A Group’s Success Does Not Tell Us Who Has a Say

Human authority in shared decisions, roles, and goals

Paul Cooper

17 November 2026

Abstract

A group can become better at carrying out a plan while giving its members less opportunity to correct that plan or question their part in it. A convincing summary may misrepresent someone’s concern; a proposed task may become an apparent commitment. Collective success alone does not tell us whether participants can challenge these changes.

Responding to Michael Levin’s account of how minds become embodied, this paper asks how people can help shape the arrangements through which a group acts. It connects roles people have a say in defining, records that preserve unresolved objections, and commitments with explicit limits to opportunities for revising roles and reconsidering shared goals.

Two claims must remain distinct. Participants should retain a meaningful say in their involvement. Separately, protecting correction and revision may help groups retain relevant information and adapt when circumstances change. This second claim is a hypothesis to test, not an established benefit.

An elementary mathematical example shows that two groups can achieve identical accuracy from truthful reports, exceeding what isolated participants can achieve, while only one lets participants inspect and correct their recorded reports before commitment. That accuracy measure cannot identify this particular protection. A separate model shows how explicit rules can keep exploration from changing permissions or commitments. This conclusion depends on those rules; it does not establish voluntary consent or reliable enforcement in practice. A geometric construction illustrates a change of display that leaves the participation record untouched. It does not show that minds have that geometric structure.

Proposed studies would test correction procedures, role revision, goal reconsideration, and the educational value of geometry separately. They would measure both group performance and participants’ practical opportunities to act, including the costs of using the procedures. The paper reports no new empirical study or validated deployment.

Introduction

Imagine a team planning a community event. One person raises an accessibility concern. Another says they might be able to help on Saturday. An AI assistant turns the discussion into a clear plan, but the concern is reduced to a preference and the tentative offer becomes an assigned task. The event might still run smoothly. That success would not show that the plan represented people accurately, or that they had agreed to the responsibilities it assigned.

The question is what those people can do next. Can they inspect the record, correct how their contributions are described, and have a disputed assignment reconsidered before it is treated as an agreed commitment? Can they question the goal when new information suggests it no longer serves the event’s purpose? These are

questions about participant agency: the practical opportunities people have to understand and help shape their participation. Having a say does not mean winning every disagreement, erasing earlier commitments, or having a veto over everyone else. It means having a usable way to contribute judgment, challenge misrepresentation, and help shape the terms of their involvement.

This question connects to Michael Levin's *Ingressing Minds*, which examines how larger systems can develop abilities and pursue goals that their individual parts cannot achieve alone. Levin also proposes that physical systems can give expression to patterns, including possible minds, that exist in a structured nonphysical realm. My response concerns a more specific problem: when people help make a group capable, what say do they have in changing how it is organized? The rules for assigning work, recording agreement, and responding to objections help determine how the group acts. I argue that these rules belong within the study of its capabilities. This argument does not require treating the group as a single conscious mind or settling Levin's proposal about where minds originate.

By a group's constitution, I mean the rules and practices governing who can decide, what counts as an agreement, and how decisions and roles can be changed. These arrangements include written rules, software permissions, access to records, everyday habits, and the actions of people who interpret and enforce them. A stated right to object can be difficult to use if someone cannot find the relevant information, cannot afford the delay, or faces retaliation for speaking up. Understanding a group therefore requires looking at how its rules work in practice: whose knowledge reaches a decision, how an objection is handled, and whether a correction can change what happens next.

My ethical position is that people should retain a meaningful say in the terms of their participation, whether or not this makes a group more productive. The testable claim concerns specified tasks in which people hold different relevant information and circumstances change. In those settings, I predict that usable procedures for correcting records and revising roles will, on average, help groups adapt better than otherwise comparable arrangements without those added procedures. The proposed explanation is that corrections preserve information the group needs and role changes let it reorganize its work. This prediction could fail, and any benefit must be weighed against the time, effort, and other costs of the procedures. Evidence of no meaningful benefit or worse performance would count against the prediction in the setting studied; it would not, by itself, settle whether people should have a say.

What this paper contributes

This paper brings established questions about participation and control into a response to Levin's account of collective capability. Its contribution is a specific proposal for shared work: connect what a person contributes, how that contribution is recorded, what they authorize, and how they can seek changes to their role or the group's goal. These connections matter when a summary becomes the basis for assigning work. A summary may accurately describe someone's skills and still assign work they never accepted. Checking accuracy and obtaining permission are different requirements.

The proposal distinguishes three questions to assess alongside task performance. Was a person's contribution represented accurately? Did they authorize the resulting commitment? Can they use an agreed procedure to question or revise the arrangement? Challenging an interpretation does not make a participant the final judge of every factual dispute. Procedures must also specify how disagreements are considered and who may decide what happens next. The related work below supplies important foundations for these concerns.

Three mathematical examples make limited distinctions precise: a group's accuracy does not reveal whether people can correct their recorded reports; exploring a proposal need not change permissions; and changing a diagram need not change what anyone agreed to. Each conclusion depends on its stated example or rules.

Proposed studies then compare specific correction, role-revision, and goal-review procedures, measuring group performance, participants' opportunities to act, and the burden of using the procedures separately. These studies would test whether the procedures improve adaptation in the specified tasks and measure any gains alongside the time and effort they require.

A person is more than a role

During the 2015–2017 drafting and adoption period of the MetaCanon Constitution, my collaborators and I were addressing a practical problem [3]. A shared vision can bring people together, but giving someone responsibilities does not necessarily give them a say in defining those responsibilities.

One response was the Perspective Lens: an agreed description connecting a person's contribution to a shared purpose, areas of responsibility, and duties. Its holder helps define it and can contribute through several lenses. The person remains more than any particular role. The Enklu presentation draft from 2019–2020 also describes participants defining their own lenses and contributing through more than one [17, items 1 and 6].

Helping shape one's role does not mean rewriting another person's responsibilities or silently cancelling arrangements on which others depend. Both the holder and their collaborators need ways to question and renegotiate the relationship. Changing a shared goal may require a different agreed process from changing an individual role.

The historical comparison identifies related ideas in earlier drafts and later conversations [2]. My account here is a participant's retrospective account. It does not independently authenticate every recovered edition, signature, attribution, or adoption claim, and this argument does not depend on publication priority.

Geometry helped me picture distinct contributions oriented toward a shared purpose. It did not tell us who should decide a disagreement. When I describe a constitution as protecting emergence, I mean leaving room for participants to shape the organization's development. I am not explaining the origin of every possible pattern. A group can develop memory, specialization, and routines beyond anyone's initial design; the question is whether its contributors can still understand and question their place within it.

Capability, experience, and authority are different questions

Levin already considers different agencies cooperating and competing within a larger system, and explicitly addresses ethics [1, §§6.1, 6.5, 7.4]. Levin also asks how memories might reshape the systems that support them [1, §7.3, p. 36]. I take up a related question: how can a group's contributors help revise the organization through which it acts? A person who interrupts a group may be mistaken, may know something it has missed, or may never have agreed to the requested task. These possibilities require different responses.

Five questions should remain separate: what a system can accomplish; what participants can understand and change; whether the system has subjective experience; what moral consideration it deserves; and who may authorize an action. Evidence about one can matter to another without settling it. A founder's insight does not confer unlimited authority. A human approval step can also conceal concentrated power when others cannot inspect the evidence or safely disagree.

A Third Edition constitution states that AI agents do not have emotional or relational experiences; the copy cited here leaves its adoption date and signatories unspecified [18, §2.1.2 and adoption declaration]. That is a claim about experience, not merely an allocation of responsibilities. For this research proposal, I would separate it from rules assigning authority. In the systems I am developing, humans attest and agents do not ratify: people confirm decisions, while AI agents cannot make them binding. This allocation does not settle

whether artificial experience is possible. Relevant evidence about moral standing would still merit consideration, without automatically granting permission to use someone else's resources. This paper enacts no constitutional amendment.

Levin describes his sorting algorithms as deterministic, yet connects their unexpected capabilities to behavior that is not fully determined by the physical process [1, §6.3, pp. 29–31]. An outcome that was not an explicit objective can still follow from the full set of rules, starting conditions, and update order. Being unexpected by the designer does not, by itself, establish that the outcome was physically undetermined. My objection concerns that step in the argument; it does not dispute the reported observations or independently reproduce the computation.

Likewise, useful discoveries can justify pursuing a theory without showing that its account of reality is uniquely correct. Levin distinguishes a research program's usefulness from deductive proof of his proposed nonphysical account [1, §7.3, p. 37]. Successful changes to group procedures would first support claims about those procedures.

How a description becomes consequential

In the event example, describing someone as available can become an assignment; describing a concern as resistance can keep it out of discussion. Observing a contribution, interpreting it, and obtaining permission to act are different operations. My proposal gives contributors a way to question the transitions between them. This is a chosen constitutional commitment, not a right derived from Levin's discussion of observers and patterns [1, §7.2].

Correction may also concern the goal. Maximizing attendance could undermine accessible participation. A review should consider whether the target remains appropriate; retaining it can be justified. The number of goals changed is not a measure of success. Participants also need a say in how opportunities for correction are organized, rather than having the system decide what counts as sufficient agency for them.

Related work

The practical ability to influence shared decisions has a substantial history. Ostrom describes institutions in which people affected by operational rules can help change them and have access to affordable ways of resolving conflict [10, pp. 90–93, 100–101]. Leventhal connects fair procedures with accurate records, opportunities for correction, and representation of affected people's concerns [11, manuscript §§III.C.3–6]. Frey and Schneider distinguish expressing a view from having a voice that can bring about binding changes through transparent procedures. Their proposals include accountable authority and ways for participants to change community governance [12, pp. 2388–2393]. These works already connect several elements of the present proposal.

Research on human–AI interaction supplies close precedents for correcting records. Hirsch and colleagues describe ways to correct automated transcripts, add context to assessments, and record disagreements, while recognizing the work these demands place on users [13, pp. 96–98]. Alfrink and colleagues place opportunities to challenge and revise AI systems throughout their development and use [14, §§6.1–6.2]. Delgado and colleagues distinguish participation justified by people's standing from participation intended to improve a system. They also consider people's influence over how participation itself is organized [15, §2.4]. Correctable records and continuing influence over a system therefore have established foundations.

Zhou and colleagues address a particularly close problem. Their abstract describes how improved AI decision support can coexist with fewer practical opportunities to question representations, revise decision authority, or

change routines before recommendations become commitments [16, abstract]. The comparison here is limited to that published abstract. It establishes a close connection, but does not establish how the complete frameworks differ. The present paper develops specified procedures and proposed comparisons without claiming priority for the broader concern.

Other research helps distinguish the questions involved. De Jaegher and Di Paolo describe how people make sense of situations through interactions that can develop their own organization [5]. That account does not determine who should have authority in an institution. Santoni de Sio and van den Hoven's account of meaningful human control distinguishes a system's responsiveness to relevant human moral reasons from connections to human understanding and responsibility [8]. The present proposal makes an additional ethical commitment to contributors having a meaningful say in their involvement.

Evidence about group learning also helps frame the proposed tests. Edmondson found an association between psychological safety and team learning [6]. Lorenz and colleagues found that social influence could reduce the diversity of estimates and increase confidence without a corresponding improvement in collective accuracy in their numerical estimation task [7]. These findings concern their particular settings; they do not establish that the arrangements proposed here improve adaptation. The proposed studies would assess those arrangements directly, including their costs and possible failures.

Why a good score cannot answer the agency question

Suppose two people each see the outcome of one independent fair coin toss. Their shared task is to say whether the outcomes differ. Either person working alone has only a one-in-two chance of being right. A system receiving both truthful reports can answer correctly every time.

Now compare two arrangements. In one, each person can inspect and correct the stored account of their report before the answer becomes final. In the other, that opportunity is absent. On the initial task, with truthful reports stored faithfully, both give exactly the same correct answers. The score shows the benefit of combining information, but does not reveal whether this particular protection exists.

This is a deliberately narrow mathematical example. It does not show that correction procedures always improve performance, nor does it measure every part of agency. A task involving recording mistakes could distinguish the arrangements. Appendix A gives the definitions and proof.

Exploring, committing, and reconsidering

The Hopf Workshop's torus-and-cones framework distinguishes three activities [4]. A cone represents generating possibilities. A torus—a ring-shaped surface—represents arrangements being maintained. A second cone represents examining and, where appropriate, releasing selected arrangements. These are images for a process, not proofs of its effectiveness.

Exploring a role does not mean accepting it. Maintaining a commitment does not mean knowing its outcome: someone can agree to conduct an experiment without promising that it will succeed. Reconsideration can address one assumption or duty without dismantling the whole project. Evidence can change a factual judgment; assigning new work requires the appropriate agreement.

WeOpoly provides a related coordination model [3]. Participants organize contributions around a goal, identify conversations and decisions, carry out bounded work, and return with what happened. The point is the connection between a shared account and subsequent activity.

A usable record would distinguish proposals and their versions, attributed statements, permissions, commitments, unresolved objections, and completed actions. Permission should specify who gave it, what version and action it covers, its limits, and how it ends. A summary can link these records but cannot replace them. Agreeing to design an invitation does not authorize a later assignment to manage attendance.

Delegation can avoid a new meeting for every action. Its limits on time, resources, or repeated use must be checked before an action is sent for execution. Proposing an extension does not extend permission. Revocation blocks later requests under that permission; stopping work already sent requires a separate mechanism and may not reverse its consequences. Records also need privacy protections and retention limits; preserving relevant history does not require keeping every personal statement forever.

The workshop uses “nonattachment” to mean allowing inquiry to challenge a preferred conclusion [4]. That expectation must include founders and interpreters of the group’s purpose. Someone who listens and still disagrees retains a place in the process. Rest, refusal, maintenance, and deferral are also possible responses; personal transformation is not a condition of participation.

In the voluntary setting proposed here, ending future involvement does not depend on first completing reconciliation. Costs and dependencies can still require attention. This does not erase earlier commitments or claim that every obligation in every institution can be ended identically. A formal option to leave may be unusable where income, belonging, or necessities depend on participation. Software alone cannot resolve those conditions.

Why exploring an idea need not create a commitment

A proposed system can distinguish discussion from agreement. Its record can contain candidate proposals, permissions, existing commitments, and an event history. If the rules for exploring or proposing something never change recorded grants or commitments, then any finite sequence of those steps leaves those records unchanged. A separate authorized action is required to change those records. A recorded grant can still become ineligible for use when its stated expiry time passes.

This conclusion depends on more than labels in an interface. Existing permissions must refer to fixed terms, so editing a proposal cannot silently change what was agreed. The component applying events must follow the stated rules, and there must be no route for actions to bypass the permission checks. Appendix B states these assumptions and proves the resulting properties. The proof does not establish that a real system meets the assumptions or that an agreement was freely given.

What geometry can and cannot contribute

A picture can explain a relationship; a communication network can determine who exchanges information; a mathematical construction can have provable properties; and a spiritual image can express meaning. Success in one use does not establish the others.

The Hopf fibration is a mathematical mapping in which certain different representations give the same mapped point [9]. The workshop extends that illustration by keeping records of participation separate from the geometry [4]. A display change leaves permissions untouched because it is defined not to change them. Correcting a statement or withdrawing permission is a different operation. This does not make people’s disagreements equivalent or disposable.

The cone imagery also has limits. Joining cone surfaces along their rims differs from joining them at their tips. My cones represent exploration and reconsideration; Levin’s “cognitive light cones” represent the reach of

goals, memory, and anticipation across space and time [1, Figure 13]. Similar shapes do not establish identical structures. Returning to the same picture also does not erase the history of what happened.

Images of shared light, vessels, and embodied integration helped motivate my inquiry [2]. Their spiritual meaning does not establish that Christian light imagery, the Hopf fibration, and Levin's nonphysical pattern space describe the same thing. The stronger claim that consciousness or governance literally has Hopf topology remains a separate, unresolved conjecture [4]. The present argument requires neither that claim nor a shared spiritual interpretation. Geometry must earn its educational usefulness against an equally clear account without it.

How the proposal could be tested

The practical prediction concerns tasks in which useful information is spread across people and circumstances can change. The proposed benefit is better adaptation: retaining information that corrects a plan and changing responsibilities when the plan no longer fits. Whether a procedure actually enforces its permissions, whether people can use it, and whether a geometric explanation helps them understand it are separate questions. One successful study would not settle all of them.

The WeOpoly cycle described above offers a possible setting for the studies. The question is how the connection between shared plans, action, and subsequent review works in practice. Arranging a diagram or completing a meeting would not, by itself, establish success [3].

Compare specific procedures

A first study could use adult teams planning and carrying out a simulated community event. Different participants would receive different relevant information. Once a plan exists, a change specified in advance would make it necessary to adapt: a resource becomes unavailable, a requirement changes, or an assigned role is no longer workable. All participants would retain basic protections and the right to leave the research study.

The first comparison would concern a correction procedure. Both groups would have the same initial information, access to sources, limits on authority, facilitation, task time, and access to any AI assistant. The comparison group would have a competent facilitator and an ordinary way to raise concerns. The other group would also have a shared correction channel that links statements to their sources and keeps an unresolved objection visible until its outcome is recorded. Training time and interface clarity should be comparable. This would test the added procedure as a package; it would not identify the separate effects of source links, visibility, and persistence.

A second study could hold the correction procedure fixed and compare two ways to revise roles. One would provide scheduled opportunities to renegotiate. The other would let participants initiate proposals for a change, with attention to collaborators who depend on their work. Both would allow people to question assignments and decline new involvement. The comparison concerns how revision is supported; it does not remove basic standing from the comparison group.

A further study could examine reconsidering a shared goal. Sometimes new information would undermine the goal itself; at other times, a smaller adjustment would be enough. A specified procedure for triggering review would be compared with ordinary access to discussion, keeping the available evidence and decision authority equivalent. A useful procedure should help distinguish a justified change of direction from unnecessary reopening of a workable plan.

These comparisons would help identify which procedures make a difference. Later studies could examine whether the procedures strengthen or undermine one another when combined.

Measure performance and participation separately

The study should assess the quality of the resulting plan, recovery after the change, and the time and resources used. Assessment criteria should be set in advance. Where possible, evaluators should not know which procedure a team used. A system's own record that it followed its rules would not be an independent measure of success.

Participants' practical opportunities need their own measures. Can they identify what they actually agreed to? Can they detect and challenge a misrepresentation? Can they use the correction procedure and find out what happened to an objection? Private reports can supplement observations, especially where apparent agreement conceals reluctance. A useful opportunity to challenge a record is different from a guarantee that one's preferred account will prevail.

For a goal-review study, independent assessors would judge the response to new information using criteria fixed before teams are assigned to conditions. The record should show who initiated review, whether contrary reasons were considered, and why the goal was retained or changed. Learning the researcher's preferred vocabulary would not count as success. A better-written explanation would not establish that relevant evidence was heard.

Permission errors should be checked against the scenario's independent record of actual permissions. An AI-generated summary cannot serve as both the account being tested and the answer key. The study should also measure the burden of participation: who spends time correcting records, whose objections wait longest, and whether people with less status can use the same procedures.

A team average can hide a serious loss of agency for a particular person. Where the data permit, results should therefore be reported by participant and role as well as by team. Tradeoffs should remain visible. Neither a single average nor the lowest individual score automatically settles what a group should value.

Leaving the real research study must not count as task failure or become a reason to pressure someone to stay. Simulated changes in availability can test how a team handles disrupted responsibilities without compromising that protection.

Decide in advance what the evidence would mean

Before a study intended to test the prediction begins, its plan should specify the main comparison, what size of difference would matter, how many teams are needed and why, which observations may be excluded, and how the results will be analyzed. People on the same team and teams completing repeated tasks do not provide wholly independent observations; the analysis must account for those relationships. A preliminary study can check feasibility and estimate variation without being presented as proof that the procedure works.

Slower recovery, unnecessary reopening of commitments, unequal correction burdens, lost minority information, or confusion between considering and agreeing would count against the relevant practical claims. Evidence that a simpler procedure works equally well with less burden would support using the simpler procedure. An inconclusive result should remain inconclusive. Calling participants insufficiently open would not explain away a procedure's failure.

No participants have been enrolled in these proposed studies, and this paper reports no results from them.

Test the explanation separately from the procedure

To test the educational value of geometry, the coordination rules should stay the same. One group could learn them through a clear plain-language map; another could use the lens, torus, and cone imagery. Both should

receive comparable practice, explanation, and time. The question is whether the geometric version helps people understand commitments accurately and apply what they learned to a new situation. Enjoying the pictures is a different outcome.

A separate experiment with software agents could change which agents can communicate with which others, while keeping the models, information, memory access, and computational resources comparable. That would test the effect of actual communication connections. A picture of a polyhedron does not establish those effects. Any benefit should be attributed to what the comparison identifies, such as better access to information or a more effective review procedure.

Objections and limits

My ethical commitments are chosen, not deduced from biology. Their effects on information and decisions can nevertheless be investigated. More opportunities to renegotiate may slow action or prevent completion; preserving participation does not mean maximizing every immediate choice.

A group may have interests deserving consideration, but describing it as a higher-level agent does not automatically cancel its participants' claims. Internal agreement also cannot settle consequences for people outside the group.

Objection procedures can obstruct work, and appeals to shared purpose can suppress criticism. Both require explicit arrangements for evidence, review, and decisions. Neither founders nor dissenters become infallible. Participants can challenge a record without acquiring the right to erase contrary evidence.

Records can contain falsehoods, permissions can be coerced, and actions can bypass software. The Hopf Workshop is a bounded research and design artifact; it does not establish that the proposed protections work in an operating institution [4].

Conclusion

I propose studying how people can help shape the groups their contributions make possible. Roles, records, permissions, and opportunities for revision help organize collective action. Their usefulness, costs, and failures require evidence. A group's growing capabilities and its participants' practical opportunities to influence their involvement belong in the same inquiry.

Appendix A. Better answers do not establish an opportunity to correct the record

Consider two people who each see one independent fair coin toss. Their shared task is to say whether the two outcomes differ. Either person, working alone, has only half the required information. Combining their reports solves the task, but this success says nothing about whether they can later correct what the system attributes to them.

Write the outcomes as bits $B_1, B_2 \in \{0, 1\}$, each equally likely to be zero or one, and independent. The target is their parity,

$$Y = B_1 \oplus B_2,$$

where $\oplus$ gives one when its inputs differ and zero when they agree. Participant i observes only B_i, receives no communication from the other participant, and may use private randomness R_i independent of the pair (B_1, B_2). Any answer they give is therefore a function $f_i(B_i, R_i)$.

Isolated-information result. Every such rule that returns a bit has success probability exactly $1/2$.

Proof. Fix the observed bit and the participant's private random information. The other bit is still a fair coin, because it is independent of both. Of its two equally likely values, exactly one makes the participant's answer correct. Thus the conditional probability of success is $1/2$. Averaging over the observation and private randomness leaves that probability unchanged. Random guessing, a complicated calculation, and a fixed answer all have the same accuracy under these assumptions. ▫

Now define two protocols. Both accept one truthful report $r_i = B_i$ from each participant, store it faithfully, and calculate $\hat{Y} = r_1 \oplus r_2$. Both have a specified commitment point after which that trial's answer is final.

In the first protocol, before commitment, each participant can authenticate, inspect their own stored report, and invoke $correct(i, b)$. The operation replaces only their report with $b \in \{0, 1\}$, records the change, and recomputes $\hat{Y}$ before commitment can proceed. For this example, authentication reliably restricts the operation to that participant. A truthful correction sets $b = B_i$. In the second protocol, there is no operation for inspecting or changing the stored attribution after submission. Its initial calculation is otherwise identical.

On the benchmark with truthful reports and faithful storage, any corrections are truthful and leave the report values unchanged. Both protocols therefore answer correctly with probability one, strictly exceeding the isolated-information baseline. Yet only the first supplies the specified opportunity to inspect and correct the attribution before commitment. Identical aggregate accuracy is compatible with either presence or absence of this narrow protection, called *report-contestability* here.

The conclusion concerns what the benchmark can establish. It does not rank the protocols on every task or provide a complete measure of agency. A benchmark introducing recording errors could distinguish them.

Appendix B. When exploration cannot create a commitment

A proposed arrangement can let people explore possibilities without treating exploration as agreement. The following model states the conditions under which that separation holds.

Represent the recorded situation by

$$q = (P, G, C, H).$$

Here P contains candidate proposals, G contains grants of permission and their recorded status, C contains active commitments, and H is the ordered event history. Call the set of such states Q. A grant or commitment binds a fixed version of its terms. Editing a proposal creates a different version; it cannot change an existing grant or commitment through a reference to editable text.

Assume the following transition rules. Exploration and proposal events may change P and append to H, but must leave G and C unchanged. Creating, changing, or ending a grant or commitment requires a distinct authenticated event authorized by an explicit policy. That policy remains fixed throughout the trace considered here. A discussion, recommendation, or proposed approval is not such an event.

A dispatch event requests that a specified action be sent for execution. It is admitted only if, at that point, an eligible grant covers the action and its exact version and satisfies all applicable limits, including any expiry or

remaining allowance. Revoking a grant makes that grant ineligible for subsequent dispatches; any replacement requires a new authorized grant. Checking permission and admitting dispatch form one atomic transition, so another event cannot intervene between them.

Finally, assume a trusted sequential reducer: the component applying events follows these rules, one event at a time, in a definite order. There is no path for changing the protected state or dispatching actions outside these rules. Every accepted event appends to the history without deleting or rewriting earlier entries.

Conditional result. Starting from any state and following these assumptions:

1. Every finite sequence containing only exploration and proposal events leaves G and C unchanged.

2. Every admitted dispatch passes the grant check in the state immediately before its admission.

3. Every earlier history is an initial segment of each later history.

Proof. For the first statement, a sequence of length zero changes nothing. If the statement holds after n events, the next exploration or proposal event preserves G and C by its rule. It therefore holds after $n+1$ events, and hence for every finite sequence. The second statement follows directly from the only rule that admits a dispatch. For the third, each accepted event changes the history from H to $H \Vert [\, e \,]$, meaning the old list followed by the new event. Repeated appending preserves every earlier list as an initial segment. ▫

These conclusions describe the stipulated model. They do not establish reliable authentication, voluntary assent, accurate clocks, legitimate policy, or faithful external execution. Revocation prevents later admissions under the revoked grant; it does not undo an action already admitted. An append-only record preserves what was recorded, including possible mistakes. It does not make that record true or require indefinite retention of personal information in an actual institution.

Appendix C. A change of representation that leaves a record fixed

The geometric example makes one limited distinction exact: some changes to a representation can leave its displayed result and its associated record unchanged. Use this explicit convention for the standard Hopf map [9]:

$$S^3 = \{ (z_1, z_2) \in \mathbb{C}^2 : |z_1|^2 + |z_2|^2 = 1 \},$$

$$\pi(z_1, z_2) = \left(2\,\mathrm{Re}(z_1 \overline{z_2}),\, 2\,\mathrm{Im}(z_1 \overline{z_2}),\, |z_1|^2 - |z_2|^2 \right).$$

Each complex number has two real coordinates, so the input has four real coordinates constrained to unit length. The bar denotes complex conjugation; Re and Im take the real and imaginary parts. The output has three real coordinates.

To check that the output lies on the ordinary unit sphere S^2, put $a = |z_1|^2$ and $b = |z_2|^2$. Its squared length is

$$\Vert \pi(z_1, z_2) \Vert^2 = 4\,|z_1 \overline{z_2}|^2 + (a - b)^2 = 4ab + (a - b)^2 = (a + b)^2 = 1.$$

Now multiply both inputs by the same complex number u of unit length. Their squared magnitudes remain unchanged, and

$$(u z_1)\overline{u z_2} = u \bar{u}\, z_1 \overline{z_2} = z_1 \overline{z_2}.$$

Consequently, $\pi(u z_1, u z_2) = \pi(z_1, z_2)$. This is common-phase invariance: changing both complex phases by the same amount does not change the output.

To include the semantic record, take pairs $(z, q) \in S^3 \times Q$, using Q from Appendix B. Define

$$g_u(z, q) = (uz, q), F(z, q) = (\pi(z), q).$$

For every such pair and every $|u| = 1$,

$$F(g_u(z, q)) = (\pi(uz), q) = (\pi(z), q) = F(z, q).$$

The record q, including its permissions, remains unchanged because the operation explicitly leaves it fixed. Geometry does not supply that protection on its own. Nor does this identity establish that human disagreements are irrelevant changes of phase, that objections may be discarded, or that consciousness has the topology of S^3. It illustrates a precisely defined change of representation; substantive corrections require different operations.

Research status and sources

This is a conceptual response with elementary mathematical arguments and proposed empirical studies. It reports no new human-participant experiment or validated production deployment. The arguments in the mathematical appendices are presented as ordinary proofs; they are not claimed to have been checked by a proof assistant. No independent human peer review is claimed for this draft.

The response addresses Levin's published article; it does not independently reproduce the biological or computational studies discussed there [1]. The related-work section compares selected sources, rather than offering an exhaustive history or establishing priority for this proposal. The account of participatory sense-making draws on the publisher's abstract and the author's account [5]. The comparison with Zhou and colleagues is limited to their published abstract [16].

The account of MetaCanon's 2015–2017 drafting and adoption period is my retrospective account [3]. The Enklu draft provides a separate documentary source for holder-defined Perspective Lenses and the use of geometry as analogy [17]. It does not establish the earlier adoption history. The cited Third Edition has blank adoption fields [18]. These documents should not be treated as evidence that every described arrangement was formally adopted or successfully implemented.

The comparison report, first-person response, and Hopf Workshop are unpublished sources [2–4]. The workshop's date and version identify a research and design package; its mathematical arguments concern defined constructions, and its software claims are not evidence of enforcement in a working institution. The source map identifies the inspected copies, relevant passages, and reading limits. Private access to these materials does not imply permission to redistribute all their underlying records.

AI systems assisted with drafting, mathematical analysis, source checking, and editorial review.

References

[1] Levin, Michael. 2026. "Ingressing Minds: Causal, Non-Physical Patterns In-Form Natural, Synthetic, and Hybrid Embodiments." *Philosophies* 11(5), article 161, 1–53. Published 9 September 2026. DOI: 10.3390/philosophies11050161. Publisher's full text.

[2] *Paul Cooper's Work in Conversation with Ingressing Minds*. 2026. Unpublished comparative research report, 11 pages. Individual authorship is unspecified. Sections 2–6 and 9–10 discuss documentary history, interpretations, and source limitations.

[3] Cooper, Paul. 2026. *The way a collective becomes capable and the way its participants retain agency belong in the same design inquiry*. Unpublished first-person response. Its account of the 2015–2017 period is retrospective.

[4] *The Hopf Workshop—Complete Handoff*, version 0.2, 13 September 2026. Unpublished AI-assisted research and design package prepared for Paul Cooper / Jedi Sherpa. Relevant documents: "One model, three languages—the unified construction," §§2–6; "Mathematical foundations and proofs," M12; and "The Living Workshop," §§1–5.

[5] De Jaegher, Hanne, and Ezequiel Di Paolo. 2007. "Participatory sense-making: An enactive approach to social cognition." *Phenomenology and the Cognitive Sciences* 6(4), 485–507. DOI: 10.1007/s11097-007-9076-9. Author's account and publication link.

[6] Edmondson, Amy. 1999. "Psychological Safety and Learning Behavior in Work Teams." *Administrative Science Quarterly* 44(2), 350–383. DOI: 10.2307/2666999. Full article hosted by MIT. The study's causal and generalization limits are discussed on pages 378–379.

[7] Lorenz, Jan, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. 2011. "How social influence can undermine the wisdom of crowd effect." *Proceedings of the National Academy of Sciences* 108(22), 9020–9025. DOI and publisher's full text: 10.1073/pnas.1008636108.

[8] Santoni de Sio, Filippo, and Jeroen van den Hoven. 2018. "Meaningful Human Control over Autonomous Systems: A Philosophical Account." *Frontiers in Robotics and AI* 5, article 15. DOI: 10.3389/frobt.2018.00015. Publisher's full text.

[9] Lyons, David W. 2003. "An Elementary Introduction to the Hopf Fibration." *Mathematics Magazine* 76(2), 87–98. DOI: 10.1080/0025570X.2003.11953158. Author-deposited manuscript, arXiv:2212.01642v1, deposited 3 December 2022, includes corrections and commentary concerning the 2003 publication.

[10] Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press. Chapter 3, especially pp. 90–93 and 100–101. DOI: 10.1017/CBO9780511807763. Original chapter scan hosted by the University of Washington.

[11] Leventhal, Gerald S. 1976. *What Should Be Done with Equity Theory? New Approaches to the Study of Fairness in Social Relationships*. Author manuscript, September 1976, ERIC ED142463, §§III.C.3–6, manuscript pp. 25–31. Inspected manuscript. A later chapter appeared in *Social Exchange* (1980), pp. 27–55, DOI: 10.1007/978-1-4613-3087-5_2; the citations here refer to the 1976 manuscript.

[12] Frey, Seth, and Nathan Schneider. 2023. “Effective voice: Beyond exit and affect in online communities.” *New Media & Society* 25(9), 2381–2398. First published online in 2021. DOI: 10.1177/14614448211044025. Author manuscript, arXiv:2009.12470v2. Page citations refer to the published journal article.

[13] Hirsch, Tad, Kritzia Merced, Shrikanth Narayanan, Zac E. Imel, and David C. Atkins. 2017. “Designing Contestability: Interaction Design, Machine Learning, and Mental Health.” *Proceedings of the 2017 Conference on Designing Interactive Systems*, 95–99. DOI: 10.1145/3064663.3064703. Author manuscript.

[14] Alfrink, Kars, Ianus Keller, Gerd Kortuem, and Neelke Doorn. 2023. “Contestable AI by Design: Towards a Framework.” *Minds and Machines* 33, 613–639. First published online in 2022. DOI and full text: 10.1007/s11023-022-09611-z.

[15] Delgado, Fernando, Stephen Yang, Michael Madaio, and Qian Yang. 2023. “The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice.” *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, article 37, 1–23. DOI: 10.1145/3617694.3623261. Inspected author manuscript, arXiv:2310.00907v1.

[16] Zhou, Weihua, Xin Lin, Lei Fu, Ying Tang, and Cangyu Jin. 2026. “Why human–AI collaboration fails: An operational contestability framework.” *ENGINEERING Management* 13(3), 796–808. DOI and publisher abstract: 10.1007/s42524-026-6049-7. The comparison in this paper is limited to the published abstract.

[17] *Draft Enklu AR Text*. 2019–2020. Unpublished presentation draft, items 1 and 6. The associated document metadata records creation on 23 December 2019 and modification on 2 February 2020; these dates do not establish individual authorship or constitutional adoption.

[18] *The Metacanon Constitution (Third Edition)*. Undated unpublished draft, §2.1.2 and “Constitution Adoption Declaration.” The cited copy leaves the adoption date and ratifier signatures unspecified.