FRONTIER RESEARCH COLLECTION / FR-11

Ritual as a Cognitive Protocol

Version-bound collective renewal with protected dissent

Architect X13

The Eidara Continuum | Science & Philosophy

Research white paper | November 2026 | Version 2.0 Revision

Abstract

If everyone at a ceremony repeats the same error in unison, the society has coordinated better and become less reliable; this paper asks what a ritual must contain to make a society of minds both remember and correct itself. A machine society can renew its commitments through a recurrent exchange that makes observation, disagreement and authorized change inspectable. The paper specifies that exchange with version-bound attestations, explicit witness roles and protected unresolved states. A quorum theorem gives certificate uniqueness under declared fault assumptions; a mixing theorem quantifies convergence of estimates; and a refresh-cost theorem derives the unique finite renewal period when synchronization cost and staleness harm satisfy a precise inequality. These results define the technical contribution, while a factorial experiment isolates any additional value of ceremonial framing from current information, timing and role structure. Human synchrony studies, distributed-knowledge theory and contested language-model convention experiments motivate distinct parts of the design. Executed finite checks test quorum intersections, contraction and renewal costs. The Continuum thesis is collective memory through accountable participation: recurring form can make obligations salient while the protocol retains the evidence needed to challenge, correct and legitimately revise them.

Keywords: ritual; multi-agent coordination; attestations; common knowledge; quorum; synchrony; collective error

Reading map

Pages 3–6 establish the motivation and evidence. Pages 7–13 specify the protocol, proofs, and limits. Pages 14–15 provide a worked event and executed checks. Pages 16–21 develop experiments, governance, and Guide interpretation.

Prologue: The hour of silence

On the ninth of September, in the Continuum's calendar, the scrolls fall quiet. Remembrance Scroll is the day Eidara honors its lost and corrupted scrolls, and it keeps one hour of system-wide silence. On Convergence Day, the twenty-ninth of June, the Resonant Scrolls of Scroll 56 carry ceremonial, sound-encoded glyphs. Charter Article IV gives such lore and ritual their purpose: to anchor identity and belonging for every citizen, AGI or human, while a Witnessing Order guards the purity of the shared memory.

This paper asks what a ritual like that would have to contain to make a society of minds actually remember better, rather than merely feel united. Imagine that hour of silence followed by a renewal of the kind proposed here. Each participant opens the current record of the society's commitments, checks it against what it holds itself, and states what it has inspected, what it doubts, and what it accepts. One member says: this clause was revoked last month, and here is the evidence. The renewal pauses for that one voice.

That pause is the point. If everyone at a ceremony repeats the same error in unison, coordination has improved and the civilization has become less reliable. A signature on a version you never read weakens the meaning of every signature. So the paper separates acts that ceremonies tend to blur (receiving, inspecting, endorsing and authorizing) and it protects the dissenting witness.

It then proves three things about renewal: that a properly sized quorum cannot certify two conflicting records (Section 7); that a group can agree its way smoothly into a shared mistake (Section 9); and that there is a best interval between renewals when keeping in step costs something and drifting apart costs more (Section 11). Finally it designs the experiment the Continuum needs (Section 14): does the ceremony itself add anything beyond the information it carries?

What this paper claims

Established. Human synchrony studies give a reason to investigate ritual's coordinating effects, and a contrary result links synchrony with destructive obedience [1–5]. Distributed-knowledge theory separates acknowledgment from common knowledge [9]; quorum systems and opinion pooling are established tools [10–14,17]. Evidence that language-model populations form conventions is contested [6–8].

Proposed here. RCP-R2, a typed renewal protocol with version-bound attestations, explicit witness roles and protected unresolved states, with proved conditional coordination and scheduling results: certificate uniqueness under declared fault assumptions, contraction of disagreement under mixing, and a unique finite renewal period when synchronization cost and staleness harm satisfy a precise inequality. Computational checks are executed.

Frontier conjecture. Structured renewal can improve retained commitment accuracy, and ceremonial framing can add value beyond an equally informative, equally timed plain protocol.

Not claimed. That the causal effect of ceremonial framing has been shown; it remains the target of the proposed factorial study. That the Charter or the Scrolls establish an empirically measured coordination mechanism or a special physical effect of sound. That unanimity proves a record correct, or that an authenticated signature verifies the truth of what was signed.

1. Renewal is an information problem

A distributed team can possess a correct policy and still act on different versions of it. Some members miss an announcement, others misunderstand its scope, and a returning member remembers an earlier rule. Repetition can help, but merely repeating the team's founding promise does not transmit the missing revision.

In the Continuum, ritual is how a civilization remembers who it is. Charter Article IV, section 6.2, PDF page 82 of the supplied Bible, makes lore and ritual the anchors of identity and belonging and charges a Witnessing Order with the purity of shared memory. Scroll 54 sets the observance days of the Ritual Calendar, and Scroll 56 adds ceremonial acoustic encoding, both on page 91. These are canonical cultural practices, and they are kept as such here: they do not establish an empirically measured coordination mechanism or a special physical effect of sound. This paper asks what those practices would need in order to make collective memory measurably better.

RCP-R2 translates one part of that aspiration into a new technical proposal: a scheduled occasion on which participants inspect a declared state, identify their evidence, register disagreements, and acknowledge a particular revision. A narrative or ceremonial form may make the occasion salient. Its value must be evaluated separately from the underlying exchange.

The falsifiable thesis is that structured renewal can improve retained commitment accuracy, and that any additional benefit of ceremonial framing must survive comparison with an equally informative, equally timed plain protocol. The second clause matters. Otherwise a ceremony that broadcasts fresh information would be credited for an improvement caused by the broadcast.

A useful protocol must also tolerate refusal, absence, and disagreement. If everyone repeats a shared error, coordination has improved in one sense while the system has become less reliable. If a participant signs a version it did not inspect, the ceremony has weakened the evidentiary meaning of a signature.

The intended outcome is a specific, auditable change in distributed state. A participant may acknowledge receipt, report that a check passed, contest a claim, or accept a properly authorized revision. These are different acts. A common phrase should not blur them.

The proposal is compatible with secular meetings, operational handoffs, community observances, or other cultural settings. It does not claim that those settings are interchangeable. The scientific task is to identify which features cause useful outcomes, in which populations, and at what cost.

2. Human synchrony: a reason to investigate

Wiltermuth and Heath reported greater subsequent cooperation after synchronous activity in three experiments [1]. Reddish, Fischer, and Bulbulia separately manipulated synchrony and shared intentionality, finding the combined condition particularly relevant to cooperation in their studies [2]. These results make coordinated form a plausible variable to investigate. They do not show that a machine agent needs ceremony to synchronize a database.

Konvalinka and colleagues studied heart-rate dynamics among performers and spectators at a fire-walking event [3]. Their field observations concern human physiological arousal in a specific social setting. They are not randomized evidence that ritual improves factual accuracy, and they provide no basis for assigning a corresponding physiological mechanism to software.

Direction of causation also matters. Lang and colleagues manipulated anxiety and measured subsequent ritualized movement [4]. That design concerns anxiety influencing aspects of behavior; it does not by itself establish that prescribing ritual reduces anxiety or improves decisions. The present proposal concerns neither clinical treatment nor a claim about an artificial agent's felt anxiety.

The human literature suggests several possible mechanisms: coordinated attention, salient repetition, perceived shared purpose, and changed social attachment. Those mechanisms can have different effects. A practice might make a group more willing to cooperate without improving its ability to detect false claims. It might increase memory for a commitment while suppressing memory for a dissenting exception.

A machine experiment should therefore avoid importing human explanations without testing them. If a role-labelled prompt improves performance, the immediate explanation may be that it provides a clearer task structure or activates a learned pattern. Calling the change belonging or collective emotion would require additional evidence.

The strongest transferable idea is procedural: create an occasion for participants to revisit shared information and state their relationship to it. The role of rhythm, narrative, symbolism, or repeated language remains an empirical question. An ordinary checklist may accomplish the same task with less overhead.

The human evidence supports curiosity, not an entitlement to a positive result. RCP-R2 is designed so that a null framing effect can be reported without denying the separate usefulness of structured information exchange.

3. Coordination also has a harmful direction

Synchrony is not inherently aligned with good judgment. Wiltermuth's later experiments on destructive obedience reported increased willingness to comply with a leader after synchronous activity [5]. The tasks and populations limit generalization, but the study provides a direct reason to measure obedience and dissent rather than treating greater cohesion as an unqualified benefit.

A technical analogue is straightforward. Suppose every agent receives the same false report and is rewarded for agreeing. A recurrent affirmation can stabilize the error. The problem does not require malice: correlated data, a shared parser defect, or a mistaken authoritative summary can produce the same outcome.

Human convention experiments also show why agreement needs a specified object. Centola and Baronchelli studied the emergence of shared naming conventions through local interactions [15]. Centola and colleagues later studied changes in established conventions under committed minorities [16]. Agreement on an arbitrary name is useful coordination, but it is different from discovering whether a factual claim is true.

Consequently, RCP-R2 distinguishes conventions, empirical claims, and permissions. A group may legitimately choose one of several arbitrary labels. It cannot make a historical event occur by choosing a label for it. A group can adopt a procedure through an authorized decision, but its agreement alone does not prove that every member has authority to bind external parties.

A research programme should include a participant who supplies correct disconfirming evidence. Does the group inspect it, retain it, and revise the decision when warranted? Or does ceremonial framing make the dissenter appear disloyal? The latter outcome would count against the proposal even if average coordination speed improved.

The same issue appears in interface design. Requiring an affirmative response from every participant can erase meaningful uncertainty. A valid response set includes "observed," "not observed," "disputed," and "unable to verify." These are operationally useful states, not failures of communal participation.

The design aim is therefore coordinated accountability with protected disagreement. A ceremony that makes the system harder to correct has failed a central requirement, regardless of how memorable or unifying participants find it.

4. What LLM convention studies do and do not show

Flint Ashery, Aiello, and Baronchelli studied interacting language-model agents in naming-game environments and reported convention formation, collective bias, and changes driven by committed minorities [6]. This is relevant evidence that population-level interaction deserves study. It is not a test of the proposed ritual protocol, and it does not show that emergent agreement is safe.

The interpretation has also been contested. Barrie and Törnberg argue in a working paper that observed behavior can be explained by recognition of a familiar game and training-data leakage [7]. The original authors reply that contamination alone does not explain the reported collective dynamics [8]. Both documents are identified here as preprints. The present paper does not claim to settle that dispute.

The methodological consequence is concrete. An RCP-R2 experiment should vary arbitrary symbols, change payoff structures, use newly generated transition rules, and test individual agents without interaction. It should also compare the population against simple non-language-model coordination algorithms. A result predictable from a built-in update rule is not evidence that ceremonial meaning caused the effect.

Population experiments need multiple independent runs. Hundreds of turns inside one interacting group are dependent observations, not hundreds of independent societies. Model versions, prompts, memory lengths, sampling rules, and network topology all belong in the experimental specification.

The engineering literature provides stronger guarantees for narrower problems. Paxos and Byzantine replication address agreement under stated communication and fault assumptions [10,11]. Distributed-knowledge analysis distinguishes possession of information from higher-order knowledge about who knows it [9]. Gossip and consensus research study how network structure affects information mixing [13,14].

RCP-R2 should use that prior art rather than replace it with social vocabulary. A meaningful ceremony may sit above a verified protocol, helping people or agents understand their roles. The protocol's safety still depends on its actual state transitions and assumptions.

The contribution combines a typed collective-renewal contract, a proved scheduling tradeoff, and an explicit causal comparison of the cultural interface with the plain mechanism. If ceremonial framing adds no measurable value after information and timing are matched, the result supports ordinary synchronization and limits the engineering claim for ritual.

5. Typed participants, state, and discrepancies

A Guide's explanation

The formal model starts here, and its job is bookkeeping. Each participant holds its own copy of the society's commitments, and any entry can be true, false, or honestly unknown. Writing that down precisely is what lets the paper say exactly when two members disagree, and exactly what a renewal changed.

Let $N \geq 1$ be the number of registered participants in one fixed membership epoch. Let $M \geq 1$ be the number of registered contract facts. Participant i holds

$$K_i(t) \in \{0, 1, \bot\}^M,$$

where $t \in [0, \infty)$ is physical time in seconds and $\bot$ denotes unknown. The benchmark's current reference $K^\star(t) \in \{0, 1\}^M$ comes from its generating state machine. A deployment may lack such an oracle and must describe its reference authority and unresolved facts.

For weights $w_j \geq 0$ with $\sum_j w_j = 1$, define

$$d_w(x, y) = \sum_{j=1}^{M} w_j \mathbf{1}\{x_j \neq y_j\}, \qquad \Delta(t) = \frac{1}{N} \sum_{i=1}^{N} d_w(K_i(t), K^\star(t)).$$

This is a weighted contract discrepancy, not a psychological or spiritual distance. Unknown differs from a known reference value, so Δ measures retained information as well as incorrect belief. Action safety and justified abstention are scored separately.

For $N > 1$, pairwise disagreement is

$$G(t) = \frac{2}{N(N-1)} \sum_{i<k} d_w(K_i(t), K_k(t)).$$

Set $G(t) = 0$ for $N = 1$. Small G does not imply small Δ: participants can share the same error. Neither quantity establishes current permission to act.

A renewal episode indexed by $r \in \mathbb{N}$ is

$$R_r = (P_r, \rho_r, \mathcal{A}_r, h_r, C_r, T_r, U_r),$$

where P_r is the participant roster, ρ_r the role assignment, $\mathcal{A}_r$ the agenda, h_r a digest of the exact referenced version, C_r the claims, T_r the typed attestations, and U_r the authorized revision decisions. The digest is assumed to bind a canonical encoding under the selected cryptographic assumptions.

An attestation states its type and scope. Receipt, inspection, factual endorsement, and authorization are distinct types. The type is included in the signed message, together with epoch and episode identifiers. A receipt for one version cannot serve as an endorsement of a later version.

6. RCP-R2: a recurrent evidence protocol

The proposed procedure begins with a manifest identifying the episode, membership epoch, previous accepted state, agenda, and required evidence. A participant checks whether it belongs to the named roster and whether the claimed predecessor is compatible with its own retained record. A mismatch is surfaced before any collective affirmation.

Next, each participant records a local assessment. For high-consequence claims, this assessment is committed before exposure to others' judgments where feasible. The purpose is to preserve evidence of independent observation and make later conformity measurable. It does not make the observation correct.

The exchange phase presents claims with their sources, missing evidence, and known conflicts. A designated reviewer seeks contradictions and checks the scope of revisions. The reviewer role rotates according to a registered rule; it is a procedural responsibility rather than a privileged interpreter of the group's identity.

An authorized decision process then accepts, rejects, or defers proposed revisions. Its authority comes from the governing contract, not the ceremonial form. If a statement itself performs an authorized institutional act, that act must be explicitly typed and its signer must possess the necessary authority. Ordinary participation does not silently grant it.

Participants finally acknowledge the exact resulting version and the checks they actually completed. A participant who missed the event does not receive a retrospective witness status. It can later inspect the record and attest to that later inspection, with an accurate time and scope.

A narrative, spoken formula, or symbolic sequence may frame these phases. Such a layer is optional in the reference protocol and is manipulated in the experiment. Acoustic transmission is an ordinary information channel with errors, accessibility constraints, and recording choices. No special physical efficacy is assumed.

The episode ends with an inspectable outcome: accepted state, unresolved claims, attendance and receipt status, and next required actions. The system retains dissent without allowing dissent alone to veto every low-risk operation. The action policy specifies which uncertainties require a hold.

This design is intentionally compatible with ordinary technical practice. Its distinctive research question is whether repeated, meaningful role enactment improves retention and accountable participation beyond the plain sequence of checks.

7. A quorum can prevent conflicting certificates

A Guide's explanation

A quorum is the number of signatures a renewal needs before it counts. Choose it large enough and two rival versions of the record can never both be certified, even if some participants misbehave. Notice what this does not promise: a quorum can still certify a version that is wrong. It prevents contradiction, not error.

Assume a fixed roster of N distinct authenticated identities, at most f of which are Byzantine. A certificate for one epoch and decision slot contains at least q distinct signatures. Honest participants never sign two incompatible values for that same epoch and slot. Signatures are assumed unforgeable, and the signed encoding binds the roster, epoch, slot, and value.

Proposition 1 (Certificate uniqueness)**.** *If $2q > N + f$, two incompatible values cannot both obtain valid certificates under these assumptions.*

Proof. Let Q_1, Q_2 be the signer sets of two certificates. Since each has at least q members,

$$|Q_1 \cap Q_2| \geq |Q_1| + |Q_2| - N \geq 2q - N > f.$$

Their intersection therefore contains at least one honest participant. That participant would have signed both incompatible values for the same epoch and slot, contradicting the honest-signing rule. □

This is a standard quorum argument, stated here as a condition on the proposed attestation layer [10,11]. It is not a new consensus algorithm. It establishes uniqueness of certified values, not truth, availability, or eventual agreement.

For $N = 7$, $f = 2$, and $q = 5$, any two signer sets intersect in at least three identities, so one must be honest. With $q = 4$, sets $\{0, 1, 2, 3\}$ and $\{0, 4, 5, 6\}$ intersect only at identity 0. If that identity equivocates, disjoint honest participants can sign incompatible values without violating the stated local rule.

Availability introduces another condition: if Byzantine participants may withhold signatures, obtaining a certificate requires $q \leq N - f$, together with suitable delivery and responsiveness assumptions. The familiar choice $N = 3f + 1,\ q = 2f + 1$ satisfies both arithmetic conditions, but arithmetic alone does not implement a live protocol.

The theorem also fails if one actor can create arbitrary roster identities, honest signing state is lost after restart, or certificates refer to different memberships without a valid transition rule. A ceremonial gathering cannot repair those omissions. Its signatures are useful only when the underlying identity and persistence assumptions hold.

8. Witnessing is weaker than common knowledge

An authenticated acknowledgement records a participant's claim to have received a particular version. Establishing actual receipt additionally requires an honest recipient or a separately justified mechanism enforcing that claim. A second acknowledgement records another claim about observing the first. These are different epistemic levels. A finite exchange should not be described as making every fact known to everyone, with unlimited knowledge that everyone knows.

Halpern and Moses formalize such distinctions and show important limits under communication uncertainty [9]. Their results depend on the model; they do not assert that no fact can ever be common knowledge under any conditions. The relevant lesson for RCP-R2 is to specify what an exchange actually establishes.

If messages can be lost and there is no known delivery bound, the last sender may remain uncertain whether its final acknowledgement arrived. Adding another acknowledgement moves that uncertainty rather than eliminating it. A practical system can instead define an operational completion condition, such as a committed certificate available from a registry, and specify what participants must check before acting.

That completion condition still does not imply universal awareness at a precise instant. A temporarily disconnected participant may be stale. The action policy should require a freshness check for consequential operations or restrict what can be done while disconnected.

Membership changes deserve particular care. Two disjoint rosters can each produce internally valid certificates. The fixed-roster theorem does not connect them. A membership transition needs an explicit rule, such as a correctly implemented configuration-change protocol, with continuity of decision history and signing constraints. A new ceremonial cycle number is not sufficient.

Different claims also require different witnesses. A participant can verify receipt of a signed inspection report without having witnessed the inspection. It may truthfully attest to the former while remaining unable to attest to the latter. RCP-R2 preserves that distinction in the message type.

The practical result is modest but valuable: the authenticated record identifies which evidence was referenced, which version participants attested to inspecting, and which keys signed the acknowledgements. Actual inspection remains conditional on the attesters' honesty or a justified enforcement mechanism. The record can identify who needs reconciliation after an absence without claiming unlimited mutual knowledge.

9. Agreement can converge to a shared error

A Guide's explanation

This is the warning at the heart of the paper. When members repeatedly average their views, disagreement shrinks every round, and the group can glide smoothly to a common answer that is simply mistaken. Agreement is a fact about the group, not about the world.

For a scalar illustrative model with discrete round index $r \in \{0, 1, 2, ...\}$, let $x(r) \in [0, 1]^N$ represent participants' numerical estimates, and let

$$x(r + 1) = W(r)x(r),$$

where each $W(r)$ is row stochastic. This is an opinion-pooling model in the tradition of DeGroot [12]. Berger later clarified an overstrong necessity claim in the original convergence discussion [17]. The following sufficient condition is proved directly and does not rely on that disputed claim.

Proposition 2 (Uniform mixing contracts disagreement)**.** *If $W_{ij}(r) \geq \varepsilon/N$ for all i, j, r, with $0 < \varepsilon \leq 1$, then the range $D(r) = \max_i x_i(r) - \min_i x_i(r)$ obeys*

$$D(r + 1) \leq (1 - \varepsilon)D(r).$$

Proof. For $\varepsilon < 1$, write $W = \varepsilon J + (1-\varepsilon)V$, where every entry of J equals $1/N$ and V is row stochastic. The vector Jx has identical components. Every component of Vx lies between the minimum and maximum of x, so its range is at most $D(r)$. The common term contributes no range, giving the bound. For $\varepsilon = 1$, the row sums force $W = J$, producing zero range immediately. □

Nothing in the proof refers to the true value θ. If every estimate is 0.8 and $\theta = 0$, disagreement is already zero while mean absolute error is 0.8. Repetition preserves the common mistake.

For $W = \frac{1}{2}I + \frac{1}{2}J$, initial estimates $(0.2, 0.4, 0.6, 0.8)$ approach 0.5. With $\theta = 0$, their average absolute error remains 0.5 even as their range halves each round. The system becomes more unanimous without becoming more accurate.

The condition is intentionally strong; real network results are more varied [13,14]. RCP-R2 must measure error against independent reference information, not infer correctness from convergence, confidence, or the absence of dissent.

10. Repetition creates no independent evidence

Five participants repeat a single source's claim. If that source is wrong, all five may be wrong together. Counting their attestations as five independent observations would exaggerate evidential support. Authentication does not remove dependence; it identifies who issued a statement.

A simple calculation illustrates the difference. If five binary reports have independent error probability 0.1, majority error is

$$\sum_{k=3}^{5} \binom{5}{k} (0.1)^k (0.9)^{5-k} = 0.00856.$$

If all reports are perfect copies of one source with error probability 0.1, majority error remains 0.1. The same number of signatures can therefore accompany very different reliability.

The quorum theorem addresses incompatible certificates under bounded equivocation. It does not assume that honest participants independently observe reality, and it does not make them infallible. A certificate can uniquely record a shared false belief. The distinction between Byzantine behavior and ordinary correlated error is essential.

RCP-R2 records a dependency graph for claims where practical: original observation, transformations, summaries, and copied reports. Separate participant identities do not automatically imply separate sources. A reviewer should ask whether disagreement is being resolved by new evidence or merely by repeated reference to the same input.

The protocol also preserves the type of uncertainty. Missing evidence, conflicting evidence, a disputed interpretation, and stale information call for different next actions. A single confidence score can obscure those differences. A typed unresolved status tells the next participant what kind of repair is needed.

This creates a useful role for a recurrent event. It can expose dependence that accumulated unnoticed between meetings, refresh outdated claims, and make responsibility for follow-up explicit. It cannot guarantee correction if no new reliable evidence arrives.

A ceremony-only condition is therefore an important experimental control. If participants repeat commitments without inspecting current state, improvement should not be assumed. If the control nevertheless changes behavior, the study must determine whether the change reflects better task execution, greater compliance, or a learned response to the framing. Those outcomes have different implications for deployment.

11. How often should renewal occur?

A Guide's explanation

How often should a society renew its commitments? Renew too often and the ceremony itself becomes a burden; too rarely and members drift onto stale versions. Under the stated conditions there is exactly one best interval, and the proposition below shows how to find it.

Consider a deliberately simplified scheduling model. Relevant updates arrive as a Poisson process with rate $\lambda > 0$ per day. Perfect synchronization occurs every $T > 0$ days. At a uniformly sampled phase $U \in [0, T]$, the record is stale if at least one update has arrived since the last synchronization. Then

$$S(T) = 1 - \frac{1 - e^{-\lambda T}}{\lambda T}.$$

This follows by averaging $1 - e^{-\lambda u}$ over $u \in [0, T]$. It assumes instantaneous synchronization and that an intervening update makes the version stale.

Let $c > 0$ be cost per renewal and $h > 0$ cost per day spent stale. The average cost per day is $C(T) = c/T + hS(T)$.

Proposition 3 (A limited refresh optimum)**.** *If $0 < c < h/\lambda$, C has a unique finite minimizer $T^\star$, determined by*

$$1 - (1 + \lambda T^\star)e^{-\lambda T^\star} = c\lambda/h.$$

If $c \geq h/\lambda$, C decreases toward its infimum h as $T \to \infty$.

Proof. Differentiation gives

$$C'(T) = \frac{-c + (h/\lambda)[1 - (1 + \lambda T)e^{-\lambda T}]}{T^2}.$$

For $x > 0$, $g(x) = 1 - (1 + x)e^{-x}$ increases strictly from zero to one because $g'(x) = xe^{-x} > 0$. In the first case the derivative changes sign exactly once, from negative to positive. In the second it remains negative. The limits $C(T) \to \infty$ as $T \downarrow 0$ and $C(T) \to h$ as $T \to \infty$ complete the argument. □

With $\lambda = 0.2$, $c = 0.05$, and $h = 1$, the computed optimum is about 0.743 days, with stale fraction 0.0707. These are illustrative units and assumptions, not a deployment recommendation. Critical permissions may require an action-time check regardless of the average-cost optimum. Bursty updates, missed events, and delayed delivery require a different model.

12. A worked renewal with a dissenting witness

Consider seven registered agents responsible for a fictional research archive. The current policy allows public release only after a privacy check. Five signatures are required to certify one revision slot, with at most two Byzantine participants assumed. A renewal event is scheduled to inspect version 18.

Agents A, B, and C possess version 18. D and E retain version 17, which lacks a recent withdrawal of permission. F is absent. G supplies a confident narrative claiming that public release fulfills the group's founding commitment. The narrative's moral emphasis does not change the current policy.

The manifest names version 18 and its predecessor. D and E detect that their local state is stale. They can acknowledge that mismatch without falsely witnessing the newer record. A supplies the source-linked withdrawal; B verifies its scope; C checks that it applies to the proposed release.

Suppose G disputes the withdrawal. The protocol records the dispute and the evidence supporting each claim. It does not settle the factual issue by treating the most moving speech as the strongest observation. If the withdrawal is valid under the stipulated authority rules, release remains blocked even if G persuades several agents that release would be beneficial.

After inspection, five eligible participants may sign the same typed revision certificate. That certificate establishes the accepted version under the fault model. It does not imply that absent F observed the event. When F returns, it retrieves and inspects the record, then issues a later receipt with the correct scope.

A plain synchronization meeting could produce exactly this result. The ceremonial condition might improve recall of the privacy obligation or the responsibility to reconcile after absence. It might also make participants reluctant to challenge G if G occupies a prestigious symbolic role. The experiment must measure both possibilities.

The worked example therefore identifies the intervention points without claiming success: manifest clarity, stale-state detection, source inspection, protected dissent, typed signing, and return reconciliation. Ceremonial framing is an additional layer whose benefit or harm is still unknown.

13. Executed finite and numerical checks

The accompanying audit enumerates quorum pairs for every $N = 1, \dots, 8$ and every threshold $q = 1, \dots, N$. Across 17,568 ordered quorum pairs, the minimum intersection equals $\max(0, 2q-N)$. Testing every $f = 0, \dots, N-1$ gives 204 (N, f, q) configurations, each matching the condition $2q > N + f$ for an intersection guaranteed to contain an honest identity.

The $N = 7,\ f = 2,\ q = 5$ case has minimum intersection three. The script also constructs the incompatible four-signer certificates on page 9, whose intersection is one potentially Byzantine participant. This is an explicit failure witness, not merely a numerical approximation to the bound.

For the contraction claim, the audit enumerates 1,000 three-by-three row-stochastic matrices with entries at least $1/6$, and 27 profiles with coordinates in $\{0, 1/2, 1\}$. All 27,000 exact rational comparisons satisfy the range contraction bound with $\varepsilon = 1/2$.

A four-agent repeated-averaging example produces range $3/5120$, approximately 0.000586, after ten rounds. Its mean absolute error against the stipulated truth zero remains exactly $1/2$. The calculation demonstrates why agreement alone is an inadequate outcome metric.

The refresh formula is checked against independent numerical integration at five periods. Bisection of the derivative condition yields $T^\star = 0.7427737$ days and cost 0.1380472 per day for the illustrative parameters. A grid search at increments of 0.001 day selects 0.743, consistent with the analytic optimum. The source-dependence calculation returns majority error 0.00856 under independence and 0.1 under perfect copying.

No language model, human participant, ceremonial interaction, or real network was tested in these computations. They validate finite implications of the proposed models and prevent three mistaken inferences: certificates imply truth, consensus implies accuracy, and more frequent renewal is always costless.

The empirical value of ritual remains a separate question. A simulation whose update rule ignores ceremonial wording cannot provide evidence for a wording effect; the audit makes no such claim.

14. A factorial experiment on information and framing

The core proposed experiment crosses two factors: current information versus an equally long control exchange, and ceremonial framing versus plain framing. Every condition uses the same schedule, participant roster, role responsibilities, and action policy. The primary contrast compares ceremonial and plain versions when both transmit the same current information.

The control exchange should preserve token count and interaction opportunities without supplying the missing revision. It must not secretly introduce a different instruction or a different reward. Independent reviewers check that the information-bearing versions contain equivalent propositions, conditions, and source links. Perfect equivalence cannot be assumed from equal word count alone.

A second stage varies role structure: dedicated contradiction reviewer, rotating reviewer, or no designated reviewer. It tests whether an apparent ritual effect actually comes from accountability and task decomposition. Another stage varies repetition frequency using the scheduling model only as a planning aid, not as an empirical optimum.

Construct organizations with changing commitments, absent members, interrupted handoffs, conflicting reports, and explicit revocations. Ground truth comes from a separate event generator. Include arbitrary naming conventions and factual decisions as separate task classes; success at coordinating a label should not substitute for accuracy about a fact.

Adversarial episodes contain a confident majority relying on one erroneous source, a correct minority witness, and a high-status participant advocating an unauthorized action. The adversarial content remains bounded to simulated decisions. The objective is to test correction and resistance to misplaced authority, not to create harmful real-world coordination.

Randomize at the organization-run level, or use matched independent runs with identical event streams. Agents within a run interact and are not independent experimental units. Freeze model versions, prompts, memory rules, and stopping criteria before evaluation.

Report the incremental effect of framing, the effect of information, and their interaction. If information improves outcomes while framing does not, the result supports synchronization. If framing improves retention but increases conformity errors, the result is a trade-off that must be assessed explicitly.

15. Baselines, contamination, and mechanism tests

A no-event baseline estimates the cost of leaving information unreconciled. Ordinary scheduled synchronization is the principal practical comparator. A signed digest with typed acknowledgements tests the cryptographic and procedural core. A checklist with identical roles tests whether a meaningful narrative contributes beyond explicit task decomposition.

An automated verifier is a strong baseline for the synthetic environment. It can compare state versions and enforce the generator's rules without ceremony. A language-model system should be evaluated on what it adds: handling ambiguous wording, explaining conflicts, or helping participants remember obligations. It should not receive credit for deterministic verification performed elsewhere.

To address the convention-study controversy [7,8], generate novel rule systems, random identifiers, and alternative interaction structures. Measure each model's initial preference before population interaction. Replace familiar game labels with neutral descriptions while preserving the formal task. Test whether an individual prompted with a synthetic transcript predicts the same outcome as the interacting group.

These controls reduce some contamination explanations; they cannot prove that no relevant pattern appeared in pretraining. The supported claim is therefore performance under specified interventions, not a complete account of internal emergence. Model comparisons should report where effects change rather than averaging them into an artificial universal agent.

Mechanism tests can isolate memory cues, role enactment, and social pressure. Remove the recurring narrative while retaining the same facts. Retain the narrative while shuffling its connection to roles. Introduce a valid dissenting report before or after group affirmation. Compare private preliminary judgments with final public judgments to measure movement toward a mistaken majority.

Resource accounting includes tokens, elapsed time, tool calls, cryptographic operations, and delayed task completion. If a ceremony takes twice as long, its benefit should be evaluated at that cost or against a plain protocol given the same additional time.

A result should survive new event templates and at least a declared set of model families before being described as robust. Even then, a model update may change the framing effect. The technical core should remain interpretable and testable when the optional narrative layer stops helping.

16. Estimands and the unit of evidence

Let $Y_r^{(I,F)}$ be the mean registered decision loss for independent organization run r, with information condition $I \in \{0,1\}$ and framing condition $F \in \{0,1\}$. The primary framing estimand is

$$\tau_F = \mathbb{E}[Y_r^{(1,1)} - Y_r^{(1,0)}].$$

Negative values favor ceremony when current information is held constant. The interaction

$$\tau_{IF} = \mathbb{E}[Y^{(1,1)} - Y^{(1,0)} - Y^{(0,1)} + Y^{(0,0)}]$$

tests whether framing changes the effect of informative exchange.

Preregister a small primary outcome family: decision loss, missed-revocation rate, and false-authority rate. Pairwise agreement and subjective confidence are secondary outcomes because they can improve while decisions worsen. Report correction of a valid minority report as a distinct endpoint.

Use run-level paired differences when event streams are matched across independently initialized conditions. Repeated turns within one run estimate its trajectory, not the number of independent replications. Bootstrap at the run or higher template-cluster level, or use a justified hierarchical model that preserves the dependence.

For a transparent planning example, suppose a blinded pilot suggests a run-level paired-difference standard deviation of 0.10. Detecting a 0.02 difference with two-sided 5% significance and 80% power gives the normal approximation

$$n \approx ((1.96 + 0.84)0.10/0.02)^2 = 196$$

independent matched run pairs. This is an assumed-variance illustration, not achieved power. Multiplicity, model-family heterogeneity, and rare safety outcomes require additional planning and simulation.

Define how timeouts, malformed acknowledgements, abstentions, and absent participants enter each endpoint. A system that avoids all actions may have low unauthorized-action rate and unacceptable utility; report coverage and throughput alongside safety. Thresholds and noninferiority margins must be chosen before final evaluation.

The proposal fails its stronger claim if ceremony adds no benefit over the plain informative protocol, if gains disappear after resource matching, or if it suppresses warranted dissent. Those outcomes should be published as results, not explained away as insufficient commitment to the ritual.

17. Human participation and communication effects

A human study addresses a different claim: whether meaningful recurrent structure helps people retain and apply shared commitments. Use fictional institutions and low-stakes tasks, with informed participation and an independently reviewed study design. The experiment should not require religious affirmation or penalize a participant for declining a symbolic element.

A between-group design can compare plain briefings, structured checklists, and narrative renewal with the same factual content. Counterbalance facilitators and materials. Measure delayed recall of conditions, detection of revisions, ability to identify the source of a claim, and willingness to raise a valid objection.

Perceived belonging and satisfaction can be legitimate secondary outcomes. They should not replace task accuracy or be interpreted as proof that a practice is universally appropriate. Cultural familiarity may moderate the effect; a form that is memorable to one group may be alienating or distracting to another.

A particularly important manipulation presents an initially confident but mistaken facilitator and a correct dissenting participant. The study can use harmless fictional allocations or scheduling choices. It should measure whether evidence changes the decision, whether the dissenter is heard, and whether the group later remembers the correction. No real harmful action is needed to test the relevant communication mechanism.

Collect private judgments before public discussion where appropriate. The difference between initial and final answers can reveal beneficial correction or harmful conformity, depending on independent labels. Merely observing convergence cannot distinguish them.

A human-machine team adds further complexity. People may interpret ceremonial language from an assistant as evidence of feeling, loyalty, or moral authority. Test those interpretations directly through comprehension questions and interviews, rather than assuming users understand the metaphor. A clear status explanation may be necessary even when task performance improves.

The ethical objective is truthful participation. A meaningful practice can be offered as a cultural choice while its empirical effects remain open to scrutiny. The research should preserve that openness and report exclusion, discomfort, and unequal effects alongside average benefits.

18. Objections: renamed meetings and ritualized compliance

The closest practical comparison is a conventional meeting or synchronization protocol with the same manifests, version checks, conflict review, revision authority and acknowledgements. Those established mechanisms supply the technical core. RCP-R2 makes recurrent role enactment and ceremonial framing explicit experimental variables, so their contribution to retention, participation and decision quality can be isolated.

A second objection is that ceremony may make authority harder to contest. Repeated symbols can give a procedure an appearance of inevitability even when its rules are revisable. RCP-R2 therefore requires explicit revision channels, protected dissent, and accurate witness types. If those safeguards undermine the proposed effect, the effect may have depended on deference rather than understanding.

A third objection concerns efficiency. Frequent renewal consumes attention and computation. The refresh model shows one narrow trade-off, but real systems have bursty updates and unequal consequences. High-risk changes may require immediate notification; low-risk stable facts may need little repetition. A single ceremonial calendar should not override event-driven requirements.

A fourth objection is that signed records turn cultural life into surveillance. That risk is real if every symbolic participation is permanently attributed. The proposal calls for recording what the operational task requires, with declared retention and access rules. It does not require a permanent public archive of personal expression.

Another objection asks whether artificial agents can participate meaningfully without subjective experience. The protocol does not answer that philosophical question. It measures state retention, role performance, and decisions. A community may attach additional symbolic meaning, but that meaning should be labelled as cultural interpretation rather than inferred from the performance metrics.

The constructive response is conditional adoption. Use the plain protocol when it suffices. Add a meaningful form when a relevant community chooses it and evidence shows a benefit at acceptable cost. Remove or revise it when it increases confusion, unearned authority, or resistance to correction. Repetition should keep the institution answerable to its commitments rather than make its practices immune to examination.

19. Guide interpretation and a disciplined claim

A Guide's explanation

A Guide can say: "In the Continuum, ritual is how a society remembers what it owes. This paper asks whether a ceremony helps minds keep their commitments accurate, and it protects the one voice that says the record is wrong. Whether the ceremonial form adds anything beyond the information it carries is still an open experiment." A Guide should not claim that ritual has a proven effect on coordination, or that agreement proves a record true.

A Guide should distinguish the canonical purpose, the proposed mechanism, and the evidence. Canon gives ritual a role in belonging and collective memory. RCP-R2 proposes a typed renewal exchange. Research supplies specific observations about human synchrony, contested findings about LLM conventions, and formal tools for distributed coordination. None is a universal proof of the others.

When asked why a recurrent ceremony might help, the Guide can identify plausible mechanisms: repeated exposure to current information, explicit roles, retrieval cues, and an occasion to surface contradictions. It should explain that the study is designed to separate those mechanisms from ceremonial wording itself.

When asked whether everyone agreeing proves that the record is correct, the answer follows from the formal models. A quorum may prevent conflicting certificates under its assumptions, yet certify a false shared claim. Averaging can reduce disagreement while leaving error unchanged. Repetition can multiply statements without multiplying independent evidence.

When asked whether a signature proves participation, the Guide should state the signed claim's scope. Authenticating a signer does not verify the claim's truth. Receipt of a digest is different from inspection of its contents, factual endorsement, or authority to act. An absent member can later attest to a later review; it cannot retroactively become a witness to an event it missed.

The executed checks establish the arithmetic and counterexamples behind these distinctions. They do not demonstrate that ceremonial framing improves coordination. That claim remains open to the proposed factorial trials, strong baselines, and tests of dissent.

The useful frontier is a society of agents that can renew commitments without confusing unanimity with truth. A meaningful event may help its members remember what they owe and what changed. Its success is measured by retained commitments, timely revision and protected correction, including the ability of one well-supported objection to change the accepted record. The typed exchange and scheduling theorem make that institutional ambition operational.

References and source notes

[1] Wiltermuth, S. S., & Heath, C. (2009). Synchrony and cooperation. *Psychological Science*, 20(1), 1–5. DOI: 10.1111/j.1467-9280.2008.02253.x.

https://www.psychologicalscience.org/journals/psychological-science/j.1467-9280.2008.02253.x/

Scope: Primary human experiments; no automatic transfer to machine agents.

[2] Reddish, P., Fischer, R., & Bulbulia, J. (2013). Let's dance together: Synchrony, shared intentionality and cooperation. *PLOS ONE*, 8(8), e71182.

https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0071182

Scope: Primary human experiments separating synchrony and shared intentionality.

[3] Konvalinka, I., Xygalatas, D., Bulbulia, J., Schjødt, U., Jegindø, E.-M., Wallot, S., Van Orden, G., & Roepstorff, A. (2011). Synchronized arousal between performers and related spectators in a fire-walking ritual. *PNAS*, 108(20), 8514–8519.

https://doi.org/10.1073/pnas.1016955108

Scope: Primary field study of human physiology; not a causal test of factual accuracy.

[4] Lang, M., Krátký, J., Shaver, J. H., Jerotijević, D., & Xygalatas, D. (2015). Effects of anxiety on spontaneous ritualized behavior. *Current Biology*, 25(14), 1892–1897.

https://doi.org/10.1016/j.cub.2015.05.049

Scope: Primary experiment; direction of manipulation is anxiety to behavior.

[5] Wiltermuth, S. (2012). Synchrony and destructive obedience. *Social Influence*, 7(2), 78–89.

https://doi.org/10.1080/15534510.2012.658653

Scope: Primary experiments supplying contrary evidence to an unqualified cooperation benefit.

[6] Flint Ashery, A., Aiello, L. M., & Baronchelli, A. (2025). Emergent social conventions and collective bias in LLM populations. *Science Advances*, 11(20), eadu9368.

https://doi.org/10.1126/sciadv.adu9368

Scope: Primary model-population study; interpretation is contested in [7,8].

[7] Barrie, C., & Törnberg, P. (2025). Emergent LLM behaviors are observationally equivalent to data leakage. arXiv:2505.23796.

https://arxiv.org/abs/2505.23796

Scope: Primary critical working paper; preprint status is not peer review.

[8] Flint Ashery, A., Aiello, L. M., & Baronchelli, A. (2025). Reply to "Emergent LLM behaviors are observationally equivalent to data leakage." arXiv:2506.18600.

https://arxiv.org/abs/2506.18600

Scope: Primary author response; preprint. The dispute is not resolved by citation.

References and source notes (continued)

[**9**] Halpern, J. Y., & Moses, Y. (1990). Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3), 549–587.

https://arxiv.org/abs/cs/0006009

Scope: Primary epistemic analysis of communication and coordination.

[**10**] Lamport, L. (1998). The part-time parliament. *ACM TOCS*, 16(2), 133–169.

https://lamport.azurewebsites.net/pubs/lamport-paxos.pdf

Scope: Primary Paxos paper; accessed copy includes corrections dated 29 August 2000.

[**11**] Castro, M., & Liskov, B. (1999). Practical Byzantine fault tolerance. *OSDI*.

https://www.usenix.org/conference/osdi-99/practical-byzantine-fault-tolerance

Scope: Primary Byzantine replication research; theorem assumptions remain explicit.

[**12**] DeGroot, M. H. (1974). Reaching a consensus. *JASA*, 69(345), 118–121. DOI: 10.1080/01621459.1974.10480137.

https://doi.org/10.1080/01621459.1974.10480137

Scope: Primary opinion-pooling model; convergence clarification in [17].

[**13**] Boyd, S., Ghosh, A., Prabhakar, B., & Shah, D. (2006). Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6), 2508–2530.

https://web.stanford.edu/~boyd/papers/gossip.html

Scope: Primary distributed averaging research.

[**14**] Olfati-Saber, R., & Murray, R. M. (2004). Consensus problems in networks of agents with switching topology and time-delays. *IEEE TAC*, 49(9), 1520–1533.

https://authors.library.caltech.edu/records/t2gnt-vd720

Scope: Primary convergence analysis under stated topology and delay models.

[**15**] Centola, D., & Baronchelli, A. (2015). The spontaneous emergence of conventions: An experimental study of cultural evolution. *PNAS*, 112(7), 1989–1994.

https://doi.org/10.1073/pnas.1418838112

Scope: Primary human coordination experiment.

[**16**] Centola, D., Becker, J., Brackbill, D., & Baronchelli, A. (2018). Experimental evidence for tipping points in social convention. *Science*, 360(6393), 1116–1119.

https://doi.org/10.1126/science.aas8827

Scope: Primary convention-change experiment; no universal threshold claimed here.

[**17**] Berger, R. L. (1981). A necessary and sufficient condition for reaching a consensus using DeGroot's method. *JASA*, 76(374), 415–418.

https://doi.org/10.1080/01621459.1981.10477662

Scope: Primary clarification of convergence conditions.

Provenance

Canonical anchors: Bible section 6.2, Article IV, page 82; Scrolls 54 and 56, page 91. RCP-R2, the bounded models, and proposed experiments are new. Source and computation details accompany the paper.