

Facts, Trust, and Cooperation: The KFD Foundation Triad

Keren Dong
Kungfu Origin Technology Limited
keren.dong@kungfu.link

Draft: July 2026

Abstract

Persistent work must remain intelligible across participants, tools, and time. That requirement is often described with desirable values such as truth, trust, and cooperation, but a list of values does not explain what makes a shared work world possible. This paper presents the first three Kung Fu Decisions (KFDs) as an ordered systems model rather than a code of conduct: facts must not drift; trust must start from facts; and cooperation must start from trusted value. The three decisions preserve, respectively, continuity of the work world, warranted reliance on that world, and the capacity of independent participants to generate further joint action. Their composition is recursive: cooperation produces new occurrences and consequences, which must become inspectable inputs to the next fact cut. We test the model through removal and order-inversion arguments and through one cross-session continuation case. We then show how declared perspective, governed primitive discovery, and causal experience arise as the first pressures on this foundation, providing a narrow bridge to KFD-4, KFD-5, and the draft KFD-6. The contribution is a falsifiable account of the minimal roles required for a durable cooperative world, not a claim of complete ontology, moral consensus, or universal sufficiency.

1 Introduction

A conversation can end while the work continues. A maintainer leaves, an agent process stops, a repository moves, a release is promoted, or a new participant arrives with no access to the private context that made the last decision seem obvious. If the work is to survive, the newcomer must recover more than a summary. The newcomer must be able to identify what happened, decide which claims may be relied upon, and determine whether and how to continue.

This is a systems problem before it is an interface problem. Better models, larger context windows, and more persuasive explanations do not by themselves preserve the identity of the work across time. Nor does a complete event log automatically tell a participant what a claim proves or why cooperation is worth entering. Persistent work needs a shared world that can outlive the participant who currently interprets it.

The first three Kung Fu Decisions (KFDs) state a compact foundation for that world [6]:

KFD-1: facts must not drift.

KFD-2: trust must start from facts.

KFD-3: cooperation must start from trusted value.

Read as three imperatives, the statements can sound like good advice. This paper advances a stronger and more falsifiable interpretation: they name three different roles that persistent

cooperation cannot collapse into one another. KFD-1 preserves the continuity of the world about which participants reason. KFD-2 converts parts of that world into bounded permission to rely. KFD-3 converts warranted reliance and visible value into coordinated action among participants who retain independent judgment.

The dependency is ordered, but its operation is recursive:

$$\text{facts} \rightarrow \text{trust} \rightarrow \text{cooperation} \rightarrow \text{new occurrences} \rightarrow \text{successor facts.} \quad (1)$$

This loop is the paper’s central object. It explains why the triad can be foundational without being a complete ontology. The three KFDs do not name every object a cooperative system will need. They preserve the conditions under which new action can occur, its consequences can correct the shared world, and missing objects can later be discovered without replacing reality with narrative.

The paper makes four contributions:

- it defines facts, trust, and cooperation as operational roles rather than moral qualities;
- it uses removal and order-inversion tests to explain why the roles are separately necessary and directionally dependent;
- it models the triad as a recursive kernel through which a shared work world persists, becomes actionable, and produces further evidence; and
- it derives a narrow frontier from that kernel: declared perspective in KFD-4, separated genesis and qualification in KFD-5, and the causal-experience gate in the draft KFD-6.

The claims remain bounded. The paper does not prove that the triad is a mathematically unique basis for every social or technical system. It does not equate recorded claims with objective reality, trust with certainty, or cooperation with agreement. It argues that any system claiming durable cooperation must implement equivalent roles or explain how it preserves the same continuity, warranted reliance, and generative coordination without them.

Section 2 defines the continuity problem. Section 3 specifies the three roles. Section 4 applies removal and order tests. Section 5 develops the recursive model and a cross-session case. Section 6 shows how KFD-4/5/6 arise at the first pressure frontier. Sections 7 and 8 position and bound the claim.

2 The Continuity Problem

2.1 An operational shared world

We use *shared work world* in a deliberately operational sense. It is the recoverable set of facts, claims, artifacts, actions, responsibilities, constraints, and open questions through which participants identify and continue the same work. It is not a metaphysical claim that a system captures all of reality. A shared work world is sufficient when relevant participants can locate its authority boundaries, inspect what is known and unknown, and connect later action to earlier consequence.

This world is distributed. Material state may live in repositories and artifacts; causal state in runtime records and receipts; epistemic state in claims, evidence, and gaps; responsibility state in permissions and commitments. No participant directly observes the whole. Distributed cognition has long treated thinking as organized across people, artifacts, and environments rather than contained in one head [3]. Agents increase both the reach of that system and the cost of leaving its coordination implicit.

The persistence requirement can be stated as follows. Given participants p_i and p_j separated by time, tool, organization, or perspective, p_j should be able to continue the work without requiring the private memory of p_i and without silently inventing a new authoritative world. Perfect reconstruction is not required. Missing evidence, ambiguity, and disagreement may remain, but they must remain visible as such.

2.2 Three transformations

Persistent cooperation repeatedly performs three transformations:

1. occurrences, observations, and state become retained facts;
2. retained facts become warranted reliance on a bounded claim; and
3. warranted reliance and visible value become coordinated action.

Each transformation can fail independently. Records can change meaning while keeping the same name. Evidence can exist without establishing the claim for which it is cited. A correct claim can be used to obtain compliance through a hidden constraint rather than cooperation. Treating the three transformations as one generic notion of “trustworthiness” conceals these different failure modes.

The transformations are also lossy. A fact cut selects from reality. A trust assessment selects facts relevant to a purpose. A cooperation decision selects one action from available choices. A durable system does not eliminate that loss; it declares the cut, purpose, residual uncertainty, choice, and responsibility through which the loss occurred.

2.3 A running case: work after the session

Consider a task performed through an agent session. The agent edits source, runs checks, encounters one unresolved failure, and then the session ends. Three weak continuations are common:

- the transcript is treated as the work, although authoritative files and external state may differ from its narrative;
- a summary says the task is complete, but does not expose the evidence cut, failed check, or residual responsibility; or
- a new agent is instructed to continue, but cannot inspect why the next action is valuable or which constraints govern it.

The resulting failure is not merely lost context. The shared world has become dependent on private reconstruction. Kungfu’s product thesis that the session is not the work is one practical response to this continuity problem [5]. The KFD question is more basic: what must remain true for any response to work?

2.4 Why a value list is insufficient

Truthfulness, trust, and collaboration are widely desirable. Their desirability does not establish architectural necessity. A foundation claim requires counterfactual force: remove one role and identify a capability the remaining system can no longer provide; reverse their dependency and identify the unstable substitute that appears. The next sections therefore treat the triad as a candidate minimal kernel and subject it to those tests.

3 The Foundation Triad

3.1 KFD-1: continuity of the fact world

KFD-1 states that facts must not drift. A *fact* here is not whatever a system happens to store, nor a claim of observer-free certainty. It is a load-bearing reference admitted from a declared source, authority, and evidence boundary. KFD distinguishes a *fact cut*, which states what holds at such a boundary, from a *causal record*, which states what occurred between declared boundaries [6]. Either may include known gaps, uncertainty, or conflicting evidence.

Let F_t be a fact cut used at time t , and let c be its declared compatibility boundary. If the system presents the same fact identity and boundary at t' without declaring a change, its load-bearing meaning must be preserved:

$$\text{id}(F_t) = \text{id}(F_{t'}) \wedge c_t = c_{t'} \implies \text{meaning}(F_t) = \text{meaning}(F_{t'}). \quad (2)$$

The invariant does not freeze the world. It requires semantic change to cross an explicit boundary so later evidence can contradict, supersede, or narrow an earlier model without silently rewriting what the earlier participant saw.

KFD-1 therefore provides *world continuity*. It lets two participants refer to the same admitted state or occurrence even when they disagree about its interpretation. Without such a reference, disagreement cannot reliably be distinguished from drift in the object being discussed.

3.2 KFD-2: warranted reliance

KFD-2 states that trust must start from facts. Trust is not modeled as a global property of a person, model, product, or artifact. It is a bounded permission to rely on a claim for a purpose.

For claim C and purpose P , a trust assessment can be represented as:

$$T_P(C) = \mathcal{A}(C, P, F, E, R, Q), \quad (3)$$

where F is the bound fact cut, E the checked evidence, R residual risk, and Q the assignment of source, verification, and decision responsibility. The result may be pass, warning, fail, or unverifiable. Its force is limited to the stated purpose and cut.

This definition explains why more data is not the same as more trust. A log may be authentic while irrelevant to a completion claim. A check may pass while omitting the platform a consumer uses. A capable agent may produce a persuasive explanation without exposing the facts that would let another participant challenge it. KFD-2 makes reliance revisable because contrary facts can change the assessment without first defeating the authority or reputation of the claimant.

KFD-2 therefore provides *judgment continuity*. It allows a later participant to recover not only what was claimed, but why and within what boundary the claim was permitted to guide action.

3.3 KFD-3: generative coordination

KFD-3 states that cooperation must start from trusted value. Facts and trust do not themselves create joint action. A system may expose accurate evidence while hiding the value exchange, available alternatives, or constraints under which participation is requested.

Let V be a participant-relevant value claim, $T(V)$ its trust assessment, K the visible constraints, and O the meaningful available options. A cooperation decision has the form:

$$C = \mathcal{C}(V, T(V), K, O, Q_C), \quad (4)$$

where Q_C declares responsibility for the commitment and resulting action. The participant need not agree with every other participant. It must be able to understand what is offered, what is constrained, and what accepting, rejecting, escalating, or withdrawing means within the declared boundary.

This is not an absence-of-control rule. Sandboxes, permissions, refusal, revocation, and hard safety gates may be necessary. KFD-3 distinguishes explainable, reviewable constraint from hidden capture, and cooperation from mere compliance. The distinction matters operationally: a compliant component executes inside one authority; a cooperating participant can contribute new facts, objections, alternatives, and commitments to a shared world.

KFD-3 therefore provides *action continuity*. It lets independent participants extend the work without surrendering the fact and judgment boundaries that made the extension trustworthy.

3.4 Three irreducible questions

The triad can be summarized by three questions:

Layer	Role	Question preserved across participants and time
KFD-1	World continuity	What state or occurrence are we talking about, and where may its meaning change?
KFD-2	Judgment continuity	Why may this claim be relied upon for this purpose, and what remains unproved?
KFD-3	Action continuity	Why and within what visible choices and constraints should independent participants act together?

Table 1: The three operational roles of the Foundation Triad.

The roles correspond to reality as admitted, belief as warranted, and action as coordinated. They are related but not interchangeable. That separation is what makes the foundation testable.

4 Necessity and Order

The triad is a foundation only if each role carries independent load and if later roles cannot safely manufacture their own earlier conditions. We apply two conceptual tests: removal and order inversion.

4.1 Removal tests

Remove KFD-1. Suppose a system retains trust assessments and cooperation records but does not preserve stable fact identities and change boundaries. A later participant cannot determine whether an assessment still refers to the same contract, artifact, occurrence, or meaning. Trust attaches to a moving object. Conflict is resolved by the latest narrative or strongest authority because no stable reference remains against which either can be challenged.

Remove KFD-2. Suppose facts are stable and participants are willing to cooperate, but no purpose-bound assessment connects facts to reliance. Every participant must either repeat the complete verification, accept another participant’s reputation, or act without knowing which gap matters. The archive may be excellent, but reliance does not scale. Stable facts are necessary evidence; they are not self-interpreting permission to act.

Remove KFD-3. Suppose facts and assessments are inspectable, but the system has no account of participant value, choice, or visible constraint. It can support audit, certification, and command execution. It cannot distinguish durable cooperation from compliance obtained by monopoly, hidden workflow capture, or unreviewable pressure. The world may be correct and trustworthy yet remain non-generative across independent participants.

The failure modes are summarized in Table 2.

Removed role	What remains	Structural failure
Stable facts	trust language and coordination	reliance binds to mutable narrative or authority
Warranted trust	records and willing action	repeated verification, blind reliance, or reputation substitution
Trusted-value cooperation	facts and assessment	accountable control may remain, but independent joint generation collapses into compliance

Table 2: Removal tests for the three foundation roles.

These tests establish functional non-redundancy, not mathematical uniqueness. Another vocabulary may implement equivalent roles. The KFD claim is that the capabilities themselves cannot be omitted from durable cooperation.

4.2 Order-inversion failures

The arrows in the triad describe dependency, not a rigid wall-clock sequence. Participants may gather evidence, negotiate value, and act in overlapping cycles. Nevertheless, a later layer may not retroactively create the basis on which it claims to stand.

Trust before facts. When a system asks for reliance before exposing the relevant facts, identity, reputation, confidence, or persuasive narrative becomes the substitute. Facts may later be selected to defend the trust already granted. Correction then requires defeating an authority relationship rather than updating an assessment.

Cooperation before trustable value. When a system asks for commitment before value, facts, choices, and constraints can be inspected, participation becomes the evidence that the offer was acceptable. Lock-in is misread as preference and compliance as agreement. The system can grow while losing the ability to learn why participants would have chosen differently.

Successor facts admitted by the action that needs them. When completed action is allowed to rewrite its own prior fact cut, a result can erase failed attempts, costs, reversals, or disconfirming evidence. The loop becomes self-certifying: success defines which history counts. KFD-1 therefore applies again when new occurrences are admitted into the successor world.

4.3 Necessary, not sufficient

The triad does not guarantee good outcomes. Facts may be incomplete, an assessment mistaken, visible choices economically unequal, or cooperation harmful to others outside the declared boundary. Domain knowledge, ethics, law, security, governance, and empirical correction remain necessary. The

foundation claim is narrower: without continuity of fact, warranted reliance, and participant-aware coordination, those additional mechanisms cannot form a durable shared work world.

5 The Recursive Kernel

The ordered triad becomes a world-building mechanism only when it closes the loop from current facts through action to successor facts.

5.1 From fact cut to successor fact cut

Let F_t be the declared fact cut available at time t . A participant forms a purpose-bound trust assessment T_t , decides whether to cooperate through C_t , and performs an action whose real occurrence is $E_{t,t+1}$. Evidence from that occurrence is reviewed and admitted into a successor fact cut:

$$F_t \xrightarrow{\text{assessment}} T_t \xrightarrow{\text{value,choice}} C_t \xrightarrow{\text{action}} E_{t,t+1} \xrightarrow{\text{review,admission}} F_{t+1}. \quad (5)$$

$E_{t,t+1}$ is not identical to its record, and its record does not automatically become accepted fact. The distinction prevents an actor, tool, or model from defining reality solely through its own report. Admission may accept, reject, qualify, or leave evidence unresolved. What matters is that the path remains inspectable enough for later correction.

The kernel has three useful properties.

Persistence. Another participant can identify the prior cut, assessment, commitment, and result without recovering a private conversation.

Correction. New consequence can revise trust, expose a bad decision, or invalidate an object model while preserving what the earlier cut actually contained.

Generation. Cooperation does not terminate in agreement. It creates new actions, artifacts, failures, counterexamples, and commitments that expand the world available to later participants.

Facts are therefore not merely the beginning of the chain. They are also the governed return point. Trust is not merely confidence. It is the bridge that makes a fact cut actionable. Cooperation is not merely social harmony. It is the generative transition through which the world can change.

5.2 Cross-session continuation

Return to the agent task from Section 2. A strong continuation implements the loop as follows:

1. The ending session leaves authoritative source coordinates, changed artifacts, check results, the unresolved failure, known omissions, and responsibility boundaries. The transcript may be evidence, but is not silently promoted into the whole fact world.
2. A fresh participant checks a bounded completion claim against those facts. The result may be warning rather than pass because one failure remains unresolved.
3. The participant sees the value of continuing, the allowed actions, the blocked or approval-gated actions, and the party responsible for the next decision.
4. The participant chooses a bounded continuation, performs it, and leaves new source changes, observations, receipts, and verification results.

5. Review admits those results into a successor cut without rewriting the first session’s evidence boundary.

The session can disappear because it is not the identity of the work. The continuation does not require blind trust in the first agent, because trust is reconstructed from retained facts. The second participant does not merely obey a next-step prompt, because value and constraints are inspectable. The result becomes material for a third participant. Continuity is thus an outcome of the entire triad, not of storage alone.

5.3 Why this is a foundation for KFD

Every later KFD can introduce new objects, procedures, or domain profiles, but it must re-enter this kernel. A new object needs a stable fact identity. Claims about it need purpose-bound assessment. Its use with reasoning participants needs trusted value and visible constraint. Its action produces consequences that must return through review rather than self-certification.

This dependency explains foundation without requiring the first paper to catalog every later KFD. The triad is a generative kernel, not a compressed table of contents.

6 The First Derived Frontier

The recursive kernel creates new work and new evidence. That success exposes pressures the triad names but does not by itself operationalize. KFD-4, KFD-5, and the draft KFD-6 are the first derived responses [6]. They are introduced here only to demonstrate the foundation’s generative force; their full mechanism belongs to companion work.

6.1 KFD-4: facts are viewed from somewhere

KFD-1 preserves a declared fact world, but no participant consumes every fact from every position. Relevance, consequence, visible object boundaries, and available action depend on perspective. If a view hides its observer and transformation history, selection can masquerade as reality and undermine the trust assessment built on it.

KFD-4 therefore requires views to remain bound to declared perspectives. A perspective-bearing view identifies its observer position, accepted fact cut, projection boundary, consequence position, and known gaps. Changing perspective may change relevance or presentation without silently changing the source identities and causal relations required for comparison.

This is a direct pressure on the triad:

$$\text{non-drifting facts} + \text{multiple participants} \implies \text{declared perspective.} \quad (6)$$

The Observer Paper develops this problem independently. The foundation needs only the dependency: a stable world does not imply a view from nowhere.

6.2 KFD-5: the world may lack an object

Repeated action may reveal that the current ontology cannot economically describe a recurring part of reality. Participants reconstruct the same boundary, responsibility, or causal unit because no load-bearing object names it. KFD calls a candidate object with independent identity, boundary, authority, lifecycle, and operations a *Primitive*.

The triad constrains discovery. Candidate genesis may come from judgment, perspective transformation, anomaly, causal analysis, compression, or another declared method. Qualification must remain separate. It asks whether facts, alternatives, falsifiers, and real work support allowing the candidate to carry responsibility. A compelling name or coherent model cannot certify itself.

This pressure can be written as:

$$\text{successor facts} + \text{repeated reconstruction} \implies \text{candidate genesis} \neq \text{qualification}. \quad (7)$$

KFD-5 is therefore not an unrelated creativity principle. It applies KFD-1 and KFD-2 to revision of the world itself, while KFD-3 lets different participants contribute counterexamples and value judgments without fusing their authorities.

6.3 KFD-6: experience may support repeated discovery

Once causal work histories become inspectable, an agent may compare many Episodes and propose missing objects. The draft KFD-6 asks whether that process can scale while remaining grounded in causal experience and unable to certify itself. Corpus cuts, instrumentation loss, held-out evidence, intervention bounds, and independent promotion authority become load-bearing.

The companion paper *From Episodes to Primitives* develops this path [4] as:

$$\begin{aligned} \text{perspective-bound replay} &\rightarrow \text{candidate genesis} \rightarrow \text{independent qualification} \\ &\rightarrow \text{held-out evaluation} \rightarrow \text{separate promotion authority}. \end{aligned} \quad (8)$$

KFD-6 remains draft. The relevant foundation result is not that autonomous discovery already works. It is that scaling the recursive loop creates an explicit new proof obligation: generated narrative must not replace the causal world from which it claims to learn.

6.4 Evaluation program

The foundation should be evaluated by capabilities, not agreement with its language:

Fact continuity.

Give independent participants the same declared coordinate at different times. Detect whether its meaning, gaps, and authority boundary remain recoverable or change silently.

Reliance test.

Give a fresh participant a claim and its declared evidence. Test whether the participant can recover the purpose, checked facts, residual risk, and responsibility without private oral context.

Cooperation-independence test.

Remove privileged interfaces and hidden prompts. Test whether participants can still inspect value and constraints, refuse or choose a bounded path, and leave a reviewable commitment.

Loop-integrity test.

Inject failure, reversal, missing evidence, or a contradictory consequence. The successor world should expose qualification or disagreement rather than rewrite the prior cut or synthesize success.

Alternative-kernel test.

Accept a narrower or differently named model if it preserves the same world, judgment, and action continuity at lower cost.

These tests can falsify an implementation and may narrow the theoretical claim. Independent adoption and comparative measurement remain necessary before asserting cross-domain sufficiency.

7 Related Work

7.1 Distributed cognition and augmentation

Distributed cognition studies how cognitive work is organized across people, artifacts, and environments [3]. Engelbart’s augmentation program similarly framed computing as increasing collective human capability rather than merely automating isolated tasks [1]. The Foundation Triad shares this system-level view. Its specific claim is that a distributed work world needs separate mechanisms for preserving reference, warranting reliance, and enabling participant-aware action.

7.2 Situated action and coordination

Suchman argues that plans cannot be separated from situated practice [10]. Winograd and Flores place commitment and coordination at the center of work mediated by computers [12]. KFD-3 is compatible with those traditions but depends explicitly on KFD-1 and KFD-2: commitments remain revisable only when their fact and trust boundaries survive the situation in which they were formed.

7.3 Causality, event histories, and provenance

Lamport’s ordering model distinguishes causal structure from a universal physical clock [7]. Event sourcing reconstructs state from recorded events [2], while distributed tracing preserves relations among runtime operations [8]. These systems provide important material for fact cuts and causal records. The KFD model adds a separation between occurrence, evidence-bearing record, purpose-bound trust assessment, participant commitment, and admission into a successor fact world.

Software supply-chain systems make artifact origin and build process inspectable. SLSA defines provenance expectations [9], and in-toto models authorized steps and evidence [11]. These mechanisms are strong examples of trust beginning from facts. The triad generalizes the dependency beyond artifact release while preserving the same caution: valid provenance for one purpose does not automatically prove every downstream claim or create cooperation.

7.4 What remains distinct

Existing work provides rich accounts of shared cognition, situated action, commitment, causal ordering, event history, and provenance. The distinct proposal here is their composition into a small recursive kernel. Stable reference, warranted reliance, and generative coordination remain separate roles, and every action must return through consequence and review to the shared world it changes.

8 Discussion and Limitations

8.1 A systems foundation, not a moral ranking

The words *facts*, *trust*, and *cooperation* carry moral and political meanings. This paper does not rank participants by possession of those virtues. It assigns operational responsibilities to a system: preserve the reference, expose the warrant, and make the basis of coordinated action inspectable. A participant may be honest while using a drifting contract; a reputable institution may make an under-evidenced claim; competitors may interoperate reliably without mutual affection.

8.2 Facts are admitted, not total

KFD-1 does not turn a database, repository, or ledger into reality. Every fact cut has an observer, source, authority boundary, and omissions. Stability can preserve a mistaken model. The recursive kernel matters because contrary consequence can be retained without rewriting the earlier cut, making correction possible. KFD-4 makes the perspective boundary more explicit; it does not make all perspectives equally supported.

8.3 Trust is bounded, not accumulated prestige

Trust assessments are purpose-specific and revisable. A system may be trusted to preserve bytes but not to interpret their value; an agent may be trusted to run a bounded transformation but not to approve its deployment. Aggregating these assessments into a universal trust score would discard the purpose, evidence, and responsibility boundaries on which KFD-2 depends.

8.4 Cooperation is not harmony or permissiveness

Participants may disagree, compete, refuse, or operate under strict safety and legal constraints. KFD-3 applies where a system asks reasoning participants to coordinate. It requires participant-relevant value and constraints to be inspectable; it does not guarantee equal bargaining power or eliminate coercive conditions outside the system boundary. A complete evaluation must therefore examine meaningful alternatives and affected non-participants, not only interface transparency.

8.5 Minimal roles are not a complete ontology

The removal arguments establish that the three roles are non-redundant within the proposed model. They do not prove that the English terms are uniquely minimal or that every domain must expose them as first-class objects. A system that preserves equivalent continuity through another decomposition is a valid counterexample to vocabulary-level necessity. Security, incentives, law, embodiment, resource constraints, institutional history, and domain semantics remain outside the triad's complete account.

8.6 Originating-ecosystem bias

The model was developed inside the Kungfu ecosystem and is reflected in its product, release, and agent-work designs. That coherence is implementation evidence but also a source of overfitting. The cross-session case is a design trace, not proof that the same decomposition minimizes cost in healthcare, public administration, scientific collaboration, or other domains. Independent adopters may find a narrower kernel, different boundaries, or social practices that outperform formalization.

8.7 Formalization has a cost

Fact cuts, evidence bindings, responsibility assignments, and cooperation records move reconstruction cost from participants into infrastructure. The move is justified only where the recurring cost, consequence, or duration of the work exceeds the cost of maintaining the mechanism. A short-lived local task may need no formal KFD surface. The foundation is load-bearing where the work must cross a boundary and still remain the same work.

9 Conclusion

The KFD Foundation Triad is not a request to value facts, trust, and cooperation. It is an account of three roles through which persistent work can remain one shared world across participants and time.

KFD-1 preserves continuity of admitted state and occurrence. KFD-2 turns that world into bounded, revisable permission to rely. KFD-3 turns warranted value into coordinated action without reducing independent participants to captive execution surfaces. Cooperation then produces new occurrences and consequences, which return through review and admission to a successor fact cut.

The resulting loop explains both the triad’s minimality claim and its generative force:

$$\text{Facts} \rightarrow \text{Trust} \rightarrow \text{Cooperation} \rightarrow \text{New Facts.} \quad (9)$$

Remove stable facts and trust attaches to narrative. Remove warranted trust and facts remain inert or require repeated reconstruction. Remove trusted-value cooperation and an inspectable system can still obtain compliance but cannot account for durable joint generation. Reverse the order and authority, reputation, capture, or self-certification fills the missing layer.

The first derived pressures follow from operating the loop, not from expanding a list of values. Facts are viewed from declared perspectives; surprising consequences expose missing objects; repeated causal experience raises the possibility and risk of autonomous discovery. KFD-4, KFD-5, and the draft KFD-6 begin that next inquiry.

The foundation remains falsifiable. Another model may preserve the same continuities more economically, and independent evidence may show that this decomposition does not generalize. Until then, the triad provides a compact engineering test for any system that claims its work can outlive the session: can a later participant recover the world, reconstruct why it may be trusted, and choose how to extend it so that the result becomes inspectable reality for whoever comes next?

References

- [1] Douglas C. Engelbart. Augmenting human intellect: A conceptual framework. Technical report, Stanford Research Institute, 1962.
- [2] Martin Fowler. Event sourcing. <https://martinfowler.com/eaaDev/EventSourcing.html>, 2005.
- [3] Edwin Hutchins. *Cognition in the Wild*. MIT Press, 1995.
- [4] Kungfu Origin Technology Limited. From episodes to primitives: Causal historical assets for evidence-governed ontology revision. <https://github.com/kungfu-systems/paper-episodes-to-primitives/releases/tag/v0.1.0-alpha.2>, 2026. Release v0.1.0-alpha.2, accessed July 22, 2026.
- [5] Kungfu Origin Technology Limited. Kungfu: The session is not the work. <https://github.com/kungfu-systems/paper-kungfu-product-white-paper/releases/tag/v0.1.0-alpha.10>, 2026. Product white paper, release v0.1.0-alpha.10, accessed July 22, 2026.
- [6] Kungfu Systems. Kung fu decisions. <https://github.com/kungfu-systems/kfd/releases/tag/v1.0.0-alpha.42>, 2026. Release v1.0.0-alpha.42, accessed July 22, 2026.
- [7] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, 1978.

- [8] OpenTelemetry Authors. Opentelemetry documentation. <https://opentelemetry.io/docs/>.
- [9] SLSA Authors. Supply-chain levels for software artifacts. <https://slsa.dev/>.
- [10] Lucy A. Suchman. *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge University Press, 1987.
- [11] Santiago Torres-Arias, Hammad Afzali, Trishank Karthik Kuppusamy, Reza Curtmola, and Justin Cappos. in-toto: Providing farm-to-table guarantees for bits and bytes. In *28th USENIX Security Symposium*, pages 1393–1410. USENIX Association, 2019.
- [12] Terry Winograd and Fernando Flores. *Understanding Computers and Cognition: A New Foundation for Design*. Ablex Publishing, 1986.