GOT-V1-3b-5

On the Origin of Trust

Primitive Trust, Vicarious Closure, and Action Under Residual

---

Abstract

Trust is not first belief, confidence, loyalty, moral goodness, social virtue, institutional reputation, or epistemic warrant. These are later enrichments, claims, settings, or failures around trust. At its origin, trust is a finite-observer function: the reduction of relational uncertainty to satisfactory reliance closure for a locus under residual.

A finite observer-system cannot possess reality as total state. It encounters relation through boundary, interface, scale, horizon, abstraction, registration, retention, comparator pressure, closure demand, and residual limit. It must therefore operate through Soft Horizons: provisional, purpose-indexed closures that make relation usable without making it total. Trust is the reliance-form of that condition. It appears where a locus treats a relation as dependable enough to continue, expose itself, delegate, infer, coordinate, wait, or act without reopening the full relation at every moment.

The primitive advantage of trust is branch-space economy. A finite locus cannot revalidate every dependency before every movement, judgement, inference, or action. It cannot continuously recalculate surface, breath, balance, perception, memory, language, tool, social signal, instrument, institution, or record from first principles. Trust lowers that work-cost by allowing some relations to drop below immediate audit while remaining available to update. It does not remove residual. It makes continuation possible under residual.

Audit must therefore be grounded alongside trust. At the primitive level, audit is relation-revalidation: the redeployment of retained distinction against a present registration, claim, reliance relation, closure, or action outcome in order to test whether the relation still holds strongly enough for the closure demand being made. Trust and audit are not opposites. Trust backgrounds relation for action. Audit foregrounds relation for correction. Trust reduces audit so a finite locus can proceed; audit preserves the path by which trust can be reopened when anomaly, mismatch, changed condition, raised stakes, affected-locus exposure, or comparator failure makes background closure insufficient.

Primitive trust begins before explicit belief. A body trusts the floor before it forms a proposition about structural support. It trusts breathable air until signal reopens audit. It trusts balance, surface, light, familiar paths, caregiver rhythm, and environmental regularities because retained relation has reduced present uncertainty below the threshold required for continuation. This does not mean the body possesses certainty. It means the relevant relation has closed enough for the next branch.

Modern trust is a vicarious abstraction of this primitive function. A locus trusts another person, role, institution, record, procedure, credential, instrument, model, or machine output where it permits that other closure process or trust carrier to satisfy its own closure demand. Vicarious trust expands branch-space capacity by outsourcing closure-work. But the efficiency gain carries residual forward. The trusting locus now depends not only on the target relation, but on the carrier's closure process, domain, horizon, evidence path, abstraction, threshold, incentives, safeguards, and update conditions.

Within vicarious reliance, Epistemic Authority identifies the standing earned through demonstrated relation-preservation, including the grounds for preferential reliance between loci, while Structural Authority names permission or mandate within an organised information-processing structure.

Where vicarious trust relies on the output of another distinction-conditioned locus, the relied-upon relation may already carry Responsibility Provenance. Trust can background that provenance for finite use, but it does not transfer, erase, or recreate responsibility. A relying locus acquires responsibility only for the consequential branch relation attributable to its own distinction-conditioned organisation where the Responsibility discriminator is met. That relation may consist in a later transformation, a responsibility-bearing non-application, a retained or maintained condition, or material conditioning of another locus's Counterfactual Inclusion Field.

The danger begins when trust is reified. An indexed reliance closure becomes treated as a possessed thing, transferable warrant, moral debt, institutional status, brand asset, certificate, or authority token. A trust carrier is mistaken for a trust ground. A closure earned in one relation is spent in another. Residual disappears behind confidence. Questions become disloyalty. Action taken under threshold is later narrated as if uncertainty had been removed.

The governing rule is simple: trust must not carry reliance-strength beyond the recoverable conditions under which the trusted relation, carrier, or closure process satisfies the relying locus's closure demand.

Trust is necessary because finite systems cannot continuously audit total relation. It is dangerous because its function is to reduce audit. It remains legitimate only where relation, locus, carrier, closure demand, residual, threshold, safeguard where relevant, and update remain recoverable.

---

1. Introduction: Trust, reliance closure, and reduced audit

This paper concerns trust as a finite-observer closure function.

It does not offer a complete psychology of trust, a sociology of trust, a moral theory of trust, a political theory of trust, a legal doctrine of reliance, a theology of faith, or a general account of interpersonal intimacy. Those domains matter. Trust in human life is affective, developmental, institutional, ethical, legal, political, and existential. It can be wounded, cultivated, demanded, betrayed, repaired, ritualised, commodified, formalised, and abused. But those are not the primitive condition. They are later forms built on a deeper structure.

The primitive question is simpler.

How can a finite locus proceed where relation remains incomplete?

A finite observer-system does not possess the world as total state. It cannot reopen every dependency before every act. It cannot re-establish surface, breath, balance, memory, language, social expectation, instrument calibration, institutional record, and environmental regularity at full resolution before each movement, judgement, or decision. A system that attempted total audit before action would not become more rational. It would become non-functional.

The problem is therefore not only incomplete closure. It is the cost of repeated recalculation under incomplete closure. Every act selects a branch. If every supporting relation had to remain

open until fully revalidated, branch-space would not be navigable. The system would not merely lack certainty. It would lack a usable world.

Finite systems therefore live through closures. Some closures concern recognition. Some concern classification. Some concern communication. Some concern modelling. Some concern decision. Some concern action. Some concern reliance. Trust belongs to this last family. It appears where unresolved relation is treated as settled enough to reduce the work of branch-space selection for the relying locus under residual.

Trust is not closure without residual. It is not certainty. It is not proof. It is not final knowledge. It is not the possession of another person, process, object, institution, or system as fully transparent. Trust is a closure under residual by which a locus permits reliance to proceed.

This distinction matters because trust is easily inflated. Ordinary language treats trust as if it were a thing: someone has trust, earns trust, loses trust, deserves trust, breaks trust, restores trust, or asks for trust. The language is useful, but it hides the structure. Trust is not possessed in the way an object is possessed. It is not carried as a transferable substance. It does not automatically travel with a name, role, credential, record, office, model, certificate, method, or institution.

Trust is indexed. A locus trusts some relation, through some carrier, under some closure demand, for some purpose, over some horizon, with some residual, at some threshold, under some update condition. To remove these conditions is to convert trust into a floating noun. The word remains forceful, but the relation that earns it has disappeared.

The governing rule is therefore:

Trust must not carry reliance-strength beyond the recoverable conditions under which the trusted relation, carrier, or closure process satisfies the relying locus's closure demand.

This rule separates trust from warrant without opposing them. Trust may be warranted. Trust may be poorly warranted. Trust may be necessary where warrant remains incomplete. Trust may justify action under threshold without establishing epistemic closure. The error is not that finite systems trust before possessing total warrant. They must. The error is to let the trust closure narrate itself as if it had removed the residual under which it operates.

A patient may trust a doctor enough to follow a treatment while diagnosis remains probabilistic. A passenger may trust an aircraft enough to board without inspecting its engineering. A reader may trust a source enough to use its report while knowing that further verification would be required under higher stakes. A community may trust a bridge, a currency, a court, a standard, an archive, or a measurement system without possessing the full dependency chain by which each remains reliable. These trust closures may be rational, necessary, and strong for their purpose. None becomes total by being acted upon.

The scope of this paper is the condition that makes such reliance possible and dangerous. It begins with primitive trust before belief. It then grounds primitive audit as the counterfunction trust reduces. It decomposes the trust relation into locus, relation, carrier, closure demand, residual, threshold, safeguard, and update. Later sections show how retained relation supports pattern trust, how thresholds determine legitimate reliance, how vicarious trust permits one closure process to satisfy another locus's closure demand, and how trust carriers scale reliance through persons, roles, records, institutions, models, and machines. The final movement addresses reification, manipulation, betrayal, repair, and the point at which trust carriers become authority carriers requiring formalisation.

The present paper therefore occupies a specific position. Horizon analysis has already shown that finite observer-systems operate through relational availability, closure demand, Soft

Horizon, residual, and update. The discipline of the unknown has shown that residual must not be converted into false closure and that action under residual does not become certainty by being necessary. Power has established the distinction-conditioned branch potential through which a locus can make downstream realisation differ. Responsibility has then localised the consequential branch relations attributable to loci and shown how Responsibility Provenance composes across distributed paths without simply transferring between them. Those relations may include realised transformation, responsibility-bearing non-application, retained or maintained conditions, and cross-locus field-conditioning where the discriminator is met. Trust now names the reliance mechanism by which finite systems continue across retained and inherited relation without reopening every dependency at every use. This does not make responsibility a precondition of every primitive trust relation; it fixes the additional architecture wherever reliance crosses responsibility-bearing relations produced, maintained, or governed by other loci.

Trust is not the escape from finite access.

Trust is one of finite access's necessary forms.

---

2. Trust before belief

Trust is usually approached too late.

The familiar human case arrives already saturated. To trust someone appears to involve belief, expectation, vulnerability, confidence, loyalty, memory, moral evaluation, social standing, and sometimes love. A person trusts a friend, a partner, a witness, a doctor, a teacher, a colleague, an institution, or a source. The trusted party may be judged sincere, competent, discreet, reliable, caring, honest, or faithful. The trusting party may feel secure, exposed, dependent, reassured, anxious, betrayed, or foolish.

All of this is real. None of it is primitive.

The primitive structure appears before trust becomes interpersonal and before it becomes propositional. A body stands before it states that standing is safe. It steps before it proves the floor. It breathes before it verifies the air. It reaches before it calculates surface resistance. It turns toward sound before it identifies the source. It eats familiar food before it reconstructs the full causal pathway by which food becomes safe. It moves through space by treating countless relations as closed enough for action.[1]

These are not beliefs in the ordinary reflective sense. The body does not continuously assert that gravity will continue, the floor will hold, the light will disclose shape, the air will remain breathable, the hand will meet resistance, or the next step will find surface. It proceeds because the relevant relations have been backgrounded as reliable enough under current conditions.

That backgrounding is not empty. It is the outcome of inherited form, retained pattern, prior correction, biological calibration, environmental regularity, and ongoing signal sensitivity. The body has learned surface, weight, pressure, balance, fall, breath, pain, distance, grip, rhythm, and obstruction before it can articulate any of them. Some of this learning is individual. Some is developmental. Some is evolutionary. Some is social. Some is built into the structural coupling between organism and world. In every case, the system continues by not reopening the full dependency relation.

The floor is trusted when the relation between body-weight and surface-support is closed enough for stepping. The air is trusted when the relation between atmosphere and breath is closed enough for inhalation. Vision is trusted when the relation between light, object, interface, and orientation is closed enough for navigation. Balance is trusted when proprioceptive and vestibular relation is closed enough for posture. A familiar path is trusted when retained relation between route and arrival is closed enough for movement. A caregiver is trusted when the relation between need, signal, and response has closed enough for dependence.

None of these closures is total. Floors collapse. Air becomes toxic. Vision misleads. Balance fails. Food poisons. Paths change. Caregivers disappear. The point is not that trust removes such possibilities. The point is that a finite system cannot treat every possibility as equally open at every moment. It must reduce branch-space. It must let some relations become settled enough for the next act.

Primitive trust is this reduction of relational uncertainty to satisfactory reliance closure under residual.

It is a reliance closure because it permits action, exposure, delegation, waiting, inference, or coordination. It is satisfactory because the closure is indexed to a threshold, not to total certainty. It is under residual because some uncertainty remains. Trust is therefore neither ignorance nor knowledge in the strongest sense. It is not the absence of uncertainty. It is the condition under which uncertainty has been reduced enough for reliance.

A system that cannot trust cannot move. This is not a moral statement. It is a structural one. The system that must revalidate surface before every step cannot walk. The system that must revalidate breathability before every breath cannot breathe. The system that must revalidate memory before every use cannot think. The system that must revalidate every sign before every interpretation cannot speak. The system that must revalidate every social cue before every coordination cannot cooperate. The system that must revalidate every instrument before every measurement cannot investigate. The system that must revalidate every record before every decision cannot govern.

Total audit before action is not superior rationality. It is paralysis.

Trust is therefore a branch-space economy. It reduces the number of relations that must remain open for immediate recalculation. It permits a locus to act through retained closures instead of reconstructing the world from first principles. It allows the world to become usable without becoming possessed. It is one of the ways finite systems convert overwhelming relation into navigable branch-space.[2]

This does not make trust automatically good. A false closure may also reduce branch-space. A misleading source may still reduce uncertainty. A harmful institution may still stabilise expectation. A manipulator may provide precisely the cues needed to lower audit. Trust is powerful because it works before certainty. It is dangerous for the same reason.

The same structure that allows a body to rely on the floor allows a person to rely on a liar, a court to rely on contaminated evidence, a patient to rely on a miscalibrated test, a platform to rely on a biased classifier, or a public to rely on an institution whose procedures no longer preserve the relation they claim to carry. Trust does not become safe by being necessary.

Primitive trust must therefore be held below belief and below warrant. Belief may later attach to trust. Confidence may accompany it. Loyalty may form around it. Moral obligation may govern it. Institutional reputation may scale it. Epistemic warrant may strengthen it. But these are later layers. The lower function is reliance closure: relation closed enough for a locus to proceed before total closure is available.

This is why trust can operate before proposition. A body does not need to formulate “the floor will hold” in order to rely on the floor. It does not need to formulate “the air is breathable” in order to inhale. It does not need to formulate “this repeated path remains available” in order to walk it. Proposition may later describe, justify, challenge, repair, or formalise the trust relation. It is not the origin of the reliance.

Trust also operates before explicit confidence. Confidence is often a conscious or affective state. Trust may be affectively felt, but it need not be. Many trust closures are silent until they fail. The body does not usually feel confidence in gravity. It simply moves. A driver does not usually feel confidence in every road marking. The markings enter action. A speaker does not usually feel confidence in every word. Language carries relation beneath attention. Trust becomes visible when it is challenged, betrayed, demanded, or lost.

The invisibility of primitive trust is one reason it is easily misdescribed. We notice trust most where it becomes fragile: in vulnerability, promise, dependency, betrayal, risk, care, authority, evidence, expertise, and institutional action. We then mistake those later dramatic cases for the origin. But the same lower structure is already present wherever reliance proceeds under residual.

The primitive case is not smaller because it is less reflective. It is more basic because it shows why trust is needed at all.

Every finite locus must simplify its world enough to continue. It must let some relations stand as sufficiently resolved. It must carry residual without making residual blocking everywhere. It must treat some pattern, object, surface, signal, person, record, or process as dependable enough for the present branch. Trust is the reliance-form of that simplification.

The discipline is calibration. The system must be able to rely without pretending reliance is certainty. It must be able to background relation without making update impossible. It must be able to inherit prior relation without freezing it into authority. It must be able to continue under residual while preserving the conditions under which residual can return.

Trust is not the opposite of scrutiny.

Trust is the calibrated reduction of scrutiny under retained, updateable relation.

---

3. Primitive audit: the counterfunction of Trust

Trust cannot be understood without the function it reduces.

The current language of audit is culturally loaded. It suggests external accountancy, regulatory inspection, compliance review, institutional oversight, financial verification, certification, or adversarial scrutiny. Those forms matter, but they are not primitive. They are later social, institutional, legal, administrative, and technical elaborations of a lower function.

At the primitive level, audit is relation-revalidation.[3]

A locus audits when it redeploys a retained distinction, expected condition, record, calibration, pattern, signal, or comparator against a present registration, claim, reliance relation, closure, or action outcome in order to test whether the relation still holds strongly enough for the closure demand being made.

This definition keeps audit below institution and above mere comparison.

Audit is not every feedback loop. Physical persistence does not constitute audit. A physical process does not become audit merely because an external observer can compare its states. Audit requires a comparator path. Something retained, expected, registered, or available must be redeployed against a target relation. The comparison must bear on whether that relation is preserved for some closure demand. The output must be capable of preserving, weakening, reopening, revising, residualising, or rejecting the closure.

Retained distinction gives audit its comparator. Bounded registration gives audit its target. Applied distinction gives audit its operation. Comparator pressure gives audit its force. Residual and update give audit its discipline.

A simple bodily case shows the structure. A foot expects surface. Pressure, angle, resistance, and balance register whether the expected relation holds. If the surface gives way, the body reopens the relation. It shifts weight, withdraws, rebalances, freezes, or searches. The audit is not reflective in the ordinary sense. It is relation-revalidation under closure demand: can this relation still carry body-weight for this action?

Breathing provides another case. Air is normally trusted below attention. The relation is backgrounded. But smoke, smell, constriction, dizziness, or pain reopens audit. The body revalidates breathability. It coughs, holds breath, moves away, seeks ventilation, or signals distress. The trust relation was not false merely because it had been backgrounded. It became insufficient under changed condition.

Audit is therefore not the permanent opposite of trust. It is the reopening of relation when the current reliance closure is no longer enough.

This gives the central reciprocity:

Trust backgrounds relation for action.

Audit foregrounds relation for correction.

Trust reduces audit so a finite locus can proceed. Audit preserves the possibility that trust can be recalibrated when conditions change. Without trust, the system cannot move because every dependency remains open. Without audit, the system cannot correct because backgrounded relations harden into unexamined closure. Trust without audit becomes capture. Audit without trust becomes paralysis.

This reciprocity also explains why trust is not irrational by default. Trust is often rational precisely because audit has previously done enough work, or because the relation has been retained and corrected enough, or because the stakes and residual allow reliance. A body does not verify every surface because prior relation and ongoing signal are sufficient for ordinary movement. A scientist does not recalibrate every instrument from scratch before every reading because calibration records, protocols, and comparator checks reduce the need for local reconstruction. A citizen does not inspect every bridge before crossing because engineering, maintenance, material practice, public record, and institutional responsibility are trusted as carrier processes.

But this same structure explains why trust can become dangerous. If the audit path disappears, trust continues to reduce scrutiny after the relation has weakened. If comparator pressure is

absent, a carrier may retain confidence without relation-preservation. If update is blocked, trust becomes inert. If affected-locus exposure is hidden, the trusting locus may benefit from reduced audit while another locus bears the residual.

Primitive audit matters because it shows that the counterfunction of trust is not suspicion. Suspicion may trigger audit, but audit itself is relation-revalidation. It asks whether the relation still holds sufficiently for the closure demand being made. Suspicion can be miscalibrated; audit can correct it. Trust can be miscalibrated; audit can correct it. Audit is not distrust as mood. It is the disciplined return of backgrounded relation to foreground.

This matters especially when trust becomes social or institutional. Later forms of audit often become procedural. A report is reviewed. A record is checked. A witness is cross-examined. A model is validated. A bridge is inspected. A medical decision is second-read. A court decision is appealed. A machine output is monitored. These are complex, formalised, and often institutional versions of the same lower function. They redeploy retained distinction against present claim or reliance in order to test whether relation still holds.

The institutional form should not obscure the primitive function. A compliance process that does not revalidate relation may be called audit, but it is not audit in the relevant sense. A certificate that confirms form without testing relation may reduce local scrutiny, but it does not by itself restore the trust ground. A record review that checks completeness while missing target failure may preserve procedure but not relation. A machine evaluation that tests benchmark performance while missing deployment relation may create apparent audit without adequate comparator.

Primitive audit is therefore both simpler and stricter than institutional audit. Simpler, because it is relation-revalidation. Stricter, because forms, roles, certificates, reports, and procedures do not count unless they actually test the relation under the closure demand.

Trust reduces audit; audit revalidates trust. This gives finite systems a dynamic rhythm. Relation is foregrounded, retained, backgrounded, relied upon, challenged, reopened, recalibrated, and backgrounded again. The rhythm may be fast and embodied, slow and institutional, formal and mathematical, procedural and legal, or automated and machine-mediated. The architecture remains the same.

When trust is well calibrated, audit is not absent. It is latent, recoverable, and triggerable. The system need not audit at full resolution at every moment, but it must preserve the conditions under which audit can return. It must know what would reopen the relation: anomaly, mismatch, changed condition, raised stakes, affected-locus challenge, comparator failure, threshold shift, or residual becoming blocking.

A trust relation with no audit path is not merely strong trust. It is closure without correction.

That is why the trust relation must now be decomposed.

---

4. The Trust relation: locus, relation, carrier, demand

Trust is often spoken of as if it were simple and direct.

One person trusts another. A citizen trusts an institution. A patient trusts a doctor. A user trusts a system. A reader trusts a source. A passenger trusts an aircraft. The grammar suggests that trust is a relation between a trusting subject and a trusted object.

That grammar is too compressed.

Trust is not merely a relation between two named entities. It is an indexed reliance closure. It involves a relying locus, a trusted relation, the availability condition and access path under which that relation stands for the relying locus, a carrier where trust is mediated, a closure demand, a purpose, scale and horizon, residual, threshold, safeguard where consequence requires one, and update condition.

The relying locus is the locus whose branch-space is reduced by reliance. It may be a person, body, institution, instrumented practice, social organisation, machine-mediated process, or coupled observer-system. The relying locus is not always the same as the affected locus. This distinction becomes crucial in institutional and machine-mediated trust. A public body may rely on a model, while a claimant bears the consequence. A hospital may rely on a diagnostic protocol, while the patient bears the residual. A court may rely on a record, while a defendant bears the exposure. Trust must therefore identify who benefits from reduced audit and who bears the risk if the closure fails.

The trusted relation is the relation treated as dependable enough for the purpose at hand. It is not the trusted object in total. A patient does not trust a doctor in every respect. The relevant relation may be diagnostic competence, procedural care, truthful disclosure, treatment judgement, confidentiality, or referral discipline. A passenger does not trust an aircraft as total object; the relevant relations include structural integrity, maintenance, pilot competence, navigational systems, regulatory oversight, weather judgement, and operational procedure. A reader does not trust a source absolutely; the relevant relation may be factual reporting, domain expertise, honest interpretation, currentness, or methodological care.

Trust becomes distorted when the trusted relation is replaced by a global noun. “Trust the doctor,” “trust the science,” “trust the institution,” “trust the model,” or “trust the process” may be pragmatically useful, but each phrase hides the relation actually being relied upon. A doctor may be trusted for technical competence but not for communication. A scientific field may be trusted for a measurement while remaining residual for interpretation. An institution may be trusted to maintain records but not to self-audit abuse. A model may be trusted for triage but not final classification. The trusted relation must be specified because trust does not apply to objects as totalities.

The carrier is the person, role, record, institution, instrument, model, credential, procedure, interface, standard, certificate, or machine output through which the trusted relation travels. In primitive embodied trust, the carrier may be retained pattern within the organism-world relation. In social trust, the carrier may be a person’s conduct, memory, role, promise, or competence. In institutional trust, the carrier may be office, procedure, archive, credential, inspection, legal rule, or public record. In machine-mediated trust, the carrier may be dataset, architecture, output, benchmark, interface, dashboard, validation report, or deployment monitor.

The carrier is necessary where the relying locus cannot directly revalidate the target relation. But the carrier is not the ground. A credential is not competence by itself. A record is not warrant by itself. A role is not reliability by itself. A model output is not target preservation by itself. A certificate is not closure by itself. The carrier reduces local recalculation only if its relation-preserving process is recoverable enough for the closure demand being satisfied.

The closure demand specifies what trust is being asked to carry. Trust for casual use differs from trust for high-stakes action. Trust for exploration differs from trust for coercion. Trust for prediction differs from trust for explanation. Trust for communication differs from trust for evidence. Trust for provisional triage differs from trust for final judgement. A relation may

satisfy one closure demand and fail another. The system must therefore ask: trusted enough for what?

Purpose and horizon further specify the closure. Trust may be local, temporary, domain-specific, action-specific, role-specific, or scale-specific. A bridge may be trusted for ordinary crossing under current maintenance conditions, not for unlimited load. A source may be trusted for reporting within its domain, not for all interpretation. A model may be trusted in one deployment environment, not after distribution shift. A person may be trusted with one kind of vulnerability, not another. Trust that travels beyond purpose and horizon becomes overextended.

Residual remains present in every trust relation. The question is not whether residual exists, but whether it blocks the reliance being granted. A residual may be acceptable for low-stakes action and blocking for high-stakes action. It may be acceptable for temporary reliance and blocking for permanent reliance. It may be acceptable for human-assisted review and blocking for automation. It may be acceptable for investigation and blocking for sanction. Trust becomes disciplined when residual is carried at the right level of force.

Threshold determines when the residual is low enough, or the need strong enough, for reliance to proceed. A body steps because the surface relation has crossed a bodily threshold for action. A patient follows advice because the doctor relation crosses a threshold for treatment. A court relies on evidence because the evidential path crosses a legal threshold. An institution relies on a record because provenance, procedure, and context cross an administrative threshold. A machine-mediated system relies on output because performance, monitoring, and deployment conditions cross an operational threshold. Threshold is not certainty. It is the point at which reliance becomes justified for the closure demand and consequence at issue.

Safeguard constrains reliance where residual remains consequential. It may take the form of appeal, second opinion, independent replication, limited deployment, monitoring, reversibility, human review, informed consent, transparency, right of correction, fail-safe, domain restriction, or stop condition. Safeguard matters because trust lowers audit. Where audit is reduced, the system must preserve ways to reopen, constrain, contest, or update the relation if residual becomes blocking.

Update is the condition under which trust must be revised. A trust relation without update is not calibrated; it is inert closure. Update may be triggered by anomaly, failure, changed condition, raised stakes, new evidence, comparator failure, domain shift, affected-locus challenge, breach, or evidence that the carrier no longer preserves the trusted relation. Trust must be able to weaken, narrow, suspend, transfer, repair, or terminate when its conditions change.

These components show why trust is not a transferable substance. It does not move automatically from one relation to another. Trust in a doctor’s medical competence does not automatically become trust in their financial judgement. Trust in an institution’s technical capacity does not automatically become trust in its incentives. Trust in a model’s benchmark performance does not automatically become trust in deployment. Trust in a record’s existence does not automatically become trust in the record’s truth. Trust in a person’s sincerity does not automatically become trust in their accuracy.

A trust claim should therefore be reconstructable as an indexed relation: this locus relies on this relation, available through this access path and carrier where relevant, under this closure demand, for this purpose, scale and horizon, with this residual, at this threshold, under this safeguard where needed, and subject to this update condition.

This is not bureaucratic excess. It is what ordinary trust grammar compresses.

In low-stakes, local, reversible, and well-patterned contexts, the structure may remain tacit. A body does not need to recite the trust packet before stepping. A friend does not need a formal audit path for every ordinary reliance. A community does not need to formalise every shared expectation. Finite life depends on tacit trust.

But as consequence rises, as trust becomes mediated, as the relying locus and affected locus separate, as carriers become institutional or machine-mediated, as residual becomes difficult to recover, or as trust is used to justify authority, the compressed structure must become more explicit. The system must be able to show what relation is being trusted and why the reliance-strength claimed does not exceed the recoverable conditions supporting it.

This decomposition prepares the later movement of the paper. Pattern trust shows how retained relation lowers audit across recurrence. Threshold discipline shows how residual can block some forms of reliance but not others. Vicarious trust shows how one locus borrows closure-work from another. Trust carriers show how reliance scales through persons, roles, records, institutions, models, and machines. Reification, manipulation, betrayal, and repair show what happens when carrier, threshold, residual, and update fail.

The central rule remains fixed:

Trust is legitimate only as indexed reliance closure under residual.

When its conditions remain recoverable, trust makes finite life possible.

When its conditions disappear, trust becomes one of the main ways unearned closure travels.

---

5. Pattern, embodied, and environmental Trust

Primitive trust does not remain confined to single acts.

A finite locus does not meet every relation as if for the first time. It retains. It repeats. It learns regularity. It carries prior correction into present readiness. A relation that has held before can drop below immediate audit unless signal, anomaly, consequence, or changed condition reopens it. Trust therefore extends from immediate reliance into pattern.

Pattern trust appears where retained relation lowers the cost of future relation-validation. The system does not merely rely once. It relies again because a relation has become sufficiently stable across recurrence. The floor has held before. The air has remained breathable. The path has led home. The word has carried meaning. The tool has worked. The person has responded. The institution has produced usable records. The instrument has remained calibrated. The model has performed inside its envelope.

This is not proof of permanence. It is retained relation under residual.

A pattern becomes trusted when recurrence has reduced uncertainty enough for reliance under the relevant threshold. The relation is no longer treated as wholly open at each occurrence. It becomes background. That backgrounding is a finite achievement. It allows an organism to move through a world without reconstructing the world at every moment. It allows a speaker to use language without revalidating every word. It allows an institution to rely on records without re-opening every fact. It allows a machine-mediated process to use stored distinctions without recalculating every prior classification.

Embodied trust is the primitive form of pattern trust. It is not a belief added to bodily action. It is the way retained relation becomes motor, perceptual, affective, and environmental readiness.

A body does not first infer the world and then act. It acts through prior closure that has become available as posture, balance, orientation, reach, breath, grip, avoidance, and timing.[4]

The body trusts surface by placing weight before reflective proof. It trusts breath by inhaling until signal disrupts the relation. It trusts proprioception by moving limbs without visual recalculation. It trusts light by orienting toward visible structure. It trusts repeated rhythm by anticipating return, resistance, or response. The trust is not propositional, but it is not empty. It is retained relation made operational.

Environmental trust extends the same structure beyond the organism's immediate body. A finite locus inhabits a world by allowing environmental regularities to carry reliance. Paths, doors, tools, signs, rooms, sounds, schedules, interfaces, devices, and social rhythms become trusted because they reduce branch-space. They make action possible without reopening the full dependency chain.

A familiar room is not audited as a new field at every entry. A tool is not reconstructed from its material history before every use. A word is not derived from linguistic origin before every utterance. A road sign is not revalidated from the institution that installed it before every turn. These relations remain provisional, but their provisionality is usually backgrounded. The world becomes usable through retained trust.

This usability is the achievement and the danger.

The achievement is obvious. Without pattern trust, finite life would collapse into recalculation. No organism could move efficiently. No language could function. No institution could coordinate. No science could build on prior results. No machine system could retain useful states. Every relation would demand full audit before use.

The danger is that pattern can outlive its conditions. The path changes. The air becomes toxic. The floor weakens. The word drifts. The record becomes outdated. The institution loses its correction path. The model moves outside its deployment envelope. The instrument falls out of calibration. The environment that made the pattern trustworthy has shifted, but the trust closure continues.

This is why pattern trust must remain updateable. A trusted pattern is not a permanent right to ignore relation. It is a reduced-audit state that remains legitimate only while the triggers for renewed audit remain available. Signals of mismatch matter because they reopen the relation. Pain, failure, anomaly, contradiction, complaint, drift, affected-locus challenge, changed condition, raised stakes, or comparator failure can make backgrounded trust insufficient.

Pattern trust also shows why trust is often invisible until it fails. The body does not notice its trust in balance until balance is threatened. A community does not notice its trust in records until records conflict. A user does not notice trust in an interface until the interface misleads. A legal system does not notice trust in category until an affected relation no longer fits. A scientific practice does not notice trust in instruments until calibration breaks. Trust becomes visible where reduced audit must be reopened.

The discipline is not to make all pattern trust explicit. That would destroy the function trust serves. A finite system cannot formalise every retained relation. Pattern trust often works precisely because it remains tacit, fast, and embedded. But the system must know when tacit trust is no longer enough. As stakes rise, as action becomes harder to reverse, as reliance becomes mediated, as affected loci bear greater exposure, or as the carrier path becomes less visible, pattern trust must become more recoverable.

Embodied and environmental trust therefore bridge primitive reliance and vicarious trust. At the primitive level, a locus relies on retained relation in its own coupling with the world. At the

vicarious level, it relies on relation retained, processed, certified, or closed elsewhere. Both reduce audit. Both preserve branch-space. Both can mislead when residual disappears.

The core structure remains the same: trust backgrounds relation for action, and audit must be able to foreground relation for correction.

---

6. Thresholds, residual blocking, and safeguard

Trust is not all-or-nothing.

A relation may be trusted weakly, strongly, locally, conditionally, temporarily, or only for a specified purpose. A person may be trusted for sincerity but not competence. A doctor may be trusted for diagnosis but not for financial advice. A record may be trusted as evidence of registration but not as warrant for truth. A model may be trusted for screening but not for final classification. A procedure may be trusted to produce consistency but not justice. A machine output may be trusted to prompt review but not to decide.

The question is therefore not simply whether trust exists.

The question is whether the relation is closed enough for the reliance being placed upon it.

Threshold names that point. A trust threshold is the condition at which relational uncertainty has been reduced enough for reliance under a specified closure demand, consequence, horizon, and residual. It does not eliminate uncertainty. It determines whether the remaining uncertainty blocks the reliance being considered.[5]

A residual may block one form of reliance while permitting another. This is the key to calibrated trust.

A medical residual may block diagnostic certainty but permit precautionary treatment. An evidential residual may block final judgement but permit further investigation. A model residual may block automation but permit human-assisted triage. A record residual may block coercive action but permit provisional inquiry. A source residual may block citation as authority but permit use as a lead. A procedural residual may block legitimacy claim but permit emergency action under safeguard.

Trust becomes overstrong when the system fails to ask what the residual blocks.

If every residual blocked every reliance, finite systems could not act. They would be paralysed by the ordinary incompleteness of relation. But if no residual blocks reliance, trust becomes unearned closure. Every uncertainty is absorbed into continuation. Every carrier becomes usable. Every action can be justified by necessity. Residual becomes ornamental.

The disciplined question is: what does this residual prevent the relying locus from doing, claiming, delegating, exposing, automating, enforcing, or forgetting?

Threshold depends on consequence. Low-consequence reliance may tolerate broader residual. High-consequence reliance requires stronger closure, clearer residual, more robust safeguard, and stronger update. A person may trust a casual recommendation under light residual. They should not trust the same recommendation for irreversible medical intervention without stronger relation-preservation. An institution may use a model to prioritise review under one threshold. It should not use the same model to impose sanction unless the residual is lower, the safeguard stronger, and the update path more explicit.

Threshold also depends on reversibility. A reversible reliance can sometimes proceed under heavier residual because correction remains possible. An irreversible or difficult-to-reverse reliance requires stronger support. The distinction is not always obvious. A temporary classification may become practically permanent once recorded. A provisional risk marker may affect future access. A machine output may enter an archive and shape later decisions. A reversible-looking trust closure may become durable in practice.

This is why safeguard belongs inside the trust relation.

Safeguard is the constraint that allows reliance to proceed under residual without allowing that residual to become unchecked exposure. A safeguard may limit scope, duration, domain, automation, authority, discretion, or affected-locus burden. It may require review, appeal, monitoring, notice, consent, second opinion, independent check, error correction, sunset, staged use, or withdrawal. Its function is not decorative reassurance. It is what keeps reduced audit from becoming domination.

Trust lowers scrutiny. Where lowered scrutiny affects only the relying locus in a low-consequence and reversible setting, a light safeguard may be enough. Where lowered scrutiny exposes another locus, grants authority, triggers coercion, automates classification, or produces durable records, safeguard becomes structurally necessary. The relying locus may benefit from trust while another locus bears the residual. In such cases, trust cannot be judged only from the perspective of the one who finds reliance convenient.

This distinction is central to institutional trust. An institution may trust a record because it saves time. The person recorded may bear the consequence if the record is wrong. A regulator may trust a certification because it simplifies oversight. The public may bear the risk if the certification fails. A platform may trust a classifier because it scales moderation. Users may bear the burden of misclassification. A court may trust a procedure because it produces reviewable form. A defendant, claimant, child, worker, patient, or citizen may bear the residual that the procedure fails to preserve.

Threshold and safeguard therefore protect trust from becoming domination.

They do not remove the need for reliance. They discipline its consequences. Trust remains finite closure under residual, but the residual must not be hidden from those who bear it. A closure that is good enough for institutional convenience may not be good enough for affected-locus exposure. A threshold adequate for internal triage may not be adequate for final judgement. A safeguard that is symbolic may not be adequate where correction is practically unavailable.

The rule is simple.

The stronger the reliance, the stronger the recoverability required.

The wider the exposure, the stronger the safeguard required.

The harder the update, the higher the threshold required.

Trust that ignores this relation turns branch-space economy into epistemic debt. It acts now by deferring the cost of unresolved relation. That cost returns later as harm, error, dispute, opacity, betrayal, institutional failure, or loss of legitimacy. The debt was not created by uncertainty itself. It was created by allowing uncertainty to disappear inside reliance-strength.

Properly calibrated trust does not deny residual. It carries residual at the right threshold, under the right safeguard, with the right update path.

That is what makes trust usable without making it total.

7. Vicarious Trust and borrowed closure

Trust becomes more powerful when it becomes vicarious.

Primitive trust allows a finite locus to rely on its own retained relation. Pattern trust allows prior relation to lower audit across recurrence. Vicarious trust goes further. It allows one locus to rely on the closure-work of another.

This is not merely trusting another person. It is permitting another closure process to satisfy one's own closure demand.[6]

The preceding paper sharpens this relation. Where borrowed closure depends on another distinction-conditioned locus, it may arrive through an already responsibility-bearing path. Responsibility Provenance remains local to the consequential branch relations attributable to that path, including the transformations, retained or maintained conditions, and field-conditioning relations that materially shaped what became available to the relying locus. Trust allows some of that provenance to remain below immediate audit; it does not cause responsibility to transfer to the relying locus. If the relying locus then classifies, combines, thresholds, transmits, withholds, redirects, acts on the carrier, or governs a standing condition through which its own distinction-conditioned organisation changes a later branch relation, that local contribution becomes a separate responsibility relation where the Responsibility discriminator is met.

A patient cannot personally reconstruct the full medical relation among symptoms, tests, physiology, diagnosis, treatment, risk, and outcome. The patient may trust the doctor because the doctor's training, instruments, experience, institutional role, evidence path, and professional constraints reduce the patient's need to perform that work locally. A passenger cannot inspect aircraft design, maintenance, pilot competence, air traffic systems, weather modelling, and regulatory history before boarding. The passenger relies on distributed closure-work. A reader cannot re-report every event. A citizen cannot audit every bridge. A scientist cannot personally rebuild every instrument and dataset. A judge cannot personally witness every fact. Modern life depends on borrowed closure.

Vicarious trust is therefore branch-space expansion through outsourced audit.

It lets finite loci act beyond their own direct access. It enables expertise, division of labour, education, medicine, science, law, engineering, finance, governance, transport, and technology. No complex society can function without it. Every extended institution is partly a mechanism for making another locus's closure-work available to those who cannot repeat it.

But vicarious trust does not remove residual. It relocates and multiplies it.

The relying locus now depends on at least two relations. First, the target relation: the medical condition, aircraft safety, legal fact, bridge integrity, experimental result, institutional classification, or modelled pattern. Second, the carrier relation: the competence, honesty, method, incentive structure, evidence path, calibration, procedural discipline, domain limitation, and update capacity of the person or system through which closure is borrowed.

The patient depends not only on the body-treatment relation, but on the doctor's closure process. The passenger depends not only on aircraft physics, but on maintenance, regulation, pilot competence, design, and operational systems. The public depends not only on official records, but on the processes by which those records are made, stored, corrected, and used. A user depends not only on the machine output, but on the data, architecture, evaluation, interface, threshold, deployment condition, and monitoring process by which that output is produced.

Vicarious trust therefore economises branch selection while over-carrying residual risk. The relying locus gains capacity by not repeating the work locally. It also inherits the possibility that the borrowed closure was wrong, stale, outside scope, misapplied, corrupted, biased, incentive-distorted, or no longer updateable.

This is not a reason to reject vicarious trust. Rejection would be impossible. A finite system that refused all borrowed closure could not live in a complex world. It could not use medicine, law, science, engineering, infrastructure, education, governance, archives, machines, maps, measurements, or language beyond direct verification. Vicarious trust is one of the main ways finite systems escape the poverty of immediate access.

The discipline is to prevent borrowed closure from becoming invisible.

A vicarious trust relation should preserve the target relation, the carrier relation, the closure demand being satisfied, the limits of the carrier's domain, the residual being carried, the threshold for reliance, the safeguard where consequence requires one, and the update path. The relying locus need not reconstruct every dependency at full resolution in every case. That would defeat the point. But the closure-work being borrowed must remain recoverable enough for the strength of reliance being claimed.

Vicarious trust is strongest where the carrier process is competent, domain-bounded, corrigible, and exposed to comparator pressure. A doctor's judgement is stronger where training, evidence, differential diagnosis, peer correction, record review, and follow-up remain active. A bridge authority is stronger where inspection, maintenance, reporting, engineering standards, and failure investigation remain active. A scientific claim is stronger where instruments, datasets, methods, replication, peer challenge, and alternative hypotheses remain recoverable. A machine output is stronger where data provenance, evaluation, deployment context, monitoring, and human review remain visible.

Vicarious trust weakens where the carrier becomes opaque, overgeneralised, insulated, incentivised against correction, or treated as authority by status alone. A credential may hide incompetence. A role may hide conflict. A record may hide poor provenance. A procedure may hide discretion. A certificate may hide weak inspection. A model may hide proxy substitution. An institution may hide affected-locus failure. A public reputation may hide loss of relation-preservation.

The central error is to treat the borrowed closure as if it arrived without residual. The relying locus says, in effect: I trust this carrier, therefore the relation is settled for me. That may be pragmatically necessary in low-stakes settings, but when consequence rises, the borrowed closure must be reconstructed at the right resolution. The question is not whether the relying locus personally repeats all closure-work. The question is whether the conditions under which that work earns reliance-strength remain recoverable.

Vicarious trust also explains why betrayal can be severe. Where one locus borrows another's closure-work, it lowers its own audit. If the carrier exploits that lowered audit, the trusting locus is exposed at the very point where reliance was meant to save work. The harm is not mere failure. It is the misuse of reduced audit.

8. Epistemic Authority and reliance

Vicarious trust allows one locus to proceed through the distinctions or closure-work of another. This introduces a question of epistemic standing within reliance. Another locus's contribution may preserve the target relation more strongly than the relying locus can presently achieve through its own available distinctions. Where that advantage is demonstrated, preferential reliance can have a recoverable ground.

Epistemic Authority is the defeasible standing of a claim, distinction, or closure-process to guide judgement or reliance, earned through demonstrated preservation of a specified target relation under stated availability conditions and closure demand.

Between loci, Epistemic Authority supports preferential reliance where one locus's available distinctions or closure-work provide a demonstrable advantage in resolving relational uncertainty concerning a shared target relation for the use being claimed. The authority extends only as far as the demonstrated advantage supports that reliance. It does not become a general property of the contributing locus.

The comparison must preserve its reference conditions. Two loci may encounter the same object while addressing different relations, operating at different scales, or seeking different closure. An advantage concerning detection does not by itself establish an advantage concerning explanation. An advantage under one horizon does not establish the same advantage under another. Preferential reliance therefore requires a shared target relation and a comparable closure demand, with the relevant availability, abstraction, scale, and horizon conditions recoverable.

Reduced relational uncertainty must also remain distinct from confidence. A contribution does not acquire Epistemic Authority because it expresses less doubt, presents a more precise output, or closes a question more readily. A stronger contribution may expose residual that a weaker one conceals. The comparison concerns demonstrated relation-preservation, including the treatment of what remains unresolved. It need not be expressed through a single numerical measure.

Relative advantage does not establish sufficient closure for every use. One contribution may be stronger than another while neither satisfies the closure demand being made. The advantage may support further inquiry, comparison, or provisional reliance while leaving stronger action blocked. Where action proceeds under residual, threshold, safeguard, and update retain their existing functions. Actionability does not substitute for demonstration of the epistemic advantage.

The advantage must also survive the reliance path. A locus may preserve the target relation strongly while its contribution reaches another through an incomplete record, altered abstraction, stale report, or misleading interface. Epistemic standing at the contributing locus does not establish the strength of what arrives. The target relation and the carrier relation must remain recoverable enough for the reliance being placed upon them. This does not require the receiving locus to repeat all prior closure-work; it requires the borrowed work to remain answerable to the present use.

Trust is the function through which such standing becomes operative in reliance between loci. It does not create the standing merely by granting it recognition. A well-grounded contribution may remain untrusted or unused. A trusted contribution may carry less epistemic strength than the relying locus attributes to it. Epistemic Authority and actual reliance must therefore remain distinguishable even where they coincide.

Where reliance makes a distinction operative through an Information Processing Locus and produces a counterfactually recoverable difference in downstream branch-realisation, the exercise of Epistemic Authority participates in an application of power. The Power discriminator establishes the consequential difference. Demonstrated relation-preservation establishes the epistemic standing. Neither the occurrence nor the magnitude of the consequence retrospectively supplies that standing.

The authority remains corrigible. Changed conditions, failed comparison, carrier distortion, or newly available distinctions may narrow, suspend, or reverse the preferential reliance previously warranted. Trust may reduce repeated audit while the supporting relation holds. It must preserve the conditions under which audit returns.

A further distinction is required where authority concerns permission or mandate.

Structural Authority is permission or mandate to act or decide within a specified scope, conferred or recognised by an organised information-processing structure.

Structural Authority does not establish Epistemic Authority by conferring permission or mandate. Epistemic Authority does not itself confer Structural Authority. Volume 1 requires this distinction to keep authorisation separate from earned epistemic standing. The further development of Structural Authority belongs to Volume 2.

This prepares the role of trust carriers. Vicarious trust does not usually travel through abstract closure-work alone. It travels through persons, roles, records, institutions, credentials, instruments, models, interfaces, certificates, procedures, and machine outputs. These carriers make borrowed closure usable. They also make it easier to mistake the carrier for the ground.

---

9. Trust carriers, generalisation, and reification

A trust carrier is any person, role, object, record, credential, procedure, institution, instrument, model, interface, certificate, standard, archive, platform, or machine output through which reliance-strength travels.[7]

A trust carrier may therefore make a relation shaped by responsibility-bearing organisation portable without making responsibility itself portable. Upstream provenance remains attached to the recoverable transformations, retained or maintained conditions, and field-conditioning relations that produced or materially shaped the carrier's content, status, or availability; downstream attribution depends on what the receiving locus itself transforms, withholds, maintains, or governs in relation to it. The carrier carries reliance-strength. It does not carry Responsibility Provenance as a transferable token.

Carriers are necessary because trust must move. A relying locus cannot always access the target relation directly. It needs signs, roles, processes, forms, artefacts, systems, and records through which prior closure-work can be inherited. A doctor's role carries expectations of competence. A maintenance log carries a history of inspection. A court record carries procedural registration. A scientific publication carries method, evidence, and review. A certificate carries

an assertion that some process has been completed. A model output carries hidden data and abstraction. An institutional office carries delegated authority. A machine interface carries system judgement into user action.

Carriers reduce recalculation. They allow trust to scale beyond immediate relation. Without them, vicarious trust would remain fragile and local. Expertise could not travel. Institutions could not coordinate. Records could not support future decisions. Scientific results could not be reused. Infrastructure could not be depended upon. Machine systems could not assist action beyond their designers' direct presence.

But a carrier is not a ground.

The carrier carries relation-preserving work only where that work remains recoverable enough for the reliance being claimed. A credential does not by itself establish competence. A certificate does not by itself establish adequacy. A record does not by itself establish truth. A procedure does not by itself establish justice. A model output does not by itself establish target preservation. An institution does not by itself establish legitimacy. A role does not by itself establish reliability.

The carrier is a path. It is not the relation itself.

Trust carrier failure begins when the carrier is treated as if its status supplies the ground. The doctor is trusted because they are a doctor, without regard to domain, evidence, attention, conflict, or error. The institution is trusted because it is official, without regard to affected-locus relation or correction. The record is trusted because it exists, without regard to provenance. The certificate is trusted because it has been issued, without regard to what it inspected. The model is trusted because it is advanced, without regard to deployment relation. The machine output is trusted because it is generated, without regard to target preservation.

The carrier then becomes a reified trust object.

Reification occurs when indexed reliance closure becomes treated as a possessed thing, transferable warrant, moral debt, institutional status, brand asset, certificate, authority token, or social demand. Trust is no longer recoverable as relation under closure demand. It becomes something one "has," "deserves," "owes," "earns," "breaks," "restores," or "signals." Ordinary language may need these terms. The conceptual danger begins when the language makes the conditions disappear.

Reified trust often travels beyond its original domain. A person trusted for competence in one field is trusted in another. An institution trusted for record-keeping is trusted for judgement. A model trusted for ranking is trusted for explanation. A certificate trusted for process completion is trusted for substantive adequacy. A public reputation built in one relation becomes authority in another. A trust carrier generalises because it is easier to carry the carrier than to reconstruct the relation.

Generalisation is not always wrong. A carrier may legitimately support wider trust where the underlying closure-work is relevant across cases. A doctor's training may generalise across many diagnoses. A scientific method may generalise across many experiments. A regulatory regime may generalise across many products. A machine architecture may generalise across tasks if the deployment relation supports it. The problem is not generalisation. The problem is unearned generalisation.

Earned generalisation preserves domain, horizon, comparator, residual, and update. It says: this carrier has supported this class of reliance under these conditions, and the present use remains within that class or has been extended through visible justification. Unearned generalisation says only: this carrier has been trusted before, therefore it should be trusted now.

This is a route by which trust is converted into unearned authority.

A carrier that repeatedly reduces audit gains social force. People stop asking what relation it preserves. Institutions build procedures around it. Records cite it. Users defer to it. Critics are treated as disloyal, ignorant, obstructive, or unreasonable. The carrier's history of successful reliance becomes a reason not to reopen its current relation. Trust then no longer functions as calibrated reduced audit. It functions as suppression of audit.

Questions become disloyalty when trust has been reified. A patient who asks about diagnosis may be treated as lacking trust. A citizen who asks about institutional procedure may be treated as undermining public confidence. A user who asks about a machine output may be treated as resisting efficiency. An affected locus who challenges a category may be treated as refusing the process. In each case, the demand for trust protects the carrier from audit.

This is backwards.

A legitimate trust carrier should be strengthened, not weakened, by recoverability. If the relation-preserving process is sound, audit can show why. If the carrier has limits, audit can locate them. If residual remains, audit can state whether it blocks the present reliance. If the carrier has drifted, audit can update it. A carrier that cannot tolerate relation-revalidation is not asking for trust in the disciplined sense. It is asking for insulation.

Trust carriers also produce trust laundering. Borrowed closure enters a carrier, then the carrier travels without the closure-work. A record cites a prior institution. A certification cites a standard. A standard cites professional practice. A model relies on labels produced by earlier decisions. A public decision relies on the model. Each stage appears to inherit trust. But if no stage preserves the underlying relation, reliance-strength circulates without returning to ground.

The danger grows with scale. A local trust failure may harm one relation. A reified carrier can distribute failure across many relations. A flawed record system can mislead multiple institutions. A weak certification regime can scale inadequate assurance. A biased classifier can reproduce inherited misclosure across populations. A captured institution can export its trust signal into law, policy, finance, medicine, education, or governance. Carrier trust is powerful because it lets closure travel. That is exactly why carrier-path discipline is necessary.

The rule remains strict:

The carrier carries; it does not consecrate.

To trust a carrier is to rely on a recoverable process by which relation has been preserved enough for the present closure demand. The carrier may reduce audit. It may not abolish the conditions under which audit could return.

This prepares the paper's negative movement. Manipulation, trust laundering, and betrayal exploit the same structure. They work because trust lowers audit, because carriers allow trust to travel, and because reified trust turns reduced audit into unearned authority.

---

10. Manipulation, Trust laundering, and betrayal

The danger of trust is not only that it may be mistaken.

The deeper danger is that it can be exploited because it works.

Trust lowers audit. It backgrounds relation. It reduces branch-space. It allows one locus to continue without reopening every dependency. That reduction is the benefit. It is also the vulnerability. A relation that has dropped below immediate scrutiny can be used, redirected, stretched, simulated, or captured before the trusting locus revalidates it.

Manipulation exploits reduced audit by capturing another locus's closure threshold.

A manipulator does not need to create certainty. The manipulator needs only to lower audit enough for reliance to proceed. They may supply familiar cues, borrowed credentials, institutional signals, social proof, emotional reassurance, technical language, procedural form, or urgency. They may mimic the carrier through which trust usually travels. They may induce the relying locus to treat a relation as settled enough for action before the relation has earned the closure demand.

This is why manipulation is structurally tied to trust. A completely untrusted relation is difficult to exploit because audit is already foregrounded. A trusted relation is available for exploitation because the system has already reduced scrutiny. Manipulation works by entering the path through which scrutiny has been reduced.

A scammer imitates institutional authority. A platform mimics social credibility. A speaker uses professional vocabulary to borrow expertise. A vendor uses certification language to imply adequacy. A public actor invokes shared loyalty to suppress audit. A machine interface presents outputs with confidence, fluency, and apparent objectivity. In each case, the manipulator does not merely assert a false proposition. They operate on the relying locus's threshold for reopening relation.

The key failure is not simple deception. Deception concerns false content. Manipulation concerns the control of another locus's audit conditions. It shapes what the relying locus asks, what it treats as settled, what it treats as disloyal to question, what it treats as urgent, what residual it ignores, and what carrier it accepts as sufficient.

This is why manipulation can occur even where some content is true. A statement may be accurate but used to suppress the residual that matters. A credential may be real but outside domain. A record may exist but not support the claim being carried. A procedure may have been followed but not preserve the affected relation. A model may perform on a benchmark but fail deployment. A true element becomes manipulative where it is used to lower audit beyond the conditions it earns.

Trust laundering is the institutionalised version of this movement.

Trust laundering occurs when reliance-strength passes through a carrier path and emerges cleaner, stronger, or more authoritative than the underlying relation warrants. Closure borrowed from one relation is spent in another. The laundering path may include records, credentials, certificates, standards, procedures, institutions, models, dashboards, compliance marks, public statements, or machine outputs. Each stage appears to inherit trust from the prior one. But if the grounding relation is not recoverable, the trust has been laundered rather than carried.[8]

A record may cite a prior record. A certification may cite a standard. A standard may cite professional consensus. A model may rely on labels generated by institutional practice. A decision may rely on the model. A public communication may rely on the decision. The final claim appears strongly supported because many carriers stand behind it. Yet the target relation may never have been preserved. The system has accumulated carrier-status rather than relation-preservation.

Trust laundering is especially dangerous because it produces the appearance of due care. The process may be documented. The form may be complete. The credentials may be valid. The model may be evaluated. The institution may be authorised. The language may be professional. But if the relation being trusted is not recoverable through those carriers, the chain has not increased warrant. It has only moved reliance-strength through socially recognised forms.

The correction is carrier-path recovery. The system must ask what relation each carrier was trusted to preserve, under what closure demand, with what residual, and whether that carrier legitimately transfers reliance-strength to the next use. A certificate may attest process completion but not substantive adequacy. A benchmark may attest performance under test conditions but not deployment relation. A record may attest registration but not truth. A professional role may attest training but not attention in this case. A compliance process may attest form but not relation-preservation.

Betrayal is the exploitation of reduced audit after trust has been granted.[9]

Responsibility Provenance helps localise the responsibility-bearing relations involved, whether realised through transformation or carried through a governed condition of branch availability, but provenance alone does not establish betrayal. A responsibility-bearing mistake may remain mistake rather than betrayal. Betrayal requires the additional trust-level condition developed here: exploitation, concealment, abuse, or knowing neglect of the reduced audit created by the reliance relation.

Betrayal is not mere failure. A trusted relation may fail without betrayal. A doctor may make a reasonable mistake. A friend may be unable to help. An institution may face conditions beyond its control. A model may fail under an unforeseen distribution shift. Failure becomes betrayal where the carrier, person, institution, or process exploits, conceals, abuses, or knowingly neglects the reduced audit created by trust.

The betrayed locus is harmed not only by the failed outcome but by the misuse of its reliance. It lowered scrutiny because the relation appeared dependable enough. The other side then used that lowered scrutiny as opportunity. This is why betrayal has a distinctive moral, institutional, and epistemic force. It is not simply that the world failed. It is that the path through which the trusting locus economised relation was turned against it.

Betrayal can be direct. A person lies where honesty was trusted. A professional abuses a role where competence or care was trusted. A witness falsifies where testimony was trusted. A caregiver exploits dependency where responsiveness was trusted.

Betrayal can also be institutional. A body invites public reliance while hiding known failure. A regulator presents oversight while inspection is weak. A platform promises safety while incentives reward opacity. A corporation markets certification while narrowing the inspection burden. A public authority asks for trust while treating challenge as disloyalty. In such cases, betrayal may be distributed across roles, records, incentives, and procedures. No single agent may carry the whole failure, but the trust relation has still been exploited.

Machine-mediated betrayal is more complex because machines do not betray in the interpersonal sense unless agency and intention are separately established. But machine systems can participate in betrayal structures. A machine output may be used to exploit trust in technical objectivity. A model may inherit a biased institutional category and present it as neutral classification. A dashboard may hide residual behind apparent precision. A deployment system may allow users or affected loci to believe that audit has occurred when only computation has occurred. The betrayal lies in the human, institutional, or governance relation that uses the machine as carrier.

Manipulation, trust laundering, and betrayal share one structure: reduced audit is converted into unearned authority.

The response is not universal distrust. Universal distrust would destroy the branch-space economy that finite systems require. The response is calibrated recoverability. Where trust reduces audit, the path back to audit must remain available. Where carriers move trust, the carrier path must remain inspectable. Where trust is scaled through institutions or machines, residual, threshold, safeguard, and update must remain operational rather than decorative.

Trust becomes legitimate by remaining reopenable.

---

11. Distrust, mistrust, and Trust repair

Distrust is not always failure.

A disciplined system must be able to withhold, narrow, suspend, or reverse trust when relation no longer satisfies the closure demand. Distrust may be the correct output of audit. If a carrier fails comparator pressure, if a relation no longer holds, if a threshold is no longer met, if residual becomes blocking, if an affected locus exposes hidden failure, or if update conditions are refused, then distrust may preserve relation better than continued reliance.

Distrust should therefore be indexed just as trust is indexed. Distrust of what relation? For which locus? Through which carrier? Under which closure demand? With what residual? At what threshold? Under what condition could trust return? Without this indexing, distrust becomes a mood, posture, identity, political signal, or blanket refusal. It may then become as undisciplined as reified trust.

Disciplined distrust says: this relation does not currently satisfy this reliance demand for this locus under these conditions. It does not need to say that the person, institution, model, or system is globally untrustworthy. A doctor may be distrusted for one treatment judgement while still trusted for care. An institution may be distrusted for self-investigation while trusted for record access. A source may be distrusted for interpretation while trusted for primary documents. A model may be distrusted for final classification while trusted for triage. Distrust, like trust, is relation-specific.

Mistrust is miscalibrated distrust.

It appears when audit is reopened without sufficient relation-based reason, or where a carrier is rejected beyond the residual that actually blocks reliance. Mistrust may arise from prior harm, false generalisation, identity conflict, political framing, unfamiliarity, misread signal, stigma, institutional history, or manipulation by another carrier. It may feel protective, and sometimes it begins from real residual. But it becomes mistrust when the refusal of reliance exceeds the recoverable conditions.

Mistrust is not corrected by demanding confidence. A system cannot repair mistrust by insisting that trust is owed. If mistrust has roots in prior betrayal, affected-locus exposure, opaque procedure, record error, category failure, or carrier overreach, then demands for trust may worsen the failure. The proper response is relation-revalidation. What relation is being asked for? What failed? What residual remains? What carrier path is now being used? What safeguard protects the exposed locus? What update has occurred?

False suspicion is a related but narrower failure. It occurs where audit is triggered by weak or irrelevant signals and then hardens into distrust. A person reads anomaly into ordinary

difference. An institution treats challenge as threat. A public treats technical complexity as concealment. A user treats unfamiliar design as manipulation. A system treats dissent as bad faith. False suspicion can damage trust by reopening relations that did not require reopening under the closure demand being made.

But the possibility of false suspicion must not be used to suppress audit. Many systems protect themselves by describing affected-locus challenge as mistrust, suspicion, paranoia, disloyalty, or ignorance. That is another form of trust reification. The question is not whether challenge feels disruptive. The question is whether it bears on the relation, carrier, residual, threshold, safeguard, or update condition.

Trust repair begins where trust has weakened, collapsed, or become rightly suspended.

Repair is not reassurance. It is not branding. It is not apology alone. It is not a demand to move on. It is not renewed confidence produced by communication strategy. Repair is the restoration of the conditions under which reduced audit becomes legitimate again.[10]

Where the failed trust relation crossed responsibility-bearing relations, repair must also recover how responsibility and reliance became misaligned. It must distinguish upstream provenance, the carrier's role, the relying locus's own local responsibility-bearing relation — whether realised through transformation, non-application, or governed field-conditioning — and any later Responsibility Warrant. This is not a demand to collapse distributed provenance into one blame allocation; it is a demand to recover enough of the path for reliance to be recalibrated without turning role or authority into causal attribution.

The first condition of repair is failure identification. The system must say what failed. Did the target relation fail, or the carrier relation? Did the threshold prove too low? Did residual become blocking? Was the carrier overgeneralised? Was the action threshold mistaken for closure? Was update ignored? Was safeguard absent? Was affected-locus exposure hidden? Was the trust relation betrayed, merely mistaken, or subjected to changed conditions?

The second condition is residual preservation. Repair fails if it closes the relation too quickly. A harmed locus may be pressured to accept reassurance before the residual has been localised. An institution may apologise without identifying what remains unresolved. A platform may adjust messaging without changing the carrier path. A professional body may discipline one person while leaving the trust architecture untouched. Repair requires the residual to remain visible until it has been addressed or properly limited.

The third condition is threshold recalibration. The earlier trust threshold may no longer be valid. Reliance that was once acceptable may need narrowing. A carrier may return to use only under lower stakes. A role may remain trusted in one domain but not another. A procedure may require additional review. A model may be limited to advisory use. A person may regain ordinary interaction but not high-vulnerability reliance. Repair does not always restore the same trust. Often it restores a more precise trust.

The fourth condition is safeguard. Where reduced audit exposed another locus to harm, repair must alter the conditions of exposure. Safeguard may require oversight, transparency, appeal, monitoring, second review, domain limits, independent verification, record correction, or withdrawal rights. Without safeguard, repair asks the affected locus to lower audit again while the original exposure remains unchanged.

The fifth condition is update. Repair must show what would prevent repetition or force reopening if repetition occurs. An apology with no update condition is affective closure without structural correction. A revised procedure with no monitoring is administrative reassurance. A certification with no comparator is carrier repair without relation repair. Trust can become legitimate again only when the conditions that failed can answer back.

Repair may produce different outputs.

It may restore strong trust where relation-preservation is re-established and residual no longer blocks the relevant demand. It may restore limited trust where the carrier remains usable only under narrower domain, lower threshold, or stronger safeguard. It may produce disciplined distrust where the relation cannot be restored. It may produce residual deferral where more evidence, time, or comparator pressure is required. It may reveal that the trust relation was never legitimate at the strength previously claimed.

The aim of repair is therefore not confidence.

The aim is calibrated reliance.

A repaired trust relation is stronger when it knows its limits. It does not ask the harmed or relying locus to forget. It does not convert apology into closure. It does not treat renewed use as proof of restoration. It restores, narrows, or refuses reduced audit according to recoverable conditions.

Trust repair is successful only where audit can again become latent without becoming impossible.

---

12. Institutional and machine mediation

Trust changes when it becomes mediated by institutions and machines.

The underlying function remains reliance closure under residual. But the trust relation becomes more complex because the relying locus, affected locus, carrier, target relation, decision path, and update path may separate. Local, embodied trust often keeps these close together. Institutional and machine-mediated trust pull them apart.

An institution allows reliance to scale. It retains records, stabilises roles, creates procedures, distributes authority, stores memory, assigns responsibility warrants, and makes decisions across time and persons. It allows one locus to rely on work done elsewhere: inspection, certification, adjudication, review, classification, measurement, archiving, governance, or enforcement. Without institutional trust, complex coordination would collapse.

Such warrants may allocate answerability, duty, burden, liability, supervision, or decision ownership; they do not by themselves establish Responsibility Provenance. Institutional trust therefore must not let a role, office, policy, or procedure make warranted answerability appear to be causal provenance by status alone.

But institutional trust carries distinctive residual.

Institutions can preserve relation through procedure, but they can also preserve procedure after relation has failed. A form may be complete while the affected relation remains invisible. A record may be current but wrong. A role may be authorised but conflicted. A category may be administrable but mis-scaled. A decision may be reviewable but unjustified. A process may be

consistent while consistently losing the target relation. Institutional trust becomes dangerous when procedural survival is mistaken for relation-preservation.

The relying locus and affected locus often diverge. An institution may rely on a category because it enables administration. The person categorised bears the residual. A court may rely on a record because it supports procedure. The party named in the record bears the exposure. A regulator may rely on certification because it simplifies oversight. The public bears the failure if certification is weak. This divergence means institutional trust cannot be judged only by the convenience or confidence of the institution.

Institutional carriers also create authority effects. A record, office, decision, certificate, or policy does not merely reduce audit; it may command action. It may allocate benefits, impose duties, deny access, authorise coercion, trigger surveillance, or structure later records. The carrier does not simply help a locus proceed. It changes the branch-space of others. This is why institutional trust requires stronger recoverability than ordinary low-stakes trust.

Machine mediation intensifies the same structure.

A machine system can carry trust through data, labels, architecture, objective functions, validation metrics, thresholds, interfaces, dashboards, rankings, classifications, recommendations, and automated actions. It may reduce human recalculation at scale. It may detect patterns, maintain logs, compare cases, monitor drift, flag anomalies, support triage, or enforce consistency. These are real advantages.[11]

But machine mediation can also widen the distance between relation and reliance.

The user may see an output without seeing the data path. The affected locus may experience a decision without seeing the model. The institution may rely on a dashboard without understanding the abstraction. The developer may optimise a proxy rather than the target. The evaluator may test a benchmark rather than deployment relation. The machine interface may present confidence while residual remains hidden.

Machine trust is especially prone to carrier reification because outputs look operationally complete. A score, label, rank, recommendation, or generated text appears as a finished object. It can be copied, stored, cited, integrated, automated, and acted upon. The system's fluency or precision may reduce audit even when the carrier path is opaque. The output then carries more reliance-strength than the recoverable conditions support.

The correction is not to reject machine systems in favour of human intuition. Human judgement can also be opaque, biased, status-driven, and poorly audited. The correction is to subordinate both human and machine carriers to relation-preservation. The question is not whether a human or machine produced the output. The question is whether the trusted relation, carrier path, closure demand, residual, threshold, safeguard, and update condition remain recoverable.

Machine systems may support trust where the domain has already been sufficiently formalised for the relevant distinctions, carrier paths, thresholds, safeguards, and update conditions to remain recoverable through machine-mediated use. They may help maintain audit logs, detect inconsistent category use, flag missing evidence, surface drift, enforce review dates, compare decisions, and prevent residual from being forgotten. They may strengthen trust by making some carrier paths more visible.

But where the domain is under-formalised, machine systems scale opacity. They inherit categories whose grounding is unclear. They treat historical records as training signal. They compress institutional misclosures into model behaviour. They produce outputs that institutions treat as neutral because they are technical. They increase the distance between affected loci and decision authority. They make weak trust carriers more portable.

Institutional and machine mediation therefore create the bridge to formalisation.

Local trust can often remain tacit because update is near relation. If a surface fails, the body feels it. If a conversation misfires, the participants can repair. If a familiar path changes, the user sees obstruction. But when trust is mediated by institution or machine, the relying locus may not see the target relation, and the affected locus may not see the carrier path. Update is no longer naturally close to failure.

Reduced audit then becomes load-bearing.

A word, category, record, role, procedure, certificate, model, interface, or output may now classify, govern, allocate, exclude, automate, certify, or justify action. The trust relation is no longer merely a convenience. It has become authority-bearing. In that condition, tacit trust is insufficient.

The conditions of trust must become recoverable.

---

13. Formalisation: when reduced audit must become recoverable

Formalisation begins where reduced audit becomes consequential authority.

It is not the rejection of trust. It is the discipline required when trust has scaled beyond local, reversible, tacit, and easily updateable relation. A finite system must continue to rely on carriers. It cannot personally audit every person, record, procedure, institution, model, or machine output from first principles. But where those carriers classify, govern, allocate, exclude, automate, certify, or justify action, the conditions under which they carry reliance must become recoverable.

Trust remains tacit where relation is local, low-stakes, reversible, and near update. A body does not need a formal protocol to trust a familiar surface. A conversation does not need formalisation for every word. A friend does not need an audit packet for every ordinary reliance. Informal trust is not inferior by default. It is often the appropriate form where consequence is limited and repair is close.

Formalisation becomes necessary when any of these conditions fail.

It becomes necessary where the relying locus and affected locus diverge. It becomes necessary where the carrier is opaque. It becomes necessary where action is high-consequence, difficult to reverse, coercive, automated, institutional, or durable. It becomes necessary where records persist beyond context. It becomes necessary where trust is borrowed through credentials, certificates, procedures, models, or machine outputs. It becomes necessary where questions are treated as disloyalty, where residual is hidden behind confidence, or where carrier status replaces trust ground.

Formalisation makes the reduced-audit path inspectable. It asks: what relation is being trusted; for which locus; under what availability condition and through what boundary or interface; through what carrier; under what abstraction and evidence or record path where relevant; against what comparator; under what closure demand; for what purpose, scale and horizon; with what residual; at what threshold; under what safeguard; and subject to what update condition? The point is not to abolish reliance. The point is to prevent reliance-strength from exceeding the conditions that earn it.

When trust carriers become authority carriers, this burden becomes unavoidable.

A record used for memory may need little formalisation. A record used to deny benefit, impose sanction, or justify intervention requires more. A model used for exploration may need less formalisation than a model used for automated classification. A credential used as a conversational signal may need less formalisation than a credential used as a legal threshold. A certificate used to indicate process completion may need less formalisation than a certificate used as proof of substantive adequacy. A procedure used internally may need less formalisation than a procedure that exposes an affected locus to irreversible consequence.

Formalisation therefore scales with consequence, mediation, opacity, authority, and residual.

It also protects trust itself. Without formalisation, trust may be demanded where it is no longer legitimate. Institutions may ask for public confidence without exposing carrier paths. Vendors may ask users to trust machine outputs without exposing deployment residual. Professionals may ask clients to trust expertise without preserving domain limits. Public bodies may ask citizens to trust decisions without showing affected-locus preservation. Such demands produce cynicism because they ask for reduced audit where recoverability is absent.

Formalisation can restore conditions for legitimate reliance by making trust less dependent on mood, reputation, status, or institutional reassurance. It turns hidden carrier dependence into auditable path. It allows trust to be narrowed rather than abandoned. It allows affected loci to challenge the right relation. It allows institutions to distinguish procedure from warrant. It allows machines to inherit only what has been made auditable enough to inherit.

But formalisation must not become a new trust idol.

A completed form is not trust ground by itself. A certificate is not trust ground by itself. A compliance mark is not trust ground by itself. A machine-readable schema is not trust ground by itself. Formalisation is itself a carrier of prior audit. It must not be reified into authority. It remains legitimate only where it preserves relation, residual, threshold, safeguard, and update rather than replacing them with formal completeness.

The handoff to the next paper is therefore precise.

Trust explains why formalisation is needed. Formalisation explains how load-bearing trust carriers become recoverable when tacit reliance is no longer enough.

The requirement is not total proof.

The requirement is recoverable reliance.

---

14. Conclusion: Trust without possession

Trust begins where finite systems must continue without total possession.

A finite observer-system cannot reopen every dependency before every act. It cannot revalidate every surface, breath, memory, word, tool, source, record, institution, model, or machine output at full resolution each time it moves. It must reduce audit. It must allow some relations to become dependable enough for continuation while residual remains.

Trust is the reliance-form of that reduction.

It is not first belief, confidence, loyalty, moral virtue, reputation, or warrant. These may accompany trust, enrich it, distort it, or become its later social forms. The primitive function is

reliance closure under residual: the reduction of relational uncertainty to satisfactory dependence for a locus under a closure demand.

Audit is the counterfunction trust reduces. Trust backgrounds relation for action. Audit foregrounds relation for correction. A finite system needs both. Without trust, branch-space becomes unnavigable. Without audit, reduced scrutiny hardens into capture. Trust is legitimate where audit can become latent without becoming impossible.

This is why trust must be indexed. A locus trusts a relation, often through a carrier, under a closure demand, for a purpose and horizon, with residual, at threshold, under safeguard where relevant, and subject to update. Remove those conditions, and trust becomes a floating noun. It can then be demanded as loyalty, sold as reputation, certified as status, invoked as authority, or treated as transferable warrant.

Vicarious trust makes finite social and institutional life possible by allowing one locus to borrow closure-work from another. It enables expertise, medicine, science, law, governance, transport, education, infrastructure, and technology. But borrowed closure carries residual forward. The trusting locus depends not only on the target relation but on the carrier's competence, domain, evidence path, abstraction, incentives, safeguards, and update conditions.

Where borrowed closure crosses responsibility-bearing relations, trust backgrounds but does not abolish their provenance. Upstream responsibility remains local to the consequential transformations, retained or maintained conditions, and recoverable field-conditioning relations that shaped the relied-upon relation; the relying locus acquires only the responsibility attributable to its own later distinction-conditioned contribution where the discriminator is met. Trust can therefore reduce audit without exporting responsibility.

Trust carriers make that borrowing portable. Persons, roles, records, credentials, institutions, procedures, models, certificates, interfaces, and machine outputs allow reliance-strength to travel. They are necessary. They are also dangerous. The carrier carries; it does not consecrate. Where carrier status replaces relation-preservation, trust becomes reified.

Manipulation, trust laundering, and betrayal exploit this architecture. Manipulation captures threshold. Trust laundering moves reliance-strength through carriers without grounding. Betrayal exploits the reduced audit trust created. These failures are not accidents around trust. They arise from trust's function: it works by lowering scrutiny.

Repair is possible, but repair is not reassurance. It must restore the conditions under which reduced audit becomes legitimate again. Sometimes this restores strong trust. Sometimes it restores narrower trust. Sometimes it produces disciplined distrust. The aim is not confidence. The aim is calibrated reliance.

Institutional and machine mediation make this discipline more urgent because they separate relying locus, affected locus, carrier, target relation, and update path. A local trust failure may be corrected near the relation. A mediated trust failure may scale through records, categories, certificates, policies, dashboards, outputs, and automated actions. Where trust carriers become authority carriers, tacit trust is no longer enough.

That is where formalisation begins.

Formalisation is not the rejection of trust. It is what trust requires when reduced audit becomes load-bearing. It makes recoverable the relation, locus, carrier, closure demand, residual, threshold, safeguard, and update condition under which reliance-strength is claimed.

The governing rule remains:

Trust must not carry reliance-strength beyond the recoverable conditions under which the trusted relation, carrier, or closure process satisfies the relying locus's closure demand.

This rule protects trust from both cynicism and idolatry. It refuses the fantasy that finite systems can live without trust. It also refuses the fantasy that trust, once given, becomes a substance that can travel without relation.

Trust is necessary because finite systems cannot continuously audit total relation.

Trust is dangerous because its function is to reduce audit.

Trust is legitimate where that reduction remains indexed, proportionate, safeguarded where necessary, and answerable to update.

Trust is not possession.

It is reliance under residual.

It is the finite permission to proceed without pretending the world has been finally closed.

Endnotes

[1] Wittgenstein's On Certainty provides a philosophical precursor for the point that justification and rechecking cannot be indefinitely reopened before action: at some point practical activity supplies the background against which further doubt or verification is possible. Paper 5 uses a more primitive finite-locus trust relation and does not inherit Wittgenstein's language-game epistemology or treat practical continuation as epistemic closure.

[2] Luhmann's Trust and Power treats trust as a mechanism by which social complexity is reduced enough for action. The branch-space economy developed here is broader and more primitive: it is indexed to a finite locus, closure demand, residual, threshold, audit, and update, and it applies beyond Luhmann's social-systems setting. No systems-theoretic machinery is inherited.

[3] Sperber et al.'s account of epistemic vigilance provides a communication-specific specialisation of the rechecking function: recipients evaluate communicated information and its source rather than accepting testimony without scrutiny. Primitive audit here is broader. It is relation-revalidation whenever reduced audit must be reopened, whether or not the relation is communicative, propositional, or human.

[4] Habets et al. use 'basic trust' for the certainty expressed in skilled, often unreflective and habitual engagement with a living environment. This is a close human ecological comparator for embodied and environmental trust. Paper 5 generalises below urban or human phenomenology and locates the relation in retained, updateable reliance under residual rather than in a specifically experiential account.

[5] Gambetta's influential formulation treats trust as a context-sensitive expectation concerning another agent's behaviour and gives uncertainty and probability a central role. Paper 5's trust threshold is not inherently probabilistic and is not limited to agent-to-agent relations; Gambetta is used as a comparator for thresholded reliance under uncertainty, not as a technical dependency.

[6] Alfano uses the exact term 'vicarious trust' for mediated or transitive interpersonal trust through another trusted person, with delegation and ratification conditions. Hardwig separately shows how modern epistemic practice depends on trust in others where an individual cannot reconstruct all relevant evidence and argument locally. Paper 5 generalises both relations into outsourced closure-work through persons, roles, institutions, records, procedures, instruments, models, and machines, while preserving residual and refusing automatic warrant transfer.

[7] Johansson and Baldvinsdottir use 'trust carrier' in an organisational and management-accounting setting, where an actor can carry trust within performance-evaluation relations. The prior term is therefore acknowledged. Stewart's trust-transfer model separately shows how trust concerning one organisation can influence trust in another through perceived association and linkage. Paper 5 generalises carrier status beyond persons, organisations, and accounting routines to any person, role, object, record, credential, procedure, institution, instrument, model, interface, certificate, standard, archive, platform, or machine output through which reliance-strength travels. It also separates the portability of reliance-strength from Responsibility Provenance: a carrier may transmit or amplify reliance without transferring the causal

provenance of the responsibility-bearing transformations that produced or shaped what is carried.

[8] Gilly independently uses 'trust laundering' for the conversion of credibility earned by a verification mark or trusted channel into credibility for content or conduct that did not itself earn that standing. The overlap with Paper 5's broader transfer-of-reliance-strength failure is material, but the external treatment is a recent, platform-specific preprint and no developmental dependence has been established. It is recorded as an examined comparator and independent-convergence candidate, not as inheritance.

[9] Baier, Hawley, and Walker are important adversarial comparators for the philosophical boundary around mature interpersonal trust. Baier makes vulnerability and the possibility of betrayal central; Hawley develops a commitment-based account and rejects identifying trust with mere reliance. Walker's account of moral repair places damaged trust within normative relations in which persons are responsible to one another and asks what is required to reconstruct those relations after wrongdoing. Paper 5 intentionally uses 'trust' for a lower reliance-closure function and then separately recovers betrayal, moral responsibility, and repair as richer agency- and norm-sensitive structures. The difference is therefore explicit and scope-level rather than an attempt to displace those mature accounts.

[10] Gillespie and Dietz develop a systemic, multilevel framework for organisational trust repair after organisation-level failure. Their account is a strong specialisation of the claim here that repair requires change in the conditions that generated or sustained failure rather than reassurance alone. Paper 5 states the relation more generally: legitimate repair must restore, narrow, or refuse reduced audit by recovering the failed relation, residual, threshold, safeguard, provenance, and update conditions. The organisational framework is therefore an examined comparator rather than a dependency.

[11] Lee and See provide the canonical human-automation treatment of trust as a basis for appropriate reliance, with calibration to the automation's capabilities and context. Paper 5 generalises the calibration problem into carrier, residual, threshold, safeguard, update, affected-locus, and formalisation disciplines, and does not inherit a human-factors model of trust as its primitive definition.

References

Alfano, M. (2016). The topology of communities of trust. *Sotsiologicheskoe Obozrenie*, 15(4), 30-56. https://doi.org/10.17323/1728-192X-2016-4-30-56

Baier, A. (1986). Trust and antitrust. *Ethics*, 96(2), 231-260. https://doi.org/10.1086/292745

Gambetta, D. (1988). Can we trust trust? In D. Gambetta (Ed.), *Trust: Making and breaking cooperative relations* (pp. 213-237). Basil Blackwell.

Gillespie, N., & Dietz, G. (2009). Trust repair after an organization-level failure. *Academy of Management Review*, 34(1), 127-145. https://doi.org/10.5465/amr.2009.35713319

Gilly, T. (2026). Reviewed for quality and security: Advertising injection through a verified connector, trust laundering, and the limits of directory verification. *SSRN*. https://doi.org/10.2139/ssrn.7130058

Habets, D., Rietveld, E., Kiverstein, J., & Denys, D. (2024). Trusted urban places. *Adaptive Behavior*, 32(5), 421-438. https://doi.org/10.1177/10597123241229735

Hardwig, J. (1991). The role of trust in knowledge. *The Journal of Philosophy*, 88(12), 693-708. https://doi.org/10.2307/2027007

Hawley, K. (2014). Trust, distrust and commitment. *Noûs*, 48(1), 1-20. https://doi.org/10.1111/nous.12000

Johansson, I.-L., & Baldvinsdottir, G. (2003). Accounting for trust: Some empirical evidence. *Management Accounting Research*, 14(3), 219-234. https://doi.org/10.1016/S1044-5005(03)00045-3

Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50-80. https://doi.org/10.1518/hfes.46.1.50_30392

Luhmann, N. (1979). *Trust and power: Two works* (H. Davis, J. Raffan, & K. Rooney, Trans.). John Wiley & Sons.

Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*, 25(4), 359-393. https://doi.org/10.1111/j.1468-0017.2010.01394.x

Stewart, K. J. (2003). Trust transfer on the World Wide Web. *Organization Science*, 14(1), 5-17. https://doi.org/10.1287/orsc.14.1.5.12810

Walker, M. U. (2006). *Moral repair: Reconstructing moral relations after wrongdoing*. Cambridge University Press.

Wittgenstein, L. (1969). *On certainty* (G. E. M. Anscombe & G. H. von Wright, Eds.; D. Paul & G. E. M. Anscombe, Trans.). Basil Blackwell.