

How History Repairs Orphaned Knowledge

Epistemic Commons and Protocols from Cave Art to Synthetic Media

[Response to Long Now Labs 001.2 — Epistemic Cycles](#)

Kerry Lynn · Roger Meike · Glenn Scott · Yunus Telli (Faculty Collaborator — WPI) · Diane Weiss

Premise

The call asks how earlier civilizations navigated epistemic disruptions — the breakdown of a society's perception of shared truth — and what patterns might strengthen our own commons. These disruptions come from new media, from deliberate actors, or from both at once. The commons is usually treated as a store of facts to be checked. We see it differently: as a collaborative achievement — bodies, tools, and institutions working in step to make shared meaning. Its foundation is **presence**—the link between a claim and an accountable author, an embodied human judgment, and a traceable chain of endorsement. Provenance records where something came from and who authored it; presence identifies the human chain of vouching behind it and keeps it open to evaluation.

Presence is what lets a community answer the only question that matters for trust - "who stands behind this, and how do we know?" When the link holds, people can unite around a claim. When it breaks, what remains is *orphaned knowledge*: a possibly convincing claim with no one identifiable or accountable behind it. An anonymous deepfake is the clear case: it can look and sound real, yet no one is answerable for it. The deeper problem is not that AI lies, but that it cuts the thread connecting what we see to anyone accountable for it. Repair, then, is not fact-checking but re-tying that thread — keeping claims connected to the people accountable for them as they travel.

Presence breaks in three ways. **Severance** cuts the link outright. **De-authorization** withdraws trust from the authority behind a still-clear source, because authorization rests on shared values, not on evidence alone. **Fragmentation** splits a commons into rival authorities with no arbiter between them. The present rupture is as much about de-authorization and fragmentation as it is about forgery.

Our age is not the first to face authority split among many sources, with no neutral arbiter between them. Communities have long held different beliefs while still sharing enough common ground to weigh each other's claims. Repair does not mean herding everyone back to one commons — difference is the normal condition, not the failure. What repair can offer is a shared way to test *how* a claim is vouched for, whoever makes it: does its chain of accountability hold up under scrutiny? A chain that survives hard testing, again and again, is the strongest signal of reliability history offers — not proof that everyone will agree, but common ground enough to argue across.

Instead of retelling the familiar story of one medium replacing the next — marks, writing, print, photography, networked media, generative AI — we map something else: the ways presence gets severed, and the ways it gets restored, across history. The record's most useful lesson is not that new media cause rupture but that repair is *labor* — done by communities who invent protocols that outlast the technologies that prompted them. The best of those protocols remain legible and adaptable; they are the raw material for a Long Now-scale theory of epistemic resilience.

We propose a typology that spans three eras: the far past, the recent past, and the present. In each, we trace how trust broke and how communities rebuilt it. From the patterns that emerge, we design at least three intervention recommendations — societal, technological, institutional, economic, or process-based — for the present rupture. The goal is a working tool, not just historical analysis. This is a cycle, not a one-time crisis: across human history, from the deep past into the deep future, communities establish presence, see it severed, and repair it — each return a turn of a spiral that builds. The present rupture is the open one — its repair still being built, which is where our work begins.

Theoretical Armature

Our master concept is **presencing**: the human habit of putting thoughts outside ourselves — into a form that lasts and travels beyond us. A cave painting, a written page, a photo, a posted video each holds a piece of us and carries it to people we will never meet. From Lascaux to alphabetic writing, print, the photograph, networked media, and generative AI, each medium presences the human more powerfully than the last.

But every act of presencing is *double*: it carries our thoughts to others, and it cuts those thoughts loose from us. The moment a thought leaves our hands, it can travel without us — and without anyone to answer for it. The mark outlives the maker. Plato raised the worry in the *Phaedrus*, where the written word is logos, "orphaned" from a parent who can no longer defend it; André Leroi-Gourhan traces the same externalizing of memory, and Bernard Stiegler names it the *pharmakon* — gift and poison in one act. Presencing swings between **extension**, the human amplified, and **severance**, the form detached from the embodied, accountable person.

Externalizing also gathers the crowd. A shared record draws people into communities of common interest, where they listen most to those who already think as they do. That clustering can harden into **fragmentation** — a closed loop that hears only itself, or it can become a working community that builds and maintains the protocols a commons needs to repair itself.

This tension is not a defect to be solved; it is the condition of all cumulative culture, and the standing possibility of rupture. Epistemic history moves in a **spiral, not a wheel**. Each more powerful medium opens a gap between a claim and the presence accountable for it, faster than communities can close it. Communities then invent protocols that re-anchor claims; a deeper medium reopens the gap. No repair returns the commons to a prior state — repair is (re-)construction, not restoration. Steve Fuller's *Post-Truth: Knowledge as a Power Game* (2018) follows the same logic; distrust of truth is endemic to human civilization from the Platonic dialogues forward — not an aberration to be repaired back to a baseline that never existed.

Stephen Cowley's synthesis of pragmatist philosophy, distributed cognition, and ecolinguistics ("How Expertise is Enabled," *Social Epistemology*, 2024) shows trust forming in small working communities — *corpora popularia* — on three layers at once: the trained sensibilities of the verifiers, the working "little worlds" that host the verification, and the cultural anchors that accumulate across many such acts. Durable repair acts on all three; one-layer repair fails. A purely technical fix works only alongside the social and cultural layers that sustain it. The tradition for studying how such mechanisms are built runs through Bruno Latour's sociology of laboratory practice (*Laboratory Life*, 1979; *Science in Action*, 1987) and Lorraine Daston and Peter Galison's *Objectivity* (2007). Carlota Perez (*Technological Revolutions and Financial Capital*, 2002) adds the timing: epistemic crises strike at the turning point of every technological revolution — the shift from installing a technology to deploying it. We are inside that turning point now, with AI.

This is not a story of media producing order or disorder on their own. It is an account of innovation, collective sense-making, and resilience — and its one recurring variable is whether a community can keep presence bound to its claims as fast as a new medium, or a deliberate hand can strip it away.

The Atlas

The typology takes the form of an **atlas** — a practical reference map of historical cases that a reader can browse to find a situation resembling their own. It catalogs both the way presence is severed from claims and the mechanisms by which it has been restored.

The atlas is organized along three axes, each answering a simple question. *How was trust created?* (mechanism) *Where did it take hold — in people, in groups, or in lasting cultural objects?* (layer) *And how hard was it to repair once it was broken?* (cost of severance). The three eras organize the cases for the reader; these three axes organize them for analysis. Any case can be read along both lines.

The first axis is **mechanism**: the technique that binds presence to a claim. The isna did it through chains of named, character-evaluated transmitters. The Royal Society did it through witnessed demonstration and the "virtual witnessing" of its Philosophical Transactions when readers could not attend in person. The atlas populates them with cases from Islamic, South Asian, European, and other traditions — noting where the same mechanism-shape appears independently in more than one place.

The second axis is **the layer** a mechanism acts on, following Cowley's three-part model. Apprenticeship and embodied transmission work on the layer of sensibility; peer review and examinations act on the working community; canon and the seal-as-cultural-object work on the cultural anchors. Durable repair recruits protocols on all three layers at once, and the atlas makes that layering visible.

The third axis is the **cost of severance**: how hard it is to reattach presence once it has been lost. At the catastrophic end sits the Mesoamerican conquest (see Far Past cases): presence destroyed at the root, the limit benchmark against which lesser ruptures are measured.

The open call's named inspirations describe specific **dynamics of severance**. Jacques Ellul (*Propaganda*, 1965) treats propaganda as a systemic byproduct of technological-bureaucratic integration. Timur Kuran (*Private Truths, Public Lies*, 1997) describes the preference-falsification cascades by which a manufactured consensus becomes self-sustaining. Renée DiResta (*Invisible Rulers*, 2024) shows how attention markets restructure what counts as evidence and authority. These severance dynamics — alongside counterfeiting, attention hijacking, ecosystem capture, and epistemicide (catastrophic, state-coordinated destruction, anchored by the Mesoamerican case) — name the failure modes the typology catalogs. These map onto the three ways presence breaks: severance cuts the link outright, de-authorization withdraws trust from the authority behind it, and fragmentation splits a commons into rival authorities with no arbiter between them.

Draft Initial Typology of Epistemic Mechanisms

The same cycle recurs across three eras — far past, recent past, and present. Communities develop protocols that bind claims to accountable sources. That link is then severed — by a new medium, by deliberate actors, or by both. A working community invents a protocol to restore it. The far and recent cases show the cycle complete; the present cases show ruptures whose repairs are still being built — which is where our work begins.

Far Past

Vedic oral transmission (~1500 BCE onward). Sacred hymns had to survive for centuries with zero drift, before writing — and human memory slips. The repair: each verse was recited in fixed, overlapping orders, every syllable cross-checked against its neighbors, so a single error broke the pattern and was caught. A human error-correcting code, much like a computer's checksum but running on human memory alone. (Staal, *The Fidelity of Oral Tradition and the Origins of Science* (1986); Pollock, *The Language of the Gods in the World of Men* (2006); UNESCO. Intangible Cultural Heritage status, Vedic chanting (2008)).

The Islamic *isnad* and *hadith* criticism (8th–10th c. CE). Anyone could attribute a saying to the Prophet, and forgeries spread. The repair: every report carried a chain recording who heard it from whom, back to the source, with each transmitter's reliability biographically assessed — so a claim's trustworthiness rode on a traceable line of named people. Modern academic citation is its direct descendant. (Brown, *Hadith: Muhammad's Legacy in the Medieval and Modern World* (2009); Motzki, *The Origins of Islamic Jurisprudence* (2002); Hallaq, *The Origins and Evolution of Islamic Law* (2005)).

The Mesoamerican conquest (1521 onward). Spanish authorities burned the Mexica, Maya, and Mixtec codices and killed off the scribes who could read them — severance by deliberate hands, at the root. This is the typology's limit case: not strained but destroyed. The partial repair came late and from outside — Sahagún's Nahuatl codex, alphabetic transcription of oral traditions, glyph decipherment centuries on. Its echo is live: deplatformed communities, training-data erasure, history rewritten by decree. (Mignolo, *The Darker Side of the Renaissance* (1995); Boone, *Stories in Red and Black: Pictorial Histories of the Aztecs and Mixtecs* (2000); Lockhart, *The Nahuas After the Conquest* (1992)).

The Royal Society (1660s–1700s). As natural philosophy spread, a claim made in one place couldn't be trusted in another — no one else had seen the experiment. The repair: demonstrations performed before gentlemen witnesses and written up in the *Philosophical Transactions*, so a distant reader could follow the procedure as if present. The line runs straight to today's cryptographic attestation and IEEE/IETF work on verifiable credentials. (Shapin & Schaffer, *Leviathan and the Air-Pump* (1985); Shapin, *A Social History of Truth* (1994)).

Recent Past

The mass press and the photographic image (1880s–1920s). Cheap printing, the telegraph, and easily faked photographs made it possible to manufacture "evidence" and sell it to millions — yellow journalism, the Hearst–Pulitzer era, invented wars. The repair held for a generation: wire-service fact-discipline, journalism schools, professional ethics codes, libel reform. The closest pre-digital rehearsal for the deepfake era. (Schudson, *Discovering the News* (1978); Carey, *Communication as Culture* (1989); Lippmann, *Public Opinion* (1922); Bernays, Edward. *Propaganda* (1928))

Financial audit and the independent auditor (1929–1934). Before the Crash, a company could report almost any numbers it liked, and investors had no way to know who, if anyone, stood behind a balance sheet. The repair tied financial claims back to accountable people: required disclosure, standardized accounting rules, and — above all — the independent auditor, a licensed professional who examines the books and signs their name to them and stakes their professional standing on the result. The newly created Securities and Exchange Commission enforced it. (*Securities Act of 1933. Securities Exchange Act of 1934*).

Email and web authentication (1990s–2000s). As the internet scaled, anyone could forge a sender address or impersonate a trusted website — a message that looked legitimate with nothing tying it to a real origin. The repair, which members of our own team helped build: protocols (SPF, DKIM, DMARC) that bind a message to its sending domain, and certificate authorities that bind a website to a verified identity. When those failed — forged certificates — the community built the next layer, public Certificate Transparency logs. (IETF *RFC 7208* (SPF); *RFC 6376* (DKIM); *RFC 7489* (DMARC); *RFC 6962* (Certificate Transparency))

Present

The cryptographic arms race (1970s–present). Where email and web authentication bound a message or site to a verified origin, those bindings became targets the moment they carried weight. Each defense provoked a sharper attack: spam, malware, and denial-of-service strike at identity, security, and availability — the key elements of trustworthiness — and every repair invites the next infiltration. The game has now turned from reaction to anticipation. With quantum computing poised to break the public-key cryptography most of the internet rests on, communities are now repairing a rupture that has not yet happened — standardizing post-quantum algorithms years ahead of the threat. Presence-restoration has become a chess game, played several moves before the break. (NIST *FIPS 203, 204, 205* (2024)).

Wikipedia and the open web (2000s–present). The early web made anyone a publisher and cut the link between a claim and an accountable source; by the early 2000s "you read it on the internet" was a punchline, and an open encyclopedia anyone could edit looked like the problem in its purest form. Social media deepened the rupture and changed its shape: trust didn't collapse so much as fragment — feeds tuned for speed, outrage, and reinforcement sorted people into bubbles with their own authorities and incompatible realities, which bad actors widened with confusion, impersonation, and tactical ambiguity meant less to sell one lie than to make a shared reality feel impossible. Yet Wikipedia became one of the web's more trusted references — and a rare shared one across communities that otherwise disagree. It restored presence not by gatekeeping but by making accountability **legible**: every claim cites a source, every edit is logged to an account in a permanent public history, and disputes are settled on talk pages under enforced policies. The neutral-point-of-view rule is the move against fragmentation — it forces rival factions to negotiate a single contestable article rather than retreat to separate ones, rebinding trust across bubbles instead of within them. Repair runs on all three layers at once: trained editorial judgment, the community that adjudicates, and the cultural anchor of the policies themselves. (Reagle, *Good Faith Collaboration: The Culture of Wikipedia* (2010)).

Art and synthetic media (2022–present). Art is a story told by an artist to their audience - eternalizing their presence. Generative AI models sever that tie at both ends. On the output side, they produce artifacts no one made and no one signs: authorship without an author. On the input side, they are trained on the work of millions of artists, absorbed without name, credit, or consent. The repair is only now being assembled: provenance standards that cryptographically bind an asset to a record of how it was made and whether a model touched it (C2PA Content Credentials); artists' shielding or poisoning their work against unconsented training (Glaze, Nightshade); and creative communities, devising the norms, disclosures, and pedagogies that re-attach human judgment and accountability to work made with machines. The anonymous deepfake is the pure case: flawless to the senses, answerable by no one. (Shan et al. "Glaze." *USENIX Security* (2023); Shan et al. "Nightshade." *IEEE Symposium on Security and Privacy* (2024); Boucher, Smith, and Telliell. "Is Resistance Futile?: Early Career Game Developers, Generative AI, and Ethical Skepticism." *CHI '24* (2024); Boucher, Smith, and Telliell (2025, forthcoming)).

We expect to refine this set during Phase 1 in conversation with the Long Now and Protocol Institute teams.

Deliverables

A long-form essay (~15,000 words) framing the typology for a general public-facing audience and positioning it within contemporary debates about synthetic media, platform governance, and institutional adaptation. A practitioner-ready typology, in the form of an atlas — a visual reference mapping the recurring severance-and-repair patterns across three eras, with companion notes, usable by platform designers, standards bodies, content moderators, and policy practitioners.

At least **two facilitated workshops** with Long Now and Protocol Institute teams; and a **public presentation** to the Long Now community in Spring 02027. The atlas is designed to be extensible: the three-axis structure functions as a framework for analyzing further cases beyond the anchor cases we develop, so that subsequent contributors can populate additional cells.

At least two intervention recommendations, drawn from the atlas's repair patterns and developed into tested, practitioner-ready materials. These span the layers on which trust is rebuilt — societal, technological, institutional, economic, and process-based — because the historical record shows that durable repair always acts on more than one layer at once. They are recommendations designed and tested for implementation, not finished products.

Two examples illustrate the forms a recommended intervention might take. The first is **Witness**, an authentication-chain protocol drawn from the atlas's mechanism patterns (potentially socialized with the Fab Foundation maker community at Fab 26, July 02026) which, if carried forward, is offered toward deployment beyond the grant period. The second is a **hands-on experience** — a small software interaction that lets a participant feel a severance-and-repair cycle firsthand: encountering a claim, watching its chain of accountability break, and working to re-establish it.

Approach and Timeline

The work moves from understanding to action in three phases. First, the team maps how trust has been severed and rebuilt across history, producing the atlas. Second, working from the atlas's repair patterns, the team designs candidate interventions for today's crisis. Third, those candidates are tested with practitioners and developed into recommendation materials others can use. The deliverable is the atlas plus at least three recommended interventions designed for real-world use — recommendations, not finished products: practical directions for repair in addition to historical analysis.

Phase 1 — Research and Sessions (July–November 02026). Synthesize the source literature, finalize the theoretical armature, and build the typology across all three eras — drafting the anchor cases and a v0 of the atlas's mechanism cells. The team assesses the recent and present eras — the internet, authentication, and AI systems we helped build — from the inside as well as the outside. Roger Meike's multi-scenario [signals dashboard](#) informs the present-day rupture axis weighting evidence by source quality and tracking convergence across scenarios for the ICT/AI transition. External collaborators across history of science, anthropology, science and technology studies, and protocol engineering to stress-test cases. The internal **workshop session** with the Long Now and Protocol Institute teams take place in the fall.

Phase 2 — Exploring and Innovating Interventions (December 02026–February 02027). From the atlas's repair patterns, the team designs and models at least three recommended interventions across the societal, economic, process, institutional, and technological layers. We will refine two examples to show the range: a technological intervention (eg *Witness*, an authentication-chain protocol drawn from the atlas's mechanisms), and an experiential intervention (eg a hand-on software interaction that lets people feel a severance-and-repair cycle firsthand). The final mix acts on multiple layers at once, as durable repair requires.

Phase 3 — Design and Development of Intervention Recommendations (March–May 02027). Test the recommendations with practitioner communities and develop them into final, practitioner-ready materials; complete the atlas and the **long-form essay**; and present the work to the Long Now community in Spring 02027 as public presentation and an interactive workshop. The recommendations are just that, not built or deployed solutions. When carried forward, a technological example like *Witness* is offered as a forward commitment for possible deployment beyond the grant period (Feb 27, July 02027).

About Appleseed Commons

[Appleseed Commons](#) is an informal group of futurists, practitioners and academics with a mission to connect technology with the people whose lives it will reshape — inspiring them to build positive futures.

Team

Kerry Lynn — Networking R&D (Apple, Cisco, Verizon Labs, Oracle); contributor to IEEE, ISO, and IETF technical standards; Fab Foundation network infrastructure lead.

Roger Meike — Technical executive (Sun Labs, Xerox PARC); 25+ issued patents in AI, ML, distributed systems, and IoT; deep AI strategy and infrastructure experience.

Glenn Scott — Computer science researcher (System Development Corp, Sun Microsystems Laboratories, Xerox PARC), 120 issued patents spanning computer networking, chemistry, and distributed systems.

Yunus Telliel — Faculty Collaborator (WPI). Assistant Professor of Anthropology and Rhetoric (PhD CUNY, postdoc UC Berkeley); Director of WPI's Great Problems Seminar. Author of *Languages of the Qur'an: Science, Translation, and Religion in Turkey* (Indiana UP, forthcoming) and currently writing on generative AI and the human; co-chair, IEEE Brain Neuroethics Subcommittee work-and-employment working group.

Diane Weiss — Product leadership across enterprise AI and emerging tech (Intuit, Apple, HP, Oracle, Sun); Berkeley SkyDeck advisor; Santa Clara University lecturer.

Risks and Mitigations

Atlas coverage. A typology spanning many civilizations could become unwieldy or omit important traditions. *Mitigation:* the atlas anchors on deep cases across three eras and several traditions (Islamic, Western, South Asian, Mesoamerican), expand laterally as the pattern stabilizes, then distill; an explicit geographical-coverage check is performed at each phase boundary.

Scholarly depth. The team is practitioner-led rather than academy-led. *Mitigation* — *and our advantage*: the team has spent careers building the internet, authentication, and AI systems at the center of the present rupture, and can both diagnose it and design for it in ways purely historical scholarship cannot. Durable repair needs all three of Cowley's layers, and the team attends to all three. Telliel anchors the historical-comparative scholarship, and Phase 2 invites three external collaborators across history of science, anthropology, STS, and protocol engineering to extend it.

Scope. Designing interventions alongside the typology expands the work. *Mitigation*: the interventions are *recommendations*, designed and tested, not products the team commits to ship; examples such as *Witness* are illustrations of an approach, bounded and optional, not deliverables on a fixed timeline.

Synthetic-media moving target. Generative AI evolves faster than scholarly cycles. *Mitigation*: the typology is designed to be protocol-agnostic; Telliel's active AI book project keeps the scholarly thread current; the team's AI practitioners keep the practitioner thread current; Roger's signals dashboard tracks the contemporary axis live.

Repair-claim overpromising. Recent putative commons-repair technologies (cryptocurrency, blockchain-for-everything) have largely underdelivered relative to their marketing. *Mitigation*: the typology evaluates proposed contemporary repairs against the historical pattern rather than advocating for any specific tool; the team's posture is skeptical-by-default about claims that a new protocol resolves epistemic disorder without sustained institutional work on all three layers. Any candidate intervention we surface, such as *Witness*, is offered as an illustration to be tested for its failure modes, not as a finished standard.

Closing

We do not propose to restore a lost equilibrium. We are not trying to turn the clock back to a more trustworthy past that never really existed. We propose to recover protocol designs already stress-tested at the civilizational scale, and to build the next equilibrium with them. The cycle does not end with us. The medium and the mechanism change; the pattern of establish, sever, and repair runs on, from the long past into the long future. Our work is to read that pattern closely enough to chart the next turn deliberately.

This team has spent careers building infrastructure and technical protocols — not just social norms. The open call's framing is a direct invitation. By Spring 02027 we will deliver an atlas that explains how epistemic trust breaks and heals, and a tested set of intervention recommendations for repairing today's commons — usable by the platform designers, standards bodies, and content moderators the call names. We would welcome a further opportunity to discuss shape and scope.