$,

$$\pi_{i,e}(\gamma_i \cdot x) = \pi_{i,e}(x) \qquad \forall\, x \in S_i,\ \forall\, \gamma_i \in \Gamma_i.$$

Gauge changes alter hidden local representations but do not alter overlap data.

Gauge transformations change hidden local representations while leaving overlap data fixed. Write

$$q : \Sigma \to \Sigma/\Gamma, \qquad q(s) = [s],$$

for the gauge-orbit map. A **physical repair law** is a family of local maps

$$\overline{T}_i^{\lambda} : \Sigma/\Gamma \to \Sigma/\Gamma$$

on the overlap-invariant quotient. A **representative repair family** is any choice of lifts

$$T_i^{\lambda} : \Sigma \to \Sigma$$

such that

$$q \circ T_i^\lambda = \overline{T}_i^\lambda \circ q.$$

This is the finite patch-net form of saying that the repair step is defined first on overlap-invariant physical data and only then lifted to hidden representatives. The separate OPH problem of deriving the physical repair law itself from recovery dynamics remains upstream; the point here is only that no extra gauge-covariance axiom is needed once repair is formulated on the quotient.

Theorem 5.2 (Gauge quotient theorem)**.** *Under the gauge action of Definition 5.1, any representative lift of a physical repair law as just defined, and the hypotheses of Theorem 3.5,*

$$q\big(\mathrm{nf}_\lambda(\gamma \cdot s)\big) = q\big(\mathrm{nf}_\lambda(s)\big) \qquad \forall\, \gamma \in \Gamma,\ \forall\, s \in \Sigma.$$

Hence the normal-form map descends to the quotient:

$$\overline{\mathrm{nf}}_\lambda : \Sigma/\Gamma \to \Sigma/\Gamma, \qquad [s] \mapsto [\mathrm{nf}_\lambda(s)].$$

Proof. Suppose $s \to_i t$, so $t = T_i^\lambda(s) \neq s$. If $s' = \gamma \cdot s$, then

$$q(T_i^\lambda(s')) = \overline{T}_i^\lambda(q(s')) = \overline{T}_i^\lambda(q(s)) = q(T_i^\lambda(s)).$$

Thus gauge-equivalent inputs induce the same repaired orbit, and by induction the orbit reached after any repair sequence depends only on the initial orbit. Every maximal repair sequence from s ends at $\mathrm{nf}_\lambda(s)$ by Theorem 3.5, so the terminal orbit $q(\mathrm{nf}_\lambda(s))$ depends only on $q(s) = [s]$. This makes

$$[s] \longmapsto [\mathrm{nf}_\lambda(s)]$$

well-defined on Σ/Γ. □

Corollary 5.3 (Gauge-invariant law)**.** *If $M : \Sigma \to Y$ is gauge-invariant ($M(\gamma \cdot s) = M(s)$ for all γ), then $M(\mathrm{nf}_\lambda(s))$ depends only on the gauge orbit $[s]$, not the representative.*

Proof. Because M is gauge-invariant, it factors through the orbit map: $M = \overline{M} \circ q$ for some $\overline{M} : \Sigma/\Gamma \to Y$. Theorem 5.2 gives

$$q\big(\mathrm{nf}_\lambda(\gamma \cdot s)\big) = q\big(\mathrm{nf}_\lambda(s)\big),$$

hence

$$M\big(\mathrm{nf}_\lambda(\gamma \cdot s)\big) = \overline{M}\big(q(\mathrm{nf}_\lambda(\gamma \cdot s))\big) = \overline{M}\big(q(\mathrm{nf}_\lambda(s))\big) = M\big(\mathrm{nf}_\lambda(s)\big).$$

□

Corollary 5.4 (Inert ancillary refinement does not change physical law)**.** *Let $K = \prod_i K_i$ be a finite ancillary state space and define $\Sigma^\eta := \Sigma \times K$. Lift the repair maps by*

$$T_i^{\lambda,\eta}(s,k) := (T_i^\lambda(s), k).$$

If $M^\eta : \Sigma^\eta \to Y$ ignores the ancillary factor,

$$M^\eta(s,k) = M(s),$$

with M gauge-invariant on Σ, then

$$M^\eta(\mathrm{nf}_\lambda^\eta(s,k)) = M(\mathrm{nf}_\lambda(s))$$

for all $(s,k) \in \Sigma^\eta$.

Proof. Because the ancillary factor is inert,

$$\mathrm{nf}^{\eta}_{\lambda}(s,k) = (\mathrm{nf}_{\lambda}(s), k).$$

Therefore $M^{\eta}(\mathrm{nf}^{\eta}_{\lambda}(s,k)) = M(\mathrm{nf}_{\lambda}(s))$, and Corollary 5.3 supplies gauge-orbit independence. □

Physical uniqueness therefore holds on the quotient by gauge or implementation hiding. The same statement remains unchanged under inert ancillary stabilization.

6 Classical Records as an Eventually Consistent Layer

Definition 6.1 (Record lattice)**.** *Let L be a finite distributive lattice. A map $F : L \to L$ is a **closure operator** if it is inflationary ($x \leq F(x)$), monotone ($x \leq y \implies F(x) \leq F(y)$), and idempotent ($F(F(x)) = F(x)$). Let $F_1, \ldots, F_n$ be pairwise commuting closure operators on L.*

Theorem 6.2 (Asynchronous convergence of the record layer)**.** *For any initial record state $x_0 \in L$, define $x^\star := F_1F_2 \cdots F_n(x_0)$. Then:*

1. *$x^\star$ is independent of the order of composition.*
2. *$x^\star$ is a common fixed point of all F_i.*
3. *$x^\star$ is the least common fixed point above x_0.*
4. *Any asynchronous application schedule in which each F_i is applied at least once converges in finitely many steps to $x^\star$.*

Proof. Commutativity makes the composition order-independent, so $x^\star$ is well-defined.

For any i, commutativity and idempotence give

$$F_i(x^\star) = F_iF_1 \cdots F_n(x_0) = F_1 \cdots F_i^2 \cdots F_n(x_0) = x^\star.$$

So $x^\star$ is a common fixed point.

If $y \geq x_0$ satisfies $F_i(y) = y$ for all i, monotonicity gives $x^\star = F_1 \cdots F_n(x_0) \leq F_1 \cdots F_n(y) = y$. So $x^\star$ is the least common fixed point above x_0.

Under any asynchronous schedule in which each operator is applied at least once, the total effect after the first full pass is $F_1 \cdots F_n$ regardless of order and repetitions (by commutativity and idempotence). The state reaches $x^\star$ at that stage and stays there. □

The record layer is therefore a finite closure system: classical records are the order-independent least common fixed point generated from the initial record state.

7 Law-Space Selection and Observer Emergence

This section studies a simple meta-selection model on law space. The aim is to formalize one criterion for favoring schedule-robust, observer-supporting, and simple laws; the replicator dynamics below is part of the model, not a claim about literal cosmological dynamics.

We begin by defining what it means for a law to support observers. An observer is treated operationally as a **persistent predictive module**: a subgraph that maintains stable records and uses them to predict its boundary's future behavior.

Definition 7.1 (Schedule robustness)**.** *Fix distributions μ over initial conditions and ν over asynchronous schedules, and a gauge-invariant observable M. For a law λ, define*

$$\mathcal{R}_M(\lambda) := \Pr_{s\sim\mu,\,\sigma,\tau\sim\nu}\big[M\big(\mathrm{nf}^\sigma_\lambda(s)\big) = M\big(\mathrm{nf}^\tau_\lambda(s)\big)\big]\,.$$

If Theorem 3.5 holds for λ, then $\mathcal{R}_M(\lambda) = 1$.

Definition 7.2 (Observer yield)**.** *Let X^λ_t denote the stationary process obtained by repeated local perturbation plus reconciliation under law λ. A subgraph $U \subseteq V$ with record map r_U is (η, ε, h)-observer-like if it is record-stable:*

$$\Pr\big[r_U(X^\lambda_{t+h}) = r_U(X^\lambda_t)\big] \geq 1 - \eta,$$

and predictive:

$$I\big(r_U(X^\lambda_t);\ X^\lambda_{\partial U,\,t+1:t+h}\big) \geq \varepsilon.$$

Define

$$\mathcal{O}_{\eta,\varepsilon,h}(\lambda) := \mathbb{E}[\#\,\{U \subseteq V : U \text{ is } (\eta,\varepsilon,h)\text{-observer-like}\}]\,.$$

Definition 7.3 (Law fitness)**.** *Let $K(\lambda)$ be a description-length penalty. Define*

$$f(\lambda) = \alpha\,\mathcal{R}_M(\lambda) + \beta\,\mathcal{O}_{\eta,\varepsilon,h}(\lambda) - \gamma\,K(\lambda),$$

with $\alpha, \beta, \gamma > 0$.

Theorem 7.4 (Replicator monotonicity on law space)**.** *Let $\Lambda = \{\lambda_1, \ldots, \lambda_m\}$ be candidate laws with population weights $x_i(t)$ under replicator dynamics:*

$$\dot{x}_i = x_i(f_i - \bar{f}), \qquad f_i := f(\lambda_i), \qquad \bar{f} := \sum_{j=1}^m x_j f_j.$$

Then

$$\frac{d}{dt}\bar{f} = \mathrm{Var}_x(f) \geq 0.$$

Mean fitness is nondecreasing, and strictly increasing unless all extant laws have the same fitness.

Proof. Direct computation:

$$\frac{d}{dt}\bar{f} = \sum_i \dot{x}_i f_i = \sum_i x_i(f_i - \bar{f}) f_i = \sum_i x_i f_i^2 - \bar{f}^2 = \mathrm{Var}_x(f) \geq 0.$$

Equality iff all f_i on the support of x are equal. □

This theorem records the monotonicity property of the meta-selection model.

8 Connection to Observer-Patch Holography

The formalism above is the computational skeleton of Observer-Patch Holography (OPH). In OPH, the observer patches carry von Neumann algebras on a holographic screen S^2, the overlap projections are restrictions to shared subalgebras, and the consistency condition is algebraic state agreement on overlaps.

The bridge to physics works as follows. In the full OPH framework:

- The patch net becomes a net of subregion algebras on S^2.
- Overlap Consistency, one of the canonical OPH axioms, is the algebraic version of Definition 2.1.
- The recoverability clause of the canonical Recoverable Generalized Entropy axiom provides the collar factorization $\rho_{ABD} = \bigoplus_\alpha p_\alpha \, \rho^{(\alpha)}_{Ab^\alpha_L} \otimes \rho^{(\alpha)}_{b^\alpha_R D}$ that makes observer checkpointing and restoration well-defined (see Appendix A).
- Gauge symmetry as implementation hiding (Theorem 5.2) becomes the edge-sector fusion structure from which OPH conditionally reconstructs a compact gauge group and, under the MAR admissibility package, selects the Standard Model gauge group $\mathrm{SU}(3)\times\mathrm{SU}(2)\times\mathrm{U}(1)/\mathbb{Z}_6$.
- Stable defects (Corollary 4.3) become the topologically protected excitations that OPH identifies with particles.
- The record-layer theorem (Theorem 6.2) provides the formal basis for the classical observation layer in OPH, where records are implemented by approximately commuting projectors in overlap centers.

The OPH manuscript develops a conditional gravity branch from entanglement equilibrium and modular geometry, the SM gauge-group closure from edge-sector admissibility plus MAR, and a worldsheet/string reorganization from large-N heat-kernel asymptotics. Input-dependent cosmological statements and later phenomenological continuations are kept explicitly separate from that recovered-core claim set.

This paper provides the finite patch-net foundation for that OPH interpretation. When the broader suite is described using labels such as $D3$–$D5$ or $D7$–$D9$, those are internal documentary node labels from the companion OPH derivation ledger [1]; they are not external references, and that companion ledger is the place to look up the definitions.

9 Discussion and Open Problems

This paper proves the fixed-point consensus spine of OPH: schedule-independent normal forms, holonomy obstructions, gauge-quotient invariance, and record-layer convergence. It also gives a clean law-selection meta-model. The conditional relativity chain and the realized Standard Model structural chain remain the recovered core. The capacity relation is separate and input-dependent. Downstream phenomenology requires additional premises beyond the consensus results proved here.

Six directions stand out as immediate next steps:

1. **Computational complexity.** What is the complexity of computing $\mathrm{nf}_\lambda(s)$? The termination bound from Φ-decrease is a worst-case measure, but typical instances may converge much faster. Is there a polynomial bound for natural classes of patch nets?
2. **Coarse-graining.** When does reconciliation commute with renormalization? If you coarse-grain the patch net and then reconcile, do you get the same result as reconciling first and then coarse-graining? The conditions under which these operations commute would connect the discrete model to continuum field theory.

3. **Defect classification and refinement-limit transportability.** The fixed-cutoff hierarchy extends from abelian frustrations to crossed-module classes $q \in \check{H}^2(N, H \to G)$. The remaining task is to connect those higher-gauge defect sectors to the refinement-stable transportable sector category used in the broader compact-gauge reconstruction lane.

4. **Observable-level confluence.** When the raw state is not unique because local confluence fails, can gauge-invariant macrostates still be unique? This weaker form of objectivity may hold under broader conditions than the strict normal-form theorem.

5. **Quantum lift.** How much of the patch-net theorem survives when local state spaces become operator algebras and repair maps become quantum channels? The Markov-collar splice theorem in the appendix suggests that the algebraic version is natural.

6. **Scaling-limit bridge to full OPH.** Under what additional regularity conditions does the consensus protocol emit the prime geometric cap pair and support ordered cut-pair rigidity on that extracted limit, so that the D3–D5 gravity chain and the D7–D9 gauge/matter chain become available? This is also the point at which the broader Phase-II and Phase-III branches must remain disciplined: dark-sector, baryogenesis, spectroscopy, and string/worldsheet topics require extra premises beyond what this paper itself proves.

The last item is the bridge to the broader OPH branches. Failure of that bridge would leave the fixed-point theorems proved here unchanged while revising the downstream gravity, gauge, or continuation sectors.

9.1 Conditional BFT and QECC Extensions

The consensus formalism of OPH has natural analogies to classical and quantum distributed Byzantine agreement. Observer patches correspond to protocol nodes, overlap repair corresponds to a quorum vote, and the repair fixed-point corresponds to a consensus state. Under explicit structural assumptions (quorum size $\geq 2f+1$, partial synchrony, one-vote-per-view, certificate semantics, and DLS-style view-change), a QBFT-style interpretation of OPH repair satisfies safety and liveness (Appendix C, Theorem C.2). The repair map admits a quantum-channel formulation via the Petz recovery map (Definition C.3), but the CPTP property on all inputs requires either full-rank $\mathcal{N}(\sigma)$ or an explicit domain restriction; trace-preserving completion is not automatic when $\mathcal{N}(\sigma)$ has a non-trivial kernel (Proposition C.5). A quantum error-correcting interpretation is possible under additional topological structure, but the code-distance / min-cut equality requires explicit conditions on logical-operator homology and boundary geometry (Claim C.9). All of these extensions are conditional or conjectural and are not part of the core theorem package of Paper 4.

A Quantum/Algebraic Lift: Markov-Collar Splice Theorem

This appendix records the algebraic splice statement used to relate the finite patch-net model to the OPH collar formalism.

In OPH, an observer is written as

$$O = (P, \mathcal{A}(P), \rho, R),$$

where P is the screen patch, $\mathcal{A}(P)$ the local von Neumann algebra, ρ the local state, and R the record algebra.

Theorem A.1 (Markov-collar splice theorem, exact and controlled)**.** *Suppose a collar tripartition A-B-D has exact Markov decomposition*

$$\rho_{ABD} = \bigoplus_{\alpha} p_\alpha \, \rho^{(\alpha)}_{Ab_L^\alpha} \otimes \rho^{(\alpha)}_{b_R^\alpha D}.$$

Let $\sigma^{(\alpha)}_{b_R^\alpha D'}$ *be any family of normalized environment states compatible with the same right-boundary sectors. Define*

$$\rho'_{ABD'} = \bigoplus_{\alpha} p_\alpha \, \rho^{(\alpha)}_{Ab_L^\alpha} \otimes \sigma^{(\alpha)}_{b_R^\alpha D'}.$$

Then for every observable X *supported on* $A \cup b_L$*,*

$$\mathrm{Tr}(X\rho'_{ABD'}) = \mathrm{Tr}(X\rho_{ABD}).$$

Now fix one finite-dimensional collar model and let

$$\mathfrak{M}_{A:B:D} := \{\tau_{ABD} : I(A:D \mid B)_\tau = 0\},$$

with exact-Markov distance modulus

$$\delta^{\mathrm{M}}_{A:B:D}(\varepsilon) := \sup\left\{ \inf_{\tau \in \mathfrak{M}_{A:B:D}} \|\omega - \tau\|_1 : I(A:D \mid B)_\omega \le \varepsilon \right\}.$$

Then

$$\delta^{\mathrm{M}}_{A:B:D}(\varepsilon) \to 0 \qquad (\varepsilon \downarrow 0).$$

Hence if $I(A:D \mid B)_\omega \le \varepsilon$ *and* $\widetilde{\omega}_\varepsilon \in \mathfrak{M}_{A:B:D}$ *is chosen so that*

$$\|\omega - \widetilde{\omega}_\varepsilon\|_1 \le \delta^{\mathrm{M}}_{A:B:D}(\varepsilon),$$

the corresponding exact splice $\widetilde{\omega}'_\varepsilon$ *satisfies*

$$\left|\mathrm{Tr}(X\omega) - \mathrm{Tr}(X\widetilde{\omega}'_\varepsilon)\right| \le \|X\|_\infty \, \delta^{\mathrm{M}}_{A:B:D}(\varepsilon)$$

for every observable X *supported on* $A \cup b_L$*.*

Independently, if $I(A:D \mid B)_\omega \le \varepsilon$*, then there exists a recovery map* $\mathcal{R}_{B\to BD}$ *such that*

$$\|\omega_{ABD} - (\mathrm{id}_A \otimes \mathcal{R}_{B\to BD})(\omega_{AB})\|_1 \le 2\sqrt{1 - e^{-\varepsilon}} \le 2\sqrt{\varepsilon}.$$

Proof. The exact splice statement is the usual blockwise factorization argument:

$$\mathrm{Tr}(X\rho'_{ABD'}) = \sum_{\alpha} p_\alpha \, \mathrm{Tr}\Big(X \, \rho^{(\alpha)}_{Ab_L^\alpha}\Big) \, \mathrm{Tr}\Big(\sigma^{(\alpha)}_{b_R^\alpha D'}\Big).$$

Each right factor is normalized, so the value agrees with the same computation for ρ_{ABD}.

For the controlled statement, compactness of the fixed finite-dimensional state space and continuity of conditional mutual information imply $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon) \to 0$: otherwise one could find a sequence with $I(A:D \mid B) \to 0$ staying a fixed trace distance away from every exact Markov state, contradicting convergence of a subsequence to an exact Markov limit point. Once $\widetilde{\omega}_\varepsilon$ is chosen, the exact splice identity for $\widetilde{\omega}_\varepsilon$ gives

$$\left|\mathrm{Tr}(X\omega) - \mathrm{Tr}(X\widetilde{\omega}'_\varepsilon)\right| = \left|\mathrm{Tr}[X(\omega - \widetilde{\omega}_\varepsilon)]\right| \le \|X\|_\infty \, \delta^{\mathrm{M}}_{A:B:D}(\varepsilon).$$

The final inequality is the standard Fawzi–Renner recovery bound [14]. ☐

This appendix therefore uses exact splice identities in only two regimes: literal exact Markovity, or a controlled collar family on one fixed finite-dimensional model for which $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon) \to 0$. Small one-shot conditional mutual information is not silently upgraded to an exact normal form.

B Fixed-Cutoff Realization, Quotient Repair, and Edge Centers

This appendix rehouses the fixed-cutoff realization and edge-center items that had been living in the former standalone technical supplement. They belong here because they sharpen the quotient-first repair interpretation used throughout the consensus paper and make the collar boundary data explicit at the same finite patch-net level.

B.1 Quotient Repair and UV Underdetermination

At fixed cutoff, each regulator cell x carries a finite-dimensional factor $\mathfrak{h}_x$, patch algebras are finite type-I algebras, and gauge-as-gluing is realized as a compact boundary redundancy action on cut data. The physical repair law therefore belongs on the overlap-invariant quotient rather than on hidden representatives. If $q : \Sigma \to \Sigma/\Gamma$ is the quotient by boundary redundancy and $\overline{T}_i$ is the physical quotient update, a representative-level map T_i is only required to be a lift satisfying

$$q \circ T_i = \overline{T}_i \circ q.$$

Hence

$$q(T_i(\gamma \cdot s)) = q(T_i(s))$$

for gauge-equivalent inputs. Quotient descent is therefore structural, while strict representative-level covariance is only implementation bookkeeping.

Proposition B.1 (Ancilla-stable UV underdetermination)**.** *Let a fixed-cutoff OPH realization be stabilized by finite ancillary factors K_P in a fixed product state, with observable patch algebras embedded as $\mathcal{A}(P) \otimes \mathbf{1}_{K_P}$ and repair dynamics acting trivially on the ancillas. Then observable expectations on the physical subalgebras, overlap data, the local-Gibbs branch, the collar conditional mutual information $I(A : D \mid B)$, the Fawzi–Renner remainder, the collar Markov modulus, and the quotient normal form are unchanged. Thus OPH determines the UV branch only modulo such ancillary stabilization together with gauge or implementation hiding, not a unique microscopic presentation.*

Proof. Product ancillas leave physical observables unchanged, cancel additively inside conditional mutual information, and are inert under the repair maps. Hence every invariant listed above is unchanged. □

B.2 Derived Boundary Data and Ordinary EC

Proposition B.2 (Derived boundary gluing datum)**.** *Choose a finite regulator chart for the patches meeting along a connected cut Σ. Because the local overlap algebras are finite-dimensional matrix algebras, any overlap-consistent recharting is an inner automorphism and is implemented by a unitary on the cut Hilbert space. The compact closure of the subgroup generated by these recharting unitaries is a compact boundary redundancy group K_Σ. If triple-overlap defects are central, the projective composition law lifts to a compact central extension $\widehat{K}_\Sigma$; on the ordinary branch one simply sets $\widehat{K}_\Sigma = K_\Sigma$. A genuinely noncentral 2-group defect is the only obstruction to reducing the overlap transition system to an ordinary compact group action.*

Theorem B.3 (Derived EC decomposition)**.** *Under the fixed-cutoff regulator realization above, and on the ordinary or central-defect branch, the collar Hilbert space is*

$$\mathcal{H}_{B_\delta} = (\tilde{\mathcal{H}}_{B_L} \otimes \tilde{\mathcal{H}}_{B_R})^{\widehat{K}_\Sigma} \cong \bigoplus_\alpha \left(\mathcal{H}_{b_L^\alpha} \otimes \mathcal{H}_{b_R^\alpha} \right),$$

and the center of the collar algebra is generated by the block projectors:

$$Z(\mathcal{A}(B_\delta)) = \bigoplus_\alpha \mathbb{C} \cdot \mathbf{1}_\alpha.$$

The right half-collar carries the contragredient representation because it sees inverse transport across the same cut.

Remark B.4. This is the finite-patch-net origin of the collar center used by the later Markov, record, and observer packages. Exact Markovity is still an additional state hypothesis; EC provides the kinematic block structure.

B.3 Higher-Gauge Replacement on the Genuinely Noncentral Branch

Proposition B.5 (Derived higher-gauge cut datum)**.** *On the genuinely noncentral branch, weak overlap gluing on a connected cut Σ is encoded by a compact crossed module*

$$\mathbb{K}_\Sigma = (H_\Sigma \xrightarrow{\partial_\Sigma} G_\Sigma, \triangleright)$$

with defect class

$$q_\Sigma \in \check{H}^2(N_\Sigma, H_\Sigma \to G_\Sigma),$$

and compact higher-gauge change system

$$\mathcal{T}_\Sigma = C^1(N_\Sigma, H_\Sigma) \rtimes C^0(N_\Sigma, G_\Sigma).$$

Theorem B.6 (Higher-gauge EC decomposition and defect transport)**.** *On the genuinely noncentral branch,*

$$\mathcal{H}^{2g}_{B_\delta} = (\tilde{\mathcal{H}}_{B_L} \otimes \tilde{\mathcal{H}}_{B_R})^{\mathcal{T}_\Sigma} \cong \bigoplus_\lambda (\mathcal{H}_{b^\lambda_L} \otimes \mathcal{H}_{b^\lambda_R}),$$

and

$$Z(\mathcal{A}_{2g}(B_\delta)) = \bigoplus_\lambda \mathbb{C} \cdot \mathbf{1}_\lambda.$$

Moreover the higher-gauge defect class q_Σ is invariant under local rechartings, vanishes iff the defect is removable, and classifies fixed-cutoff genuinely noncentral sectors.

Corollary B.7 (Exact Markov adds the state factorization)**.** *On either the ordinary/central branch of the previous theorem or the genuinely noncentral higher-gauge branch, if in addition*

$$I_\omega(A_\delta : D_\delta \mid B_\delta) = 0,$$

or one passes to the explicitly stated idealized recoverability limit that reduces to exact Markovity, then

$$\rho_{A_\delta B_\delta D_\delta} = \bigoplus_\alpha p_\alpha \left(\rho_{A_\delta b^\alpha_L} \otimes \rho_{b^\alpha_R D_\delta} \right).$$

EC therefore gives the kinematic block decomposition, while exact Markovity is the extra state input that gives the HJPW normal form.

Remark B.8. Approximate recoverability gives controlled deviations from this normal form; it is not implied by EC alone. This is the fixed-cutoff topological package behind the consensus paper's quotient and record-language surface.

C Conditional Distributed-Systems and QECC Extensions of the Consensus Formalism

Honesty labels.

[**Established**] Follows from cited prior work or a complete argument given here.

[**Conditional**] True under additional assumptions not yet derived from OPH first principles.

[**Conjecture**] A plausible open direction, not a settled result.

B.1 Theorem 1 — QBFT Safety Bound

Definition C.1 (QBFT-style protocol)**.** *A consensus protocol is* QBFT-style *in this analysis if it satisfies the following three structural properties. The safety proof of Theorem C.2 uses all three; the theorem does not hold for protocols lacking any of them without a compensating change to the argument.*

(P1) ***One-vote-per-view.*** *Each honest node casts at most one vote per view number. A node that has already voted in view v ignores any later request to vote in view v.*

(P2) ***Certificate semantics.*** *A decision requires a valid quorum certificate: $2f + 1$ distinct, unforgeable, authenticated votes for the same value in the same view.*

(P3) ***DLS-style view-change.*** *If no certificate is produced within a timeout, every honest node increments the view number by one and a new leader is selected by a fixed deterministic rule. At GST, timeouts fire correctly and the view-change terminates in bounded rounds.*

The Istanbul BFT / QBFT protocol family [5, 6] satisfies (P1)–(P3) and is the intended instance.

Assumptions A1–A6.

(A1) **Partial synchrony (DLS).** Fixed but initially unknown bounds Δ (message delay) and Φ (processing rates). *Safety* holds unconditionally; *liveness* holds after the Global Stabilisation Time (GST).

(B2) **Byzantine fault model.** At most f observers behave arbitrarily; the remaining $n - f$ are honest.

(C3) **Optimal fault bound.** $n \geq 3f + 1$ (necessary: [3]; sufficient: [3]).

(D4) **Strong quorum connectivity.** Every quorum Q with $|Q| = 2f + 1$ is strongly connected within G: for any $u, v \in Q$ there is a directed path in G contained entirely in Q. This is strictly stronger than requiring the overlap graph of quorums to be connected, and is needed to propagate signed votes within a quorum.

(E5) **Message authentication.** All messages carry unforgeable digital signatures.

(F6) **OPH quorum overlap.** Any two quorums Q_a, Q_b of size $2f + 1$ satisfy $|Q_a \cap Q_b| \geq f + 1$ (guaranteed by (A3)).

Theorem C.2 (QBFT Safety Bound [Established, conditional on A1–A6])**.** *Under assumptions (A1)–(A6), any consensus protocol satisfying (P1)–(P3) of Definition C.1 and run over the OPH observer graph satisfies:*

(i) ***Safety.*** *No two honest observers finalise conflicting patch states.*

(ii) ***Liveness.*** *After GST, every honest observer finalises within $O(f \cdot \Delta)$ wall-clock time.*

(iii) ***Optimality.*** *The bound $f < n/3$ is tight.*

Proof sketch. Safety. Suppose O_a and O_b finalise $s_a \neq s_b$ in the same view. By (P2), each required a certificate of $q = 2f+1$ votes: sets Q_a, Q_b. By (A3): $|Q_a \cap Q_b| \geq (2f+1)+(2f+1)-(3f+1) = f+1$. By (A2), at most f are Byzantine, so $Q_a \cap Q_b$ contains an honest O^*. By (A4), O^*'s signed vote is path-reachable within both quorums. By (P1), O^* voted for at most one value — contradiction.

Liveness and Optimality follow from [4] (Thm. 4.4) and [3], cited directly.

Note on FLP. Fischer, Lynch, Paterson [2] is an impossibility result for fully asynchronous systems; it does not bear on achievability under partial synchrony (A1). □

B.2 Theorem 2 — Convergence of the OPH Repair Map

Definition C.3 (OPH Repair Map — Petz form)**.** *Let $\sigma \in \mathcal{D}(\mathcal{H})$ be a full-rank reference state and $\mathcal{N} : \mathcal{B}(\mathcal{H}) \to \mathcal{B}(\mathcal{K})$ a quantum channel. The* OPH repair map *is*

$$\mathcal{R}_{\sigma,\mathcal{N}}(\rho) := \sigma^{1/2}\, \mathcal{N}^{\dagger}\big(\mathcal{N}(\sigma)^{-1/2}\, \rho\, \mathcal{N}(\sigma)^{-1/2}\big)\, \sigma^{1/2},$$

where $\mathcal{N}^{\dagger}$ is the adjoint channel and inverses are taken on $\mathrm{supp}(\mathcal{N}(\sigma))$*.*

Remark C.4 (Petz map vs. trace-distance projection)*.* The closest-point trace-distance projection $\mathcal{P}_{\mathcal{S}}(\rho) := \arg\min_{\tau \in \mathcal{S}} \frac{1}{2}\|\rho - \tau\|_1$ is a different object from the Petz map: it is defined by a variational problem in trace-norm geometry and is not CPTP in general. The two coincide only in very special cases not automatic in the OPH setting. All subsequent properties refer exclusively to Definition C.3.

Proposition C.5 (Petz map CPTP — domain-restricted statement [Established, subject to domain restriction])**.** *Let σ have full support on $\mathcal{H}$.*

(a) $\mathcal{R}_{\sigma,\mathcal{N}}$ is completely positive.

(b) $\mathcal{R}_{\sigma,\mathcal{N}}$ is trace-preserving on $\mathrm{supp}(\mathcal{N}(\sigma))$*, i.e., on inputs ρ for which $\mathcal{N}(\sigma)^{-1/2}\rho\mathcal{N}(\sigma)^{-1/2}$ is well-defined.*

(c) If additionally $\mathcal{N}(\sigma)$ has full rank on $\mathcal{K}$, then $\mathcal{R}_{\sigma,\mathcal{N}}$ is CPTP on all of $\mathcal{B}(\mathcal{K})$.

If $\mathcal{N}(\sigma)$ is not *full rank on $\mathcal{K}$, then either (i) the domain must be restricted to* $\mathrm{supp}(\mathcal{N}(\sigma))$*, or (ii) pseudoinverses must replace the inverses (generalised Petz map; cf. [9]), or (iii) a regularisation $\mathcal{N}(\sigma) \mapsto \mathcal{N}(\sigma) + \varepsilon\mathbf{1}$ must be introduced. Note that full-rank σ does* not *prevent $\mathcal{N}(\sigma)$ from being rank-deficient: the channel may map the support of σ into a strict subspace of $\mathcal{K}$. In the OPH setting, whether $\mathcal{N}(\sigma)$ is full rank depends on the specific overlap channel and must be verified separately (open issue #62).*

Proof. Complete positivity follows from composing three CP operations:

(i) sandwiching by $\mathcal{N}(\sigma)^{-1/2}(\cdot)\mathcal{N}(\sigma)^{-1/2}$ on $\mathrm{supp}(\mathcal{N}(\sigma))$;

(ii) $\mathcal{N}^{\dagger}$;

(iii) sandwiching by $\sigma^{1/2}(\cdot)\sigma^{1/2}$.

Trace preservation in the full-rank case: Petz [7]; Fagnola–Umanità [8]. □

Proposition C.6 (Contraction [Conditional]). *Suppose $\mathcal{N}$ is strictly contractive with spectral gap $\lambda \in (0,1)$:*

$$\sup_{\rho\neq\tau} \frac{\|\mathcal{N}(\rho)-\mathcal{N}(\tau)\|_1}{\|\rho-\tau\|_1} \le \lambda < 1.$$

Then $\mathcal{R}_{\sigma,\mathcal{N}} \circ \mathcal{N}$ is contractive with coefficient depending on λ and $\mathrm{spec}(\sigma)$. *Establishing $\lambda < 1$ for the OPH overlap channel requires analysis of the OPH Hamiltonian (open issue #62).*

Conjecture C.7 (Spectral gap [Conjecture]). *The transfer operator $\mathcal{T}$ for iterated application of $\mathcal{R}_{\sigma,\mathcal{N}}$ has a positive spectral gap $\delta > 0$, implying exponential convergence. Proof from OPH first principles is open issue #63.*

Theorem C.8 (Exponential Convergence [Conditional on Conjecture C.7]). *Assuming Conjecture C.7 with gap $\delta > 0$,*

$$\tfrac{1}{2}\left\|\mathcal{R}^{\circ t}_{\sigma,\mathcal{N}}(\rho)-\sigma\right\|_1 \le C\, e^{-\delta t}$$

for $C > 0$ depending on ρ and $\mathcal{N}$.

B.3 Theorem 3 — QECC Correspondence

Notation. $N = \dim(\mathcal{H}) = 2^n$ for n physical qubits. Standard notation: $[[n,k,d]]$ stabilizer code; $K = 2^k$; quantum Singleton bound: $k \le n - 2(d-1)$.

Claim C.9 (Code distance and min-cut [Conditional — tightened]). *The identity* code distance = graph min-cut *holds for topological codes (surface/toric codes) whose logical operators correspond to non-contractible homological cycles in the code graph. It does* not *hold for a generic overlap graph.*

Suppose the OPH observer network is equipped with a surface-code-type construction on a planar or toroidal graph G_{OPH}, and the following conditions are satisfied:

(i) Logical X-type operators correspond to minimum-weight non-contractible cycles in the primal *chain complex of G_{OPH}; logical Z-type operators correspond to minimum-weight non-contractible cycles in the* dual *complex.*

(ii) Boundary conditions are such that no logical operator of weight strictly less than the min-cut of G_{OPH} exists.

(iii) For non-planar geometries, the relevant group is $H_1(G_{\mathrm{OPH}};\mathbb{F}_2)$ and the code distance equals the minimum over all non-trivial homology classes of the weight of a representative cycle.

Under (i)–(iii), $d = \mathrm{mincut}(G_{\mathrm{OPH}})$ [11, 12].

This claim requires an explicit construction of the topological encoding map, the primal/dual complex of G_{OPH}, and verification of (i)–(iii), none of which have been provided. The claim is conditional (open issue #113).

Conjecture C.10 (Communication complexity [Conjecture]). *The OPH consensus-repair protocol, realised as a quantum communication task, has per-round complexity $O(n\cdot\mathrm{poly}(d))$ (open issue #72; cf. [13]).*

Theorem C.11 (QECC Correspondence [Conditional / Conjecture]). *Suppose conditions (i)–(iii) of Claim C.9 hold. Then:*

(i) [Conditional] Code distance $d = \mathrm{mincut}(G_{\mathrm{OPH}})$.

(ii) [Established] The Knill–Laflamme QECC conditions [10] are satisfied for the logical subspace whenever the number of corrupted observers satisfies $t < d/2$.

(iii) [Conjecture] Per-round communication complexity is $O(n \cdot \mathrm{poly}(d))$.

B.4 Theorem 4 — Asynchronous Convergence

Why fairness alone does not give a probability-1 statement. Standard strong fairness guarantees that every enabled action fires infinitely often along any fair schedule; it does not impose a probability space on the set of schedules. A convergence statement of the form "converges with probability 1" requires a measure on schedules and does not follow from fairness alone. The FLP impossibility result [2] confirms that even strong fairness is insufficient for bounded-time consensus in a fully asynchronous system. The phrase "with probability 1" has been removed from Theorem C.12; the convergence is per-schedule and topological.

Additional assumptions for a quantitative bound.

(B1) Finite known bound Δ on message delay after GST.

(B2) Finite bound Φ on processing rates.

(B3) $f < n/3$.

Theorem C.12 (Eventual Convergence [Established under fairness only])**.** *In a fully asynchronous OPH observer network under standard strong fairness, iterated application of* $\mathcal{R}_{\sigma,\mathcal{N}}$ *converges to a consensus state* σ^* *in the following sense: for every* $\varepsilon > 0$ *and every strongly fair schedule, there exists a step* $T(\mathit{schedule}, \varepsilon) < \infty$ *such that*

$$\tfrac{1}{2}\left\|\mathcal{R}^{\circ t}(\rho) - \sigma^*\right\|_1 < \varepsilon \quad \text{for all } t \geq T.$$

This is a per-schedule topological statement. No probability measure on schedules is assumed or needed; no uniform finite bound on T *follows from fairness alone.*

Theorem C.13 (Quantitative Convergence [Conditional on (B1)–(B3)])**.** *In a partially synchronous OPH observer network satisfying (B1)–(B3), after GST every honest observer reaches consensus within* $T = O(f \cdot \Delta)$ *wall-clock time (by applying the DLS framework [4], Thm. 4.4, to the OPH repair protocol; requires (B1) and (B2) explicitly and does not follow from fairness alone).*

B.5 Open Problems

- **#62** Derive repair map from OPH dynamics; verify full-rank condition for $\mathcal{N}(\sigma)$ (prerequisite for upgrading Proposition C.6 and the domain condition of Proposition C.5).
- **#63** Prove spectral gap from an OPH Lyapunov functional (upgrades Conjecture C.7 to Theorem C.8).
- **#68** Quantum observable-level confluence.
- **#69** Continuum / refinement limit of Theorems C.2 and C.13.
- **#72** Communication complexity (upgrades Conjecture C.10).

- **#73** Re-export repair map into OPH language.
- **#113** Construct topological encoding map for Claim C.9.

References

[1] B. Müller, *Observers Are All You Need*, 2026.
Available at https://oph-book.floatingpragma.io/papers/observers_are_all_you_need.pdf.

[2] M. J. Fischer, N. A. Lynch, and M. S. Paterson, "Impossibility of distributed consensus with one faulty process," *J. ACM* **32** (1985), no. 2, 374–382.

[3] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," *ACM Trans. Program. Lang. Syst.* **4** (1982), no. 3, 382–401.

[4] C. Dwork, N. A. Lynch, and L. Stockmeyer, "Consensus in the presence of partial synchrony," *J. ACM* **35** (1988), no. 2, 288–323.

[5] H. Moniz, *The Istanbul BFT Consensus Algorithm*, arXiv:2002.03613, 2020.

[6] R. Saltini et al., *QBFT Formal Specification and Verification*. Available at https://github.com/Consensys/qbft-formal-spec-and-verification.

[7] D. Petz, "Sufficient subalgebras and the relative entropy of states of a von Neumann algebra," *Commun. Math. Phys.* **105** (1986), no. 1, 123–131.

[8] F. Fagnola and V. Umanità, "Generators of detailed balance quantum Markov semigroups," *Infinite Dimensional Analysis, Quantum Probability and Related Topics* **13** (2010), no. 3, 459–486.

[9] M. Junge et al., "Universal recovery maps and approximate sufficiency of quantum relative entropies," *Ann. Henri Poincaré* **19** (2018), no. 8, 2505–2555.

[10] E. Knill and R. Laflamme, "Theory of quantum error-correcting codes," *Phys. Rev. A* **55** (1997), no. 2, 900–911.

[11] A. Kitaev, "Fault-tolerant quantum computation by anyons," *Ann. Phys.* **303** (2003), no. 1, 2–30.

[12] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, "Topological quantum memory," *J. Math. Phys.* **43** (2002), no. 9, 4452–4505.

[13] H. Buhrman, R. Cleve, and A. Wigderson, "Quantum vs. classical communication and computation," in *Proceedings of STOC 1998*, pp. 63–68.

[14] O. Fawzi and R. Renner, "Quantum conditional mutual information and approximate Markov chains," *Commun. Math. Phys.* **340** (2015), 575–611, arXiv:1410.0664.