GOT-V1-3c-8

The Epistemic Method

Core Principles of Relational Audit Under Finite Conditions

Abstract

The Epistemic Method is a discipline of relational audit for finite observer-systems and the Soft Horizons through which they claim, classify, model, decide, rely, and act. It addresses the conditions under which claims, observations, records, classifications, explanations, institutional categories, machine outputs, action decisions, reliance closures, trust carriers, and provisional closures preserve the relations they purport to carry.

It is not a theory of everything. It is not a substitute for science, law, judgement, ethics, domain expertise, or institutional responsibility. It is not a machine-governance protocol or a general decision procedure. It is a method for preventing finite claims and the closures that support them from acquiring more strength than their recoverable conditions warrant.

The method begins from the finite-observer condition. No finite observer-system possesses reality as total state. A system encounters relation through availability condition, boundary, interface, scale, horizon, discrimination regime, abstraction, registration, retention, redeployment, comparator pressure, closure demand, and residual limit. It then operates through Soft Horizons: provisional closures that make relation usable without making it total. A claim made under such conditions may be useful, measurable, lawful, institutionalised, formalised, encoded, actionable, or widely accepted. None of those features is sufficient by itself. A claim earns strength only insofar as it preserves the relation it asserts under recoverable conditions.

The central object of the method is relational uncertainty. Relational uncertainty is uncertainty about whether a claim, closure, reliance relation, or Soft Horizon preserves the relation it purports to carry for a specified locus, under a specified availability condition, under the closure demand being made. It is not reducible to subjective doubt, missing information, probability deficit, lack of confidence, or ordinary ignorance. It can persist where confidence is high, where records exist, where measurements are precise, where procedures have been followed, where models perform, where provisional closure appears sufficient, where trust carriers are accepted, or where institutions have acted.

The method therefore audits claim-strength, and where a claim is relied upon, it audits reliance-strength. It asks where authority has attached. A familiar category may conceal abstraction failure. A record may be treated as warrant by default. A credential, role, procedure, institution, model, certificate, stewarded release, or machine output may be trusted as if the closure-work behind it had already satisfied the present closure demand. A legal or institutional term may travel beyond the conditions that grounded its use. A machine output may inherit a prior misclosure and scale it. An action taken under threshold may later be narrated as if it had closed the uncertainty under which it was taken.

Some claims impose relation-specific audit burdens. A claim about power must recover the Counterfactual Inclusion Field physically supported by the distinction-conditioned organisation of the specified Information Processing Locus under the relevant reference, comparator, thermodynamic, abstraction, and temporal conditions, rather than infer power

from outcome, force, authority, or externally imagined alternatives. A claim about responsibility must recover the consequential branch relation attributable to the specified locus under the actual conditions of the relation and distinguish Responsibility Provenance from Responsibility Warrant. Where application is claimed, the audit must recover the relevant Counterfactual Exclusion. Where responsibility is claimed through non-application, retention, maintenance, or material cross-locus field-conditioning without application, it must recover the comparator-stable difference made to the relevant Counterfactual Inclusion Field. These are conditional audit paths within the method. They do not become universal fields for claims that do not engage those relations.

The governing rule is:

No claim earns strength beyond the recoverable conditions under which it preserves the relation it purports to carry. Where the claim is relied upon, no reliance closure earns strength beyond the recoverable conditions under which the trusted relation, carrier, or closure-process satisfies the relying locus's closure demand.

The Epistemic Method does not abolish uncertainty, and it does not abolish Soft Horizons. It prevents uncertainty and provisional closure from being hidden inside authority. It permits finite systems to claim, classify, model, decide, and act where total closure is unavailable, but only by preserving the conditions under which claim-strength is earned, limited, revised, and exposed to correction.

---

1. Introduction: claim strength under finite conditions

Finite observer-systems make claims through Soft Horizons before they possess total closure.

This is not a defect that can be eliminated by better discipline, stronger institutions, more complete data, more exact models, more formal language, or more powerful machines. It is the ordinary condition of finite relation. An organism responds before possessing the total state of its environment. A scientific practice models before it exhausts its target. A court decides under evidential limits. A medical system treats before every causal pathway is known. A public body acts through records and categories that never contain the full relation they govern. A machine classifier produces outputs under training, architecture, threshold, and deployment constraints.

The question is therefore not whether finite systems should make claims under incomplete conditions. They cannot avoid doing so. The question is how those claims should carry strength.

A claim may be weak, strong, provisional, closed for a stated purpose, actionable under threshold, corrigible under update, or false. It may function as a Soft Horizon where it closes enough relation for a specified use while residual remains. It may be expressed in natural language, formal language, measurement, classification, legal standard, institutional record, model output, diagram, score, decision notice, protocol, or machine-readable field. Its form does not settle its strength.

A claim does not become strong merely because it is grammatically clear, technically expressed, measurable, recorded, repeated, official, legal, automated, professionally endorsed, statistically framed, institutionally adopted, or acted upon.

The strength of a claim depends on whether it preserves the relation it purports to carry under the conditions in which that claim is made, retained, transferred, used, and revised.

This is the problem the Epistemic Method addresses. It is a method of relational audit under finite conditions. It asks what relation is being asserted or relied upon, for what locus, under what availability condition, under what closure demand, through what access path, at what scale, over what horizon, under what abstraction, by what evidence path, record path, or carrier path, against what comparator field, with what residual, under what reliance limit where reliance is carried, under what action threshold, constrained by what safeguard, and answerable to what update.

The method inherits a layered architecture rather than collapsing its relations into one test. Relational availability and horizon grammar establish what relation can become available and under what closure demand. Unknown discipline structures residual and action before closure. Power identifies the physically supported distinction-conditioned branch-changing potential carried by an Information Processing Locus through its Counterfactual Inclusion Field. Responsibility localises the consequential branch relations attributable to loci — including realised transformation, responsibility-bearing non-application, retained or maintained conditions, and material cross-locus field-conditioning — and separates Responsibility Provenance from Responsibility Warrant. Trust explains reliance closure and carrier-mediated reduction of audit. Formalisation preserves load-bearing claim and reliance paths. Constitutional governance keeps formalisation corrigible without allowing stewardship, release, adoption, authority, or implementation to become epistemic ground. The Epistemic Method brings these relations into one audit discipline while preserving the discriminator proper to each.

The method does not begin by assuming that all claims are equally uncertain. That would dissolve the distinctions audit requires. Some claims are stronger than others. Some relations are well preserved under stable conditions. Some measurements are robust. Some classifications survive comparator pressure. Some records carry strong provenance. Some models perform reliably within a defined envelope. Some decisions are justified under urgency, precaution, reversibility, or duty even though explanatory closure remains incomplete.

Nor does the method deny closure. Closure may be earned locally. A claim may close for a specified purpose under specified conditions. A proposition may be proven within a formal system. A measurement may establish a value within tolerance. A legal decision may close a procedural matter. A biological attribution may be sufficiently grounded for the frame under review. An engineering claim may be warranted within an operating envelope. A scientific model may be reliable for prediction at one scale while remaining residual-bearing at another.

The method disciplines the conditions under which closure, Soft Horizon, strength, action, and residual are claimed.

The central danger is that claim-strength and reliance-strength can migrate. Claim-strength can detach from the relation-preserving path that earned it and attach itself to easier carriers: a name, record, procedure, institution, model, metric, professional convention, legal category, machine output, compliance statement, Soft Horizon, or action already taken. Reliance-strength can migrate in the same way when a person, role, credential, archive, certificate, institution, stewarded release, model, procedure, interface, or machine output is trusted because it appears to inherit closure-work.

Once this happens, a claim or carrier may continue to circulate with authority even when the relation it originally purported to carry has weakened, shifted, disappeared, or was never established.

A category may begin as a disciplined abstraction but later operate as a label whose cut is no longer visible. A record may arise when the result of bounded registration is retained, but later

function as warrant. A model may begin as approximation but later be treated as target possession. A risk term may begin as a structured relation among possible harm, evidence, uncertainty, horizon, and threshold, but later become a trigger for action without preserving that structure. A machine score may begin as a statistical output under a defined training and deployment regime, but later be treated as if it directly expressed the target relation. A legal or administrative term may begin as a judgement under conditions but later become a reusable instrument of closure.

In each case, a Soft Horizon that was usable under one relation hardens into authority for another.

The Epistemic Method exists to prevent this detachment.

It does so by forcing claim-strength and reliance-strength back to recoverable conditions. A claim must not travel without its relation. A relation must not travel without its reference conditions. A category must not travel without its abstraction cut. A record must not travel as warrant by default. A trust carrier must not travel as trust ground by default. A model must not travel as ontology. A threshold must not travel as proof. A certificate, compliance statement, stewarded release, institutional process, or machine output must not travel as closure. A residual must not disappear because a verdict, decision, output, reliance relation, or action has been produced.

The method therefore occupies a disciplined middle position between paralysis and overclosure.

A system that refuses all claim-making until total closure is available becomes unusable. It cannot respond, decide, learn, govern, treat, classify, investigate, or act. It mistakes finite discipline for impossible omniscience. But a system that allows action, record, procedure, confidence, or convenient provisional closure to become closure becomes dangerous. It converts unresolved relation into authority. It hides remaining uncertainty precisely where affected loci, future auditors, comparator cases, and update conditions most need to see it.

The method permits claim-making and reliance under finite conditions, but requires the incompleteness of those conditions to remain visible where it matters. It permits Soft Horizons, but only where relation, locus, availability condition, closure demand, residual, scope, tolerance, carrier path where reliance is inherited, reliance limit where relevant, threshold where relevant, safeguard where action follows, and update remain recoverable. It permits action under residual, but only by threshold, safeguard, and update. It permits records, but separates registration from warrant. It permits trust carriers, but separates carrier status from trust ground. It permits models, but separates model performance from target possession. It permits institutional decision, but separates procedure from relation-preservation. It permits machine assistance, but separates output from adequacy.

This is not a substitute for domain judgement. The method does not decide scientific theory, legal consequence, ethical priority, political value, engineering design, medical intervention, or institutional mandate by itself. It does not remove the need for specialised standards, expert knowledge, public deliberation, professional judgement, empirical testing, legal reasoning, or practical responsibility. It operates beneath and across those practices by asking whether the claim path preserves what the claim asserts strongly enough for the strength being claimed.

The Epistemic Method is therefore a discipline of finite claim-strength. It does not promise possession of reality as total state. It requires claims and the Soft Horizons that carry them to remain answerable to the relations they carry.

2. The method problem: mislocated claim-strength

The problem that requires method is not ignorance alone.

Ignorance can be simple non-possession. A system does not know something. A relation has not entered availability. A measurement has not been made. A proof has not been found. A record has not been retained. A comparison has not been run. A model has not been constructed. An event has not been observed by the relevant locus under the relevant conditions. Such absence may be consequential, but it is not the deepest problem addressed here.

The deeper problem is mislocated claim-strength, including mislocated reliance-strength.

Mislocated claim-strength occurs when a claim carries more force than its relation-preserving conditions warrant. Mislocated reliance-strength occurs when trust attaches to carrier status rather than recoverable relation-preservation. The claim or carrier may appear strong because it is familiar, named, measured, recorded, repeated, institutionalised, legally framed, scientifically styled, statistically expressed, machine-generated, professionally endorsed, procedurally compliant, credentialed, certified, stewarded, or already acted upon. Its apparent strength may be located in the form of the claim, closure, or carrier rather than in the relation it purports to carry.

A familiar category may conceal a weak abstraction. A measurement may conceal scale error. A record may conceal uncertain provenance. A model may conceal proxy substitution. A legal standard may conceal a missing affected-locus relation. An institutional procedure may conceal discretion. A machine output may conceal inherited labels. A public-interest claim may conceal moral preference. A risk classification may conceal threshold manipulation. A scientific explanation may conceal comparator weakness. A governance decision may conceal residual burial. Each is a possible reified Soft Horizon: a provisional closure whose remaining conditions no longer travel with it.

In each case, the problem is not merely that the claim might be false. A claim can be false for many reasons. It may be based on bad evidence, invalid inference, mistaken identity, insufficient measurement, corrupt data, poor reasoning, or deliberate deception. Mislocated claim-strength is more specific. It occurs when the system cannot accurately say where the strength of the claim is coming from.

The claim may gain strength from the wrong place.

It may gain strength from authority rather than evidence. It may gain strength from record-presence rather than warrant. It may gain strength from procedure rather than relation. It may gain strength from institutional repetition rather than comparator survival. It may gain strength from machine execution rather than grounded abstraction. It may gain strength from credential, certificate, stewarded status, compliance mark, or carrier familiarity rather than recoverable closure-work. It may gain strength from action already taken rather than closure earned. It may gain strength from public agreement rather than target constraint.

This is why method is needed. The method forces strength to return to the claim path.

A claim path is the recoverable sequence by which relation becomes available under stated conditions, is classed, registered, retained, formulated, provisionally closed under a stated closure demand, compared, residualised, thresholded, acted upon, and updated. Where a claim depends on inherited closure-work, the claim path includes a reliance path: the person, role, record, credential, archive, institution, procedure, model, interface, certificate, stewarded release, or machine output through which reliance is mediated.[1]

The sequence may be short or long. It may involve sensory registration, instrument calibration, data capture, witness evidence, classification, model construction, legal reasoning, institutional procedure, machine training, statistical inference, professional judgement, or vicarious closure through another carrier. It may be distributed across many systems and time periods. What matters is that the claim's strength and any reliance-strength remain tied to the conditions by which the relation was carried.

In a scientific case, the claim path may involve instruments, measurements, controls, theories, error estimates, replication, modelling assumptions, predictive success, and comparator hypotheses. The method asks what relation the scientific claim asserts, which instruments made it available, which abstraction made it measurable, which scale and horizon the claim covers, which alternatives were excluded, which residual remains, and what would force revision.

In a legal case, the claim path may involve evidence, admissibility, witness credibility, statutory categories, precedent, burden of proof, procedural standard, judgement, remedy, and appeal. The method asks what relation the legal claim asserts, which affected locus is being judged or protected, what category cut has been applied, what evidence path supports the finding, what alternatives remain live, what residual is carried, and whether action has been confused with epistemic closure.

In an institutional case, the claim path may involve policy language, records, eligibility categories, thresholds, discretion, professional judgement, oversight, appeal, and archive. The reliance path may involve roles, offices, credentials, delegated authority, procedural status, compliance marks, historical records, or institutional memory. The method asks whether the institution has preserved the relation its category claims to carry, or whether administrative form has replaced relation-preservation and carrier trust has been mistaken for ground.

In a machine case, the claim path may involve data provenance, training labels, architecture, objective function, evaluation set, threshold, deployment context, monitoring, user interface, and human review. The reliance path may involve trust in the dataset, labelling process, benchmark, validation regime, vendor assurance, interface, model family, deployment policy, or machine output. The method asks whether the output preserves the target relation or merely inherits, compresses, and scales prior classifications under conditions no longer visible.

The method does not reduce these cases to one domain. It identifies a common audit structure that prevents claims from overreaching their conditions.

This common structure is needed because finite observer-systems must operate through incomplete but necessary provisional closures. A finite system cannot wait until total knowledge is available before acting. Organisms respond before full identification. Institutions decide before all facts are known. Medical systems treat under uncertainty. Scientific practices use models before ontology is complete. Courts decide under evidential limits. Machines classify under distributional constraint. Governance systems act through partial records.

A method that required total closure before any claim or action would be unusable.

A method that allowed Soft Horizons to harden into certainty would be dangerous.

The Epistemic Method therefore permits claim-making under finite conditions. It permits provisional model use, institutional decision, legal judgement, machine assistance, scientific explanation, practical action, and other Soft Horizons where closure is incomplete. But it requires that incompleteness remain visible and operational. Residual may not disappear because action is taken. Threshold may not become proof. Record may not become warrant. Model success may not become ontology. Compliance may not become closure.

The method exists because finite observers must act, classify, model, explain, decide, and revise through Soft Horizons under conditions where closure is partial, indexed, and vulnerable to misclosure or reification.

Finite observers must act through provisional closures. The method cannot become paralysis.

Finite observers classify. The method must expose abstraction.

Finite observers model. The method must distinguish target, model, and residual.

Finite observers explain. The method must separate explanation from naming, prediction, and record.

Finite observers decide. The method must separate action threshold from closure.

Finite observers revise. The method must preserve update conditions.

Mislocated claim-strength is the failure in which these distinctions collapse. The name becomes the relation. The record becomes warrant. The model becomes target. The action becomes proof. The carrier becomes ground. The institution becomes truth. The method begins by returning each strength claim to the conditions that earn it.

---

3. The method object: relational uncertainty

The object of the Epistemic Method is relational uncertainty.

This term should not be weakened into ordinary doubt. Relational uncertainty is not simply uncertainty in the mind of a subject. It is not a feeling of hesitation. It is not a probability deficit alone. It is not merely the absence of information. It is not identical with ignorance. It is not dissolved by confidence, consensus, measurement, procedure, institutional action, or machine output.

Relational uncertainty is uncertainty about whether a claim, closure, reliance relation, or Soft Horizon preserves the relation it purports to carry under the conditions in which it is used.[2]

The uncertainty is relational because the target does not arrive as total state. It becomes available to a locus through boundary, interface, scale, horizon, discrimination regime, and abstraction; it may then be registered, retained, redeployed, compared, and evaluated under a closure demand. The same target relation may be available to one locus and unavailable to another, measurable at one scale and residual at another, closed for action and unclosed for explanation, robust under one abstraction and fragile under another. The question is not simply whether something is known. The question is whether the relation claimed is preserved strongly enough for the closure demand being made.

Relational uncertainty can persist even where confidence is high. A decision-maker may be confident because a category is familiar, a record is present, a model is precise, a professional role is trusted, or an institution has acted. Confidence may be sincere. It may even be practically

useful. But confidence is not the same as relation-preservation. A system may be confident in the wrong carrier.

Relational uncertainty can persist where records exist. A record establishes that some content was registered and retained, or that an assertion was recorded, under some conditions. It does not automatically establish that the target relation was preserved, that the record is current, that provenance is adequate, that abstraction was sound, or that the record may support the present closure demand. Record-presence is often useful. It is not warrant by default.

Relational uncertainty can persist where measurements are precise. Precision concerns the stability or granularity of a measurement under a measurement regime. It does not by itself settle whether the measured variable preserves the target relation, whether scale is appropriate, whether the abstraction is valid, or whether the measurement supports the action being taken from it. False precision is one of the strongest carriers of mislocated claim-strength.

Relational uncertainty can persist where procedures have been followed. Procedure may constrain discretion, preserve fairness, create reviewability, and improve decision quality. But procedural completion is not relation-preservation by itself. A procedure may be followed while excluding the relevant relation, hiding residual, preventing affected-locus challenge, or treating a threshold as closure.

Relational uncertainty can persist where models perform. A model may predict well within one domain while failing in another. It may optimise a proxy. It may inherit historical labels. It may compress institutional misclosure. It may preserve correlation without preserving explanation. It may support triage but not sanction, monitoring but not automation, prediction but not ontology. Model performance is therefore not target possession.

Relational uncertainty can persist where trust carriers are accepted. A person, role, credential, institution, instrument, archive, certificate, method release, vendor tool, or machine output may reduce local audit. That reduction may be necessary. But the trust carrier must not be mistaken for the trust ground. The carrier's closure process, domain, evidence path, abstraction, threshold, safeguard, and update conditions must remain recoverable enough for the reliance-strength being claimed.

Relational uncertainty can persist where institutions have acted. Action often closes a branch. It creates records, consequences, and retrospective narratives. But action is not proof. A system may have had sufficient threshold to act under residual. It may not have had epistemic closure. The fact that a decision was necessary, lawful, urgent, or already taken does not erase the uncertainty under which it was taken.

The method therefore treats relational uncertainty as the condition that appears wherever claim-strength may exceed relation-preservation.

It asks a disciplined sequence of questions.

What relation is being carried?

For which locus is the relation evaluated?

How did the relation become available?

What closure demand is being made?

Through what boundary, interface, scale, and horizon is the claim operating?

What abstraction makes the relation usable?

What evidence path, record path, or carrier path supports it?

Where a power claim is being made, what Counterfactual Inclusion Field is physically supported by the distinction-conditioned organisation of the specified Information Processing Locus under the relevant reference and comparator conditions, and what comparator shows that the claimed potential is real?

Where a responsibility claim is being made, which consequential branch relation is attributable to which locus under the actual conditions of the relation? Where application is claimed, what Counterfactual Exclusion is recoverable? Where non-application, retention, maintenance, or cross-locus field-conditioning without application is claimed, what comparator-stable difference was made to the relevant Counterfactual Inclusion Field? What part of the attribution concerns Responsibility Provenance and what part concerns Responsibility Warrant?

What comparator field has it survived?

What residual remains?

What reliance limit applies where a carrier is trusted?

What threshold permits action if closure remains incomplete?

What safeguard constrains exposure?

What update condition would force revision?

These questions do not imply that every claim is weak. They imply that strength must be located. A claim may be strong. It may survive comparator pressure. It may carry minimal residual. It may close for a purpose. It may justify action. It may support reliance. But the method requires the source, scope, and limits of that strength to remain recoverable.

Relational uncertainty also protects against false equivalence. If final closure is unavailable, it does not follow that every formulation is equally live. Some formulations may be positively excluded. Some may be weak. Some may be better supported. Some may be strong enough for action but not explanation. Some may be strong within one scale but residual at another. The method does not flatten uncertainty. It structures it.

This is why the method inherits the discipline of the unknown but operationalises it as audit. The unknown is not mere absence. Residual is not a decorative caveat. A live formulation field is not an equality field. Weighting is not closure. Action is not certainty. The method converts these distinctions into claim-path discipline.

Relational uncertainty is therefore the method’s object because it is the point at which finite systems are most likely to overclaim. A claim may be useful, but not warranted for the demand being made. A record may be available, but not sufficient. A category may be familiar, but not well abstracted. A model may perform, but not preserve the target. A procedure may complete, but not ground the relation. A threshold may justify action, but not establish truth. A carrier may reduce audit, but not remove residual.

The method asks where that uncertainty has gone.

If it remains visible, the claim can be calibrated.

If it is hidden inside verdict, record, category, procedure, model, compliance, or action, the claim has become dangerous.

Relational uncertainty is not the enemy of method. It is the reason method is needed. The aim is not to eliminate it wherever it appears. The aim is to prevent it from being converted into unearned strength.

A finite system becomes stronger not by pretending that its claims escape relational uncertainty, but by making the uncertainty localised, limited, weighted, actionable where necessary, and answerable to update.

---

4. Indexed relation: relation, locus, availability, and closure demand

A claim cannot be audited until it states the relation it carries.

This is the first operational discipline of the method. The claim may be grammatical, meaningful, measurable, formal, legal, institutional, technical, machine-generated, widely accepted, or practically necessary. None of those features is enough. They may make the claim usable. They may allow it to enter communication, record, procedure, classification, modelling, action, or machine execution. They do not by themselves make the claim answerable to the relation it asserts.

A statement such as “this is safe,” “this is risky,” “this is evidence,” “this is observed,” “this is intelligent,” “this is reasonable,” “this is harmful,” “this is compliant,” “this model is accurate,” “this institution can be trusted,” “this certificate is sufficient,” “this method was followed,” or “this action is justified” is not yet a disciplined claim. It may be a meaningful sentence. It may be a professional judgement. It may be an institutional category. It may be a legal conclusion. It may be a model output. It may be an action threshold. But until the relation it carries is explicit, its strength cannot be audited.

The claim “this is risky” may carry a relation between possible future harm, probability, uncertainty, affected locus, severity, horizon, evidence, and action threshold. It may instead carry an institutional relation between category, policy, and intervention. It may carry a machine relation between proxy variables and score. It may carry a political relation between public anxiety and administrative response. It may carry a legal relation between foreseeable harm and duty. These relations are not interchangeable. If the operative relation is not stated, the term risk may borrow strength from one relation while operating in another.

The claim “this is observed” may carry a relation between available distinction, bounded registration, and the observing locus; where the registration is retained, it may also produce an observation record. It may also be used as if it carried truth, warrant, objectivity, or unmediated access. The method must prevent that movement. Observation establishes bounded registration under conditions. It does not automatically establish that the target relation was preserved, that any resulting record is warranted, or that the interpretation is correct.

The claim “this system is intelligent” may carry a relation between applied distinction and downstream trajectory selection for a locus. It may also be used to imply consciousness, understanding, agency, usefulness, human resemblance, or moral status. Those are different relations. A claim may be valid under one and invalid under another. The word intelligent cannot be allowed to move between them while retaining unearned authority.

The claim “this locus has power” must state the branch-changing relation being asserted. Primitive Power concerns the indexed potential of an Information Processing Locus to alter downstream branch-realisation through organisation-mediated discriminable distinctions under specified reference and comparator conditions. The relevant causal possibility field is the Counterfactual Inclusion Field: the downstream branch relations physically included within

that locus's distinction-conditioned potential under the conditions being claimed. Outcome magnitude, success, authority, possession, or an alternative imagined only by an exterior analyst does not establish that relation. Nor may the Counterfactual Inclusion Field be inferred from the observer-system's live formulation field. A formulation may remain epistemically live without being a physically included causal alternative for the IPL, while a physically supported alternative may exist without having been represented as a live formulation. The method must therefore recover the physically supported field proper to the Power claim and the part of that potential actually being asserted.

The claim "this locus is responsible" must likewise state the attribution relation being asserted. Responsibility Provenance concerns the recoverable consequential branch relations attributable to loci through distinction-conditioned organisation. These may include realised transformations, responsibility-bearing non-application, retained or maintained conditions, and material conditioning of another locus's Counterfactual Inclusion Field. Responsibility Warrant concerns the social, moral, legal, organisational, procedural, or institutional allocation of answerability, duty, burden, liability, supervision, stewardship, or decision ownership. Where application is claimed, the attribution must be capable of recovering the relevant Counterfactual Exclusion. Where responsibility is claimed without application through retention, maintenance, non-application, or cross-locus field-conditioning, the attribution must instead recover a comparator-stable difference to the relevant Counterfactual Inclusion Field. Role, authority, blame, liability, proximity, or outcome cannot be allowed to move between Provenance and Warrant without the conditions that earn the claimed attribution.[3]

The claim "this decision is reasonable" may carry a relation between evidence, standard, alternatives, proportionality, institutional competence, affected locus, and legal or administrative threshold. It may also function as a procedural label applied after a decision has already been made. The method asks which relation the word reasonable actually preserves, under what conditions, and for what closure demand.

Relation specification is therefore not definitional tidiness. It is the first guardrail against false strength.

Once the relation is stated, it must be indexed.

A relation is never available, unavailable, registered, retained, deployed, closed, residual, trusted, or acted upon in the abstract. It is evaluated for a locus, under an availability condition, under a closure demand, through some boundary and interface, at some scale and horizon. Without those conditions, relation becomes a floating noun and claim-strength begins to migrate.

The locus is the unit for which the relation is evaluated. It may be a person, organism, cell, instrument, institution, court, regulator, scientific practice, model, machine system, public body, social organisation, affected group, or coupled observer-network. The same relation may be available to one locus and unavailable to another; useful for one locus and harmful to another; closed for one locus's action and residual for another locus's explanation.

Locus discipline prevents silent transfer. A risk to an institution is not automatically a risk to the person governed by it. A distinction available to an analyst is not automatically available to the system under review. A population-level regularity is not automatically an individual classification. A machine-readable category is not automatically an affected-locus reality. A regulator's visibility is not automatically public warrant. Where locus moves silently, claim-strength detaches from the relation that earned it.

The affected locus must also be identified where consequence falls on a different locus from the one relying on the claim. An institution may rely on a record; a person may bear the error.

A platform may rely on a classifier; a user may bear the restriction. A hospital may rely on a triage score; a patient may bear the risk. A regulator may rely on a certificate; the public may bear the failure. A claim that is adequate for the relying locus's convenience may not be adequate for the affected locus's exposure.[4]

The availability condition states how the relation stands for the locus. The relation may be unavailable, weakly available, available in potential, coupled, available as contrast, registered, retained, deployed, residual-bearing, provisionally closed, or closed for purpose. These distinctions matter because claims often gain false strength by sliding across them. Existence becomes access. Access becomes observation. Observation becomes record. Record becomes warrant. Deployment becomes closure. Provisional closure becomes possession.

A complaint received by an institution establishes availability to that institution, not proof. A physical effect reaching a detector path establishes access or coupling; registration requires that the detector select a classed alternative into state. A data point retained as re-identifiable state may establish record-presence, not warrant. A model output is an output produced under stated model and deployment conditions, not target possession. A legal category may close procedure without closing the lived relation. A medical classification may support treatment without closing causal explanation. The method requires the claim to state which availability condition it occupies before strength is assigned.

The closure demand states what kind of closure is being asked of the relation. Recognition closure is not explanation closure. Measurement closure is not ontological closure. Legal closure is not scientific closure. Action closure is not truth closure. Compliance closure is not relation-preservation. A claim may be closed for one demand and residual for another.

A model may be adequate for ranking but not explanation. A record may be adequate for initiating investigation but not for sanction. A diagnosis may be adequate for precaution but not for final causal attribution. A legal finding may close a procedural matter while leaving historical or moral residual. A machine output may be adequate for triage but not for automation. Without closure-demand discipline, a locally sufficient claim expands into stronger authority than it has earned.

Boundary and interface state how relation entered. A relation may enter through perception, instrument, testimony, report, complaint, measurement, legal submission, machine output, institutional form, dataset, model, archive, or operational dashboard. The access path matters because each path admits some relations and excludes others. A form can receive only what its fields permit. A dataset carries only what prior systems recorded. A sensor measures only within its regime. A court sees through admissibility. A platform sees through interface and policy. An institution sees through category, record, and procedure.

Scale and horizon state the range over which the claim is being asserted. A short-term signal may not justify a long-term status. A local measurement may not support general policy. A population pattern may not settle an individual case. A model trained in one environment may fail in another. A risk may be immediate but not durable. A historical record may not describe current condition. A claim without scale and horizon can overextend even where its local relation is strong.

Indexed relation therefore gives the method its first complete audit unit. The claim must say: this is the relation being carried; this is the locus for which it is evaluated; this is how the relation became available; this is the closure demand being made; this is the boundary and interface through which it entered; this is the scale and horizon over which it is claimed.

This indexing does not guarantee that the claim is true. It makes the claim auditable.

A claim that cannot state its relation is not ready for strength.

A relation that cannot state its locus, availability condition, and closure demand is not ready for authority.

---

5. Abstraction, evidence, records, and carriers

Indexed relation is still not enough.

A finite system does not carry relation whole. It makes relation usable through abstraction. It selects, classes, partitions, measures, names, thresholds, compares, simplifies, models, records, and redeploys. That activity is necessary. Without abstraction, no observer-system could recognise, communicate, measure, decide, model, or act. Relation would remain unorganised variation.

The problem is not abstraction.

The problem is hidden abstraction.

A category is an abstraction cut. It treats some differences as relevant and others as background. It defines sameness and difference. It establishes inclusion and exclusion. It sets thresholds, tolerances, and edge conditions. It may compress many cases into one label. It may turn a continuous relation into a discrete status. It may convert lived relation into administrative category, legal standard, medical diagnosis, scientific variable, machine label, model feature, or governance metric.

Every such cut is a Soft Horizon. It closes enough relation for use while leaving something outside, below, above, beside, or unresolved. The abstraction may be strong. It may be the best available. It may be necessary for action. But it must be exposed.

A category such as risk must show what it treats as risk rather than uncertainty, harm, vulnerability, suspicion, liability, non-compliance, political concern, or operational inconvenience. A category such as harm must show what affected relation it treats as damaged, impaired, violated, restricted, or injured. A category such as evidence must show the difference among testimony, material evidence, measurement, record, inference, model output, allegation, and warrant. A category such as safety must show whether it concerns physical harm, system failure, legal compliance, reputational exposure, public reassurance, or acceptable residual.

Without abstraction discipline, a category may be mistaken for the relation. The system stops asking whether the cut preserves the target and begins asking only whether the case fits the cut. The category then becomes an authority object. It no longer appears as a finite closure under conditions. It appears as the thing itself.

This is especially dangerous in institutions and machines.

Institutions act through categories because they must. They cannot govern, allocate, adjudicate, inspect, certify, or remember without classing. But institutional categories often outlive the relations that first justified them. They become forms, records, eligibility rules, policy thresholds, sanctions, dashboards, and archive fields. Once a category becomes institutional memory, it may be treated as natural because the institution repeatedly acts through it.

Machines also act through abstraction. A training label, feature set, objective function, benchmark, threshold, embedding, score, or output class is a cut. It makes relation computable. But computability is not adequacy. A machine-readable distinction may be useful because it is

executable, not because it preserves the target relation. If the abstraction is hidden, the output appears technical while carrying an unexamined Soft Horizon.

The method therefore asks: what abstraction makes the claim usable? What category was applied? What metric was selected? What proxy stands in for the target? What threshold turns variation into status? What cases become equivalent? What edge cases reveal the cut? What affected relation may be excluded? What residual does the abstraction create?

Abstraction discipline then requires evidence, record, and carrier-path discipline.

A claim does not become strong merely because it has a record, a source, a metric, a certificate, a model output, a professional judgement, an institutional procedure, or a method-compliance form attached to it. Those may support the claim. They do not support it by default. The method must distinguish evidence path, record path, and carrier path before assigning strength.

The evidence path states what supports the claim and how that support was produced. Was the relation directly observed, instrumentally measured, reported, inferred, estimated, modelled, statistically associated, professionally judged, legally admitted, institutionally classified, or machine-generated? Each path has different limits. A direct observation may still be mis-scaled. An instrument may be calibrated for one regime and not another. Testimony may be sincere but incomplete. A statistical association may support prediction but not explanation. A model output may support triage but not sanction. A professional judgement may be strong but still residual-bearing.

The record path states how the relation was retained and redeployed. A record may be a file, dataset, case note, image, transcript, score, certificate, log, decision notice, audit report, model card, release note, court finding, or institutional archive. A record is valuable because it carries relation forward. But a record is not warrant by default. It may be incomplete, stale, misclassified, biased, created under a different closure demand, stripped of residual, or read beyond scope.

A record establishes that some content was registered and retained under conditions. It does not automatically establish that the relation was true, current, complete, relevant, or adequate for the present use. A historical risk flag may not describe current risk. A prior allegation may not establish fact. A compliance record may not establish safety. A model log may not establish target preservation. A court record may close procedure without closing explanation. The method requires the record’s provenance, scope, abstraction, residual, and update condition to remain visible.

The carrier path states what closure-work is being inherited. Many claims rely on trust carriers: persons, roles, credentials, institutions, archives, instruments, procedures, standards, certificates, releases, templates, schemas, models, interfaces, vendors, or machine outputs. These carriers reduce local recalculation. They make complex systems possible. A court cannot personally re-observe every fact. A patient cannot personally reconstruct all medical evidence. A regulator cannot inspect every process from first principles. A user cannot audit every model dependency. A public body cannot rebuild every dataset.

Carrier reliance is therefore not the error. The error is treating the carrier as ground.

A credential may support reliance on competence, but it is not competence by itself. A certificate may support reliance on completed process, but it is not substantive adequacy by itself. An institution may support reliance on procedure, but it is not relation-preservation by itself. A machine output may support reliance on a trained system, but it is not target possession by itself. A method release may support reliance on a stewarded formulation, but it is not epistemic authority by itself.

The method asks what the carrier is trusted to have done. Did it observe, measure, classify, verify, compare, certify, preserve, calculate, judge, or monitor? Under what domain? With what abstraction? Against what comparator? With what residual? Under what reliance limit? Under what update condition? If those questions cannot be answered, reliance-strength has detached from recoverable closure-work.

This matters because evidence, records, and carriers are the main routes through which claims travel. A claim may begin as observation, enter a record, become institutional memory, train a model, produce an output, inform a decision, generate a new record, and return as evidence in a later process. At each stage, abstraction may become invisible. Residual may be stripped away. Carrier trust may be inherited without carrier limits. Record presence may become warrant. The claim becomes stronger in practice while weaker in recoverability.[5]

Formalisation must interrupt that drift.

It must preserve the difference between evidence and record, between record and warrant, between carrier and ground, between abstraction and target, between model output and relation, between certificate and adequacy, between method-compliance and closure.

Where recovery fails, the audit should localise the failure as far as the evidence permits. The selective-preservation and receiving-side account in Part 2a, Paper 1c distinguishes a relation that never entered the relevant path from one that entered but lost discriminability through transformation, and both from a preserved relation that the receiving locus cannot presently access or stabilise. A relation already formed but insufficient for the closure demand presents a further, different limitation. These differences matter to correction: a distinction genuinely absent from the retained path requires another evidential route rather than more interpretation of the same record, while preserved but inaccessible relation may become available through a changed interface, translation, or intermediate scaffold. Formed but insufficiently supported relation instead requires the relevant comparator or closure work. Where the failure cannot itself be localised, that uncertainty must remain explicit; failed recovery alone does not show which condition failed. This is conditional diagnosis within the existing evidence, record, and carrier paths, not a new universal checklist or a requirement to reconstruct the whole causal history.

Some claim paths require a specialised recovery path within this general structure. Where power is claimed, the audit must recover the relevant Information Processing Locus, operative distinctions, physically supported branch possibility space, reference and comparator conditions, temporal conditions, and the distinction among potential, application, and consequence. Where responsibility is claimed, the audit must recover the responsibility-bearing transformations, relevant alternatives and operative frame, and distinguish Responsibility Provenance from any Responsibility Warrant allocated by role, law, procedure, or institution. Neither path is universal. Each is activated by the relation under audit.

The method does not require every claim to reconstruct all upstream dependencies at full resolution every time. That would destroy the finite economy the method is meant to discipline. It requires recoverability proportionate to the strength being claimed. A low-stakes reliance may carry light provenance. A high-consequence action requires stronger path reconstruction. Automation, coercion, durable record, affected-locus exposure, and irreversible consequence raise the burden.

The rule is simple:

Expose the abstraction.

Preserve the path.

Do not let record, carrier, or form become warrant by default.

---

6. Comparator discipline and residual

A claim is not strong because it is available.

It is not strong because it is named, recorded, repeated, measurable, institutionalised, formalised, encoded, adopted, certified, or acted upon. These features may make the claim usable. They may make it visible. They may make it portable. They may make it easier to rely on. They do not by themselves show that the claim has survived the alternatives relevant to the closure demand being made.

Strong claim-strength requires comparator discipline.

Comparator discipline asks what the claim has been tested against. What alternatives were live? What neighbouring formulations were possible? What baseline was used? What error modes were considered? What rival explanations remain? What weaker claims were positively excluded? What affected-locus challenge was heard? What scale shift was tested? What deployment condition was compared? What would distinguish this formulation from nearby but weaker formulations?[6]

A claim about risk must be distinguished from mere uncertainty, suspicion, institutional inconvenience, remote possibility, precautionary concern, or political anxiety. A claim about harm must be distinguished from discomfort, offence, disagreement, administrative burden, ordinary risk, or speculative future injury. A claim that something is evidence must distinguish allegation, record, material evidence, inference, correlation, testimony, and proof, and state what evidential strength the route carries. A claim about model adequacy must be distinguished from benchmark success, proxy performance, deployment reliability, explanation, and target preservation. A claim about intelligence must be grounded in applied distinction and distinguished from usefulness, imitation, complexity, prediction, or success alone. A claim about method-compliance must be distinguished from relation-preservation.

A power claim requires comparator pressure against a condition in which the proposed distinction-conditioned organisation cannot make the claimed branch difference. Consequence may demonstrate that potential was real, but consequence magnitude, successful outcome, institutional authority, or externally imagined possibility does not substitute for the branch relation itself.

A responsibility claim requires event-specific comparator and localisation discipline. The audit must ask which transformation materially conditioned the branch, which alternatives were available or recoverable within the locus's operative frame, and what temporal conditions apply. Where non-application is attributed, the mere existence of unused power is insufficient; the relevant alternative must have been available or recoverable and its non-application must have materially conditioned the realised branch. Causal proximity, outcome valence, role, authority, and allocated answerability do not establish Responsibility Provenance.

Where preferential reliance on another locus is claimed, comparison must establish the asserted advantage for the shared target relation and closure demand, including its preservation through the carrier path.

The purpose of comparator discipline is not to demand impossible opposition for every claim. It is to ensure that strength is earned by surviving relevant contrast. The relevant comparator

depends on the closure demand. A scientific model requires different comparators from a legal judgement. A legal judgement requires different comparators from an emergency action. A machine classifier requires different comparators from a medical diagnosis. A governance decision requires different comparators from a formal proof. But in every case, a claim becomes stronger by showing what it has survived.

Comparator discipline protects the method from first-formulation capture. The first available formulation often feels strong because it reduces uncertainty. It names the relation. It permits action. It fits the available category. It enters the record. But a formulation that has not survived comparator pressure may only be the first usable Soft Horizon. It may be necessary for temporary orientation. It is not yet strong merely because it is usable.

Comparator discipline also protects against inherited branch loss. A stabilised closure may have once survived alternatives, but later users may inherit only the winning label. They may not inherit the excluded branch-space. A theory, legal standard, policy category, machine benchmark, institutional procedure, or method release may travel without the cases that tested it. Later use then treats the surviving branch as inevitable. The method asks which alternatives were excluded, whether the exclusion was earned, whether the present use remains within the domain of survival, and what residual the old comparison left behind.

A model may have survived benchmark comparison but not deployment challenge. A policy may have survived internal consultation but not affected-locus challenge. A category may have survived administrative use but not edge cases. A legal standard may have survived procedure but not moral explanation. A machine output may have survived validation but not distribution shift. A certification may have survived process review but not substantive adequacy. The comparator must match the claim's current strength.

Comparator theatre is therefore a failure. It occurs when alternatives are listed but not allowed to test the claim. The institution names options already excluded by design. The model report includes benchmarks that miss the deployment relation. The policy review cites objections but does not let them alter abstraction. The legal decision notes alternatives but does not carry their residual. The audit process includes comparator fields but no meaningful comparator pressure. The appearance of comparison then becomes another Soft Horizon.

Real comparator discipline may produce several outputs.

It may strengthen the claim because relevant alternatives are excluded. It may weaken the claim because a neighbouring formulation survives. It may leave multiple live formulations. It may show that the claim is adequate for one closure demand but not another. It may disclose that the relation cannot yet be audited because the evidence path or abstraction is too weak. It may identify residual that must travel with the claim. It may justify a Best Available Formulation without closure.

Comparator discipline therefore leads directly to residual.

Residual is the localised remainder that prevents overclosure. It is what remains unresolved after relation, locus, availability condition, closure demand, abstraction, evidence path, record path, carrier path, and comparator pressure have been stated. It is not vague uncertainty. It is not a decorative caveat. It is not an apology for weak analysis. It is the method's protection against false possession.

Residual may concern missing evidence. The claim may lack a record, measurement, testimony, observation, calibration, comparison, or affected-locus report required for stronger closure.

Residual may concern abstraction. The category may be useful but lossy. The proxy may not preserve the target. The metric may compress too much. The legal standard may fit ordinary cases but fail a boundary case. The machine label may capture correlation while missing meaning. The institutional category may treat unlike cases as equivalent.

Residual may concern scale or horizon. A claim may be strong locally but weak globally, strong short-term but weak long-term, strong in aggregate but weak individually, strong in one deployment and weak in another.

Residual may concern evidence path. A record may exist but have weak provenance. A testimony may be credible but incomplete. A model may be evaluated but not deployed. A certificate may attest process but not substance. An archive may preserve decision but not dissent.

Residual may concern carrier path. A trusted role, institution, model, release, certificate, or machine output may reduce local audit but remain limited by its domain, incentive, abstraction, evidence, comparator, safeguard, or update condition.

Residual may concern affected-locus exposure. The relying locus may find the claim sufficient while the affected locus bears unresolved risk. A classification may be convenient for an institution but under-grounded for the person classified. A machine output may support internal triage but not external consequence. A legal threshold may permit temporary action but not final judgement.

Residual may concern update. The system may not yet know what would force revision, or the update condition may be known but inaccessible to the affected locus, user, auditor, or downstream system.

Residual must state what it blocks.

This is crucial. Residual does not always block all use. A claim may be insufficient for final judgement but sufficient for investigation. Insufficient for sanction but sufficient for monitoring. Insufficient for automation but sufficient for human review. Insufficient for explanation but sufficient for prediction. Insufficient for permanent policy but sufficient for temporary precaution. Insufficient for public assertion but sufficient for internal hypothesis.

Residual becomes disciplined only when it states the limits of permitted strength.

An evidential residual may block proof but not inquiry. A model residual may block deployment but not research. A medical residual may block diagnosis but not precaution. A legal residual may block final decision but not interim measure. An institutional residual may block automation but not escalation. A trust residual may block high-consequence reliance but not low-stakes provisional use.

Residual must also state the update condition. What would reduce, transform, or remove the residual? Further evidence? Stronger comparator? Affected-locus challenge? New measurement? Deployment monitoring? Record correction? Independent review? Changed scale? Failed prediction? New domain case? If update cannot be stated, the residual risks becoming inert. If update is stated but cannot be triggered, it becomes theatre.

The method’s residual discipline prevents verdict laundering.

A verdict is any output that closes a decision path enough to be used: approved, refused, compliant, non-compliant, safe, unsafe, high risk, low risk, observed, not observed, intelligent, not intelligent, eligible, ineligible, reasonable, unreasonable, certified, uncertified, actionable, closed. Such verdicts may be necessary. But the verdict must not absorb the residual under which it was produced.

A system that says “high risk” while hiding that the risk is based on weak proxy, uncertain horizon, disputed affected locus, and incomplete comparator has laundered residual into verdict. A system that says “method-compliant” while hiding that relation-preservation failed has laundered residual into procedure. A system that says “model validated” while hiding deployment residual has laundered residual into technical closure. A system that says “action justified” while hiding threshold and safeguard has laundered residual into authority.

Residual laundering blocks correction because the system no longer knows where it is incomplete. Later evidence enters the same overstrong frame. Affected-locus challenge is treated as resistance to a settled category. Model failure is treated as anomaly rather than abstraction failure. Institutional harm is treated as implementation error rather than claim-path failure. The residual remains operative, but it is no longer visible.

The method must therefore make residual one of its strongest outputs.

Sometimes the correct output is not a stronger claim. It is residual.

Sometimes the correct output is not closure. It is comparator failure.

Sometimes the correct output is not action. It is a statement that action threshold has not been met.

Sometimes the correct output is action under residual, with threshold, safeguard, and update.

Sometimes the correct output is non-auditability because the relation, abstraction, evidence path, record path, carrier path, comparator, or residual cannot be recovered.

This is not weakness. It is the discipline of finite strength.

Comparator discipline asks what the claim has survived.

Residual discipline states what the claim has not closed.

Together, they determine whether a claim may strengthen, remain provisional, become actionable, require safeguard, or weaken. They keep Soft Horizons useful without letting them become false possession.

---

7. Best Available Formulation without closure

The method does not end every audit with closure.

Often the strongest disciplined output is not a final claim, but a Best Available Formulation.

A Best Available Formulation is the strongest currently admissible formulation of the relation under review, selected under stated conditions, after available comparator pressure, while carrying residual openly. It is not final truth. It is not mere opinion. It is not equal possibility. It is not action-certainty. It is the best current way of carrying the relation without pretending that the relation has been closed beyond the conditions that support it.

This concept is necessary because finite systems must often proceed where closure remains incomplete. A scientific practice may use a model before ontology is settled. A medical system may treat under uncertainty. A court may decide under evidential limits. A public body may act before all future effects are known. A machine-assisted process may classify under deployment residual. A governance body may adopt a provisional policy where delay itself carries risk. In each case, the system may need a formulation strong enough to guide action, review, or communication, but not strong enough to erase residual.

The Best Available Formulation preserves this distinction.

It says: this is the strongest formulation available now; this is why weaker alternatives have been excluded or downgraded; this is the evidence path and comparator field supporting it; this is the residual that remains; this is the closure demand for which it is sufficient; this is the closure demand for which it is not sufficient; this is the update condition that would revise it.

This structure prevents two opposite failures.

The first failure is premature closure. A formulation survives enough comparison to become usable, and then the system treats usability as settlement. The model becomes the target. The category becomes the relation. The record becomes warrant. The policy phrase becomes proof. The action threshold becomes truth. The machine output becomes decision. The Best Available Formulation becomes final claim. This is overclosure.

The second failure is false equality. Final closure is unavailable, so the system treats all remaining formulations as equally live. Weak alternatives gain status from the fact that certainty is absent. Speculation, residual, anomaly, and serious counterformulation are flattened into one field. The system refuses to weight because it cannot close. This is not discipline. It is another way of losing relation.

The Best Available Formulation avoids both failures. It permits weighting without closure. It permits present use without final possession. It permits selection without forgetting what selection did not settle.

A Best Available Formulation should be produced only after the prior audit burdens have been satisfied at the required level. The relation must be stated. The locus must be identified. The availability condition must be clear. The closure demand must be specified. The abstraction must be exposed. Evidence, record, and carrier paths must be distinguished. Comparator discipline must have been applied. Residual must have been localised. Only then can the method say which formulation is currently strongest.

The formulation must also state what kind of strength it has. It may be strongest for recognition, but not explanation. Strongest for investigation, but not sanction. Strongest for triage, but not automation. Strongest for local decision, but not general policy. Strongest for prediction, but not ontology. Strongest for precaution, but not proof. Strongest for administrative routing, but not affected-locus classification. If that scope is not stated, the formulation will travel beyond the closure demand it can satisfy.

The Best Available Formulation also preserves branch memory. It should not merely state the selected formulation. It should preserve the alternatives that were considered, excluded, weakened, or left live. A later audit may need to know which branches were available. A future failure may show that an excluded formulation should be reconsidered. An affected-locus challenge may reveal that the selected formulation preserved institutional convenience better than target relation. A machine deployment may show that an apparently weak residual was operationally decisive.

Branch memory is not indecision. It is the record of finite selection under residual.

In high-consequence domains, the Best Available Formulation should carry an explicit reliance limit. Who may rely on it, for what purpose, and with what consequence? A formulation may be safe for internal analysis but not public assertion. It may be adequate for provisional monitoring but not enforcement. It may support emergency intervention but not durable classification. It may guide expert judgement but not automated action. Reliance limit prevents the formulation from becoming a general trust carrier merely because it is currently strongest.

The Best Available Formulation is therefore a disciplined Soft Horizon. It closes enough relation for present use while refusing to hide the conditions under which it remains incomplete.

It is best because it has survived the available comparison.

It is available because it is finite, situated, and conditioned.

It is a formulation because it remains answerable to relation rather than becoming possession of the target.

---

8. Action under residual

Finite systems often must act before closure.

This is not an exception to the method. It is one of the conditions that makes the method necessary. A body withdraws before identifying every cause of pain. A doctor treats before every causal pathway is known. A court orders interim relief before final judgement. A regulator intervenes before every future consequence is measured. A public body acts under precaution. A platform removes or limits content before all context is resolved. A machine system may route a case for review before the final status is known.

The method does not forbid action under residual. It forbids action under residual from narrating itself as closure.

Action threshold is not truth threshold.

A claim may be insufficient for explanation but sufficient for monitoring. Insufficient for final judgement but sufficient for investigation. Insufficient for sanction but sufficient for warning. Insufficient for automation but sufficient for human review. Insufficient for permanent classification but sufficient for temporary precaution. Insufficient for ontological claim but sufficient for practical response. The method asks which threshold is being used and why.

An action threshold should state the condition under which action becomes justified despite incomplete closure. That condition may be urgency, reversibility, precaution, legal duty, risk magnitude, evidential sufficiency, vulnerability, emergency, proportionality, public duty or mandate, or the cost of inaction. These may justify action. They do not remove residual. They state why action may proceed while residual remains.

The threshold must be indexed to consequence. Low-consequence, reversible action may proceed under heavier residual. High-consequence, coercive, durable, automated, or hard-to-reverse action requires stronger grounding, stronger safeguard, and clearer update. A model output may support review under one threshold but not denial. A risk flag may support monitoring but not sanction. A medical suspicion may support testing but not irreversible intervention. A legal allegation may support inquiry but not final finding. A governance concern may support consultation but not coercion.[7]

The threshold must also distinguish relying locus from affected locus. A claim may be convenient for the institution that relies on it while dangerous for the person or system exposed by it. A public body may rely on a category to act quickly. A citizen may bear the residual. A platform may rely on a classifier to scale enforcement. A user may bear the restriction. A regulator may rely on a certificate to reduce inspection burden. The public may bear the failure. Where the affected locus is not the relying locus, the action threshold must account for exposure, not merely institutional need.

Action under residual therefore requires safeguard.

A safeguard is not reassurance. It is the constraint that prevents action under incomplete closure from becoming unchecked authority. It may include review, appeal, notice, consent, human oversight, second opinion, independent audit, domain limitation, limited duration, staged deployment, reversibility, monitoring, record correction, affected-locus challenge, stop condition, or withdrawal path. Its function is to keep residual operational after action begins.

A safeguard is real only if it can constrain action. An appeal that cannot reach the relevant relation is not a safeguard. Human oversight that rubber-stamps machine output is not a safeguard. A review process that cannot change the decision is not a safeguard. A monitoring system that cannot trigger update is not a safeguard. A notice that gives no route of challenge is not a safeguard. Such devices may look responsible, but they do not discipline action under residual.

Action under residual also requires update.

Update is the condition under which the action, claim, threshold, category, model, record, or reliance relation must be revised. It may be triggered by new evidence, affected-locus challenge, failed prediction, adverse outcome, changed context, model drift, comparator failure, record correction, legal change, domain shift, or safeguard activation. Without update, action under residual becomes durable closure. The system acts before closure, then treats the action as if closure had been achieved.

This is especially dangerous because action changes the field. Once a system acts, it creates records, incentives, path dependencies, institutional commitments, affected-locus burdens, and retrospective narratives. The action itself may become evidence for later action. A temporary classification becomes archive. A provisional risk flag becomes future input. A machine-assisted decision becomes training data. A legal interim finding becomes practical stigma. A governance intervention becomes policy precedent. If residual is not carried forward, action creates the conditions for its own false confirmation.

The method must therefore preserve the difference between justified action and settled claim.

A decision may be justified and still residual-bearing. A claim may be actionable and still unclosed. A model output may be useful and still limited. An institutional category may be necessary and still provisional. A safeguard may permit action without making the relation safe. A threshold may permit response without proving the classification.

This discipline protects finite systems from paralysis without licensing overclaim.

A system that refuses action until final closure may fail ethically, legally, medically, scientifically, or practically. Delay can be harmful. Non-action can be a decision. Residual does not automatically block intervention. But a system that acts under residual must not use the necessity of action to erase the uncertainty that made action difficult.

Action and non-action may themselves enter responsibility-bearing paths, but neither Power nor Responsibility follows automatically from the fact that a formulation remained live, that

an alternative was physically included within a Counterfactual Inclusion Field, or that an outcome occurred. Power potential does not by itself establish responsibility for a realised branch or for a standing condition of branch availability. An allocated duty does not by itself establish Responsibility Provenance. Where responsibility is attributed, the audit must recover the distinctions and responsibility-relevant capacities available within the locus's operative frame under the relevant temporal conditions. Where application is claimed, it must recover the relevant Counterfactual Exclusion. Where non-application, retention, maintenance, or cross-locus field-conditioning is claimed without application, it must recover the comparator-stable difference made to the relevant Counterfactual Inclusion Field. The live formulation field may establish what formulations remained epistemically available to the observer-system; it does not by itself establish what causal alternatives were physically included for the responsible locus. Action threshold may justify proceeding under residual; it does not predetermine responsibility attribution.

Action under residual should therefore be stated in the following form: this relation remains unclosed in these respects; action is justified under this threshold; this residual blocks these stronger claims; this safeguard constrains exposure; this update condition will force revision.

Where this structure is absent, action becomes authority without memory.

Where this structure is present, finite systems can act without pretending to possess total closure.

---

9. Method-compliance is not closure

The method must not become its own idol.

A system may complete the method's form and still fail to preserve relation. It may state a relation weakly. It may identify the wrong locus. It may confuse availability with warrant. It may define a closure demand too narrowly. It may expose an abstraction while using a poor one. It may record evidence without testing it. It may name a carrier path while overtrusting the carrier. It may list comparators without real comparator pressure. It may state residual without allowing residual to block anything. It may define safeguards that do not constrain action. It may state update conditions that no one can trigger.

In such a case, the method has been complied with formally but not substantively.

Method-compliance is not closure.

This principle is necessary because every formal discipline can become a trust carrier. A completed claim packet may become a defence. A template may become proof of diligence. A certificate may become assurance. A schema may become the practical method. A compliance statement may become a shield. A court may treat method-use as reasonableness. A regulator may treat completion as adequacy. A vendor may treat implementation as reliability. A public body may invoke method alignment as legitimacy.

The method was introduced to prevent ungrounded authority. It must not become another source of it.

The method is itself capable of becoming a formalised authority carrier. Its release, stewardship, adoption, technical expression, certification, implementation, and domain application therefore remain answerable to the same governance discipline. Custodial and operational authority may stabilise and apply the method, but they do not create epistemic

authority. No current technical form becomes final merely through publication, adoption, implementation, or repeated use. Governance must preserve disciplined dissent, comparator alternatives, serious forks, residual, and branch-memory across Live Field Contraction so that a stabilised release does not acquire apparent inevitability merely because earlier formulations have disappeared from active use. That branch-memory records an epistemic and historical formulation field; it does not establish that archived alternatives belonged to the Counterfactual Inclusion Field of any Information Processing Locus. Where demonstrated relation-preservation exposes failure, the relevant operational, technical, or proposed constitutional level must remain open to correction.

Completing the method does not prove truth, safety, legality, justice, objectivity, intelligence, observation, model adequacy, institutional competence, or moral legitimacy. It makes the claim path more visible. It makes residual harder to hide. It gives auditors, affected loci, reviewers, courts, regulators, scientists, domain experts, and future users a structure for challenge. That is valuable. It is not proof.

The method's output should therefore be read as an audit state, not as a certificate of final adequacy.

A claim may be method-audited and strong. It may have preserved relation under stated conditions and survived comparator pressure with limited residual. A claim may be method-audited and weak. The method may expose poor abstraction, weak evidence, missing comparator, overtrusted carrier, or residual that blocks action. A claim may be method-audited and actionable under threshold. It may justify temporary action while preserving residual. A claim may be method-audited and non-auditable because the necessary path cannot be recovered.

All of these are legitimate outputs.

The method does not exist to approve claims. It exists to locate their strength.

This point prevents compliance closure. Compliance closure occurs when the fact that the method was used becomes a substitute for relation-preservation. The institution says the audit was performed. The vendor says the tool followed the method. The regulator says the template was completed. The court says the process was reasonable. The public body says the decision was method-aligned. Each claim may matter procedurally. None establishes that the relation was preserved.

The correct question remains: what did the method show?

Did it show strong claim-strength? Did it show residual? Did it show that the claim is sufficient only for a weaker closure demand? Did it show that action is justified only under safeguard? Did it show that the record path is inadequate? Did it show that carrier reliance exceeds recoverable conditions? Did it show that comparator pressure was missing? Did it show that the claim cannot be audited?

A method that always produces approval has been captured.

A method that cannot produce refusal, residual, non-auditability, weakened claim-strength, action limits, or update obligations is not a method of audit. It is a compliance ritual.[8]

Method-compliance also must not become a substitute for domain judgement. The method can expose whether claim-strength is earned, but it does not decide all legal, scientific, medical, ethical, political, or engineering questions by itself. A legal system must still decide legal consequence. A medical system must still decide intervention. A scientific community must still test theory. A public body must still deliberate value and mandate. A machine-governance

system must still allocate deployment duties, authority, and decision ownership within its governing regime. Where such a domain makes a responsibility claim, the method may audit the distinction between Responsibility Provenance and Responsibility Warrant; it does not itself supply the moral, legal, organisational, or institutional warrant. The method disciplines the claim path; it does not replace the domain.

This protects the method from both overreach and trivialisation. If the method claims to settle every domain question, it becomes false authority. If it is treated as a mere form, it becomes decorative. Its proper role is narrower and stronger: to prevent claims, classifications, records, models, decisions, trust carriers, and actions from carrying more strength than their recoverable conditions support.

Method-compliance is therefore evidence of attempted recoverability.

It is not evidence of closure unless the audit itself shows that closure was earned.

---

10. The governing structure of the method

The method has a dependency order.

It begins with a claim, but it does not trust the claim's form. The claim must state the relation it carries. The relation must be indexed to locus, availability condition, closure demand, boundary, interface, scale, and horizon. The abstraction that makes the relation usable must be exposed. Evidence, record, and carrier paths must be distinguished from warrant. Comparator discipline must test the claim against relevant alternatives. Residual must be stated and must say what it blocks. If closure is unavailable, a Best Available Formulation may be selected without becoming closure. If action follows under residual, threshold, safeguard, and update are required. Method-compliance itself must not become closure.

This order matters because later steps cannot repair missing earlier steps.

Comparator discipline cannot test a relation that has not been stated. Residual cannot be localised if closure demand is unclear. Safeguard cannot constrain action if the action threshold is unstated. Update cannot revise a category whose abstraction was never exposed. A carrier path cannot be evaluated if the relying locus and affected locus are confused. A Best Available Formulation cannot be selected if the live formulation field was never compared or if its contraction history has become unavailable where that history is material to the audit. Method-compliance cannot strengthen a claim whose relation cannot be recovered.

The governing sequence is therefore:

A claim must state its relation.

The relation must be indexed.

The abstraction must be exposed.

The evidence path, record path, and carrier path must remain paths, not warrant by default.

Where a power claim is made, the Counterfactual Inclusion Field physically supported by the distinction-conditioned organisation of the specified Information Processing Locus must be recoverable under the relevant reference and comparator conditions.

Where a responsibility claim is made, the attributable consequential branch relation must be localised and Responsibility Provenance must remain distinct from Responsibility Warrant;

application claims require recoverable Counterfactual Exclusion, while non-application or field-conditioning claims require a comparator-stable difference to the relevant Counterfactual Inclusion Field.

Comparator discipline must test the formulation.

Residual must remain visible.

Best Available Formulation may be selected without closure.

Action under residual requires threshold, safeguard, and update.

Method-compliance is not closure.

This sequence is not a mechanical checklist. It is a dependency structure. Domains may express it differently. A scientific paper, legal judgement, clinical record, public-administration decision, machine-governance process, safety certification, institutional review, or philosophical argument will not use identical forms. The method does not require one bureaucratic dialect. It requires that the claim-strength be recoverable.

This also prevents over-formalisation.

The method does not require every ordinary claim to carry full audit burden. Casual conversation, metaphor, early exploration, low-stakes judgement, and reversible local action may operate with light structure. The burden rises with consequence, authority, mediation, automation, irreversibility, affected-locus exposure, residual weight, and downstream reliance. The method is proportional because formalisation itself can become a burden and a misclosure if imposed indiscriminately.

But proportionality is not permission for under-formalisation. Where claims classify, govern, allocate, exclude, sanction, automate, certify, or justify consequential action, the claim path must carry enough structure for audit. The more authority a claim carries, the more recoverability it owes.

The method’s structure therefore guards against two failures.

Under-formalisation occurs where a claim carries authority without recoverable relation, abstraction, evidence path, carrier path, comparator, residual, threshold, safeguard, or update. It produces epistemic debt, institutional opacity, trust laundering, machine amplification, and false closure.

Over-formalisation occurs where the form becomes stronger than the relation it was meant to preserve. It produces box-ticking, compliance closure, schema reduction, false precision, and method idolatry. It treats completed structure as if it were truth. It prevents judgement from answering to relation because the form has replaced relation.

The method avoids under-formalisation by requiring recoverability.

It avoids over-formalisation by refusing to treat recoverability as possession.

The governing structure is therefore strict but finite. It does not promise final ontology. It does not abolish residual. It does not remove the need for trust. It does not remove the need for action under uncertainty. It does not replace domain judgement. It does not make machines epistemically adequate by execution. It does not make institutions legitimate by procedure. It does not make records warrant by existence.

It gives finite systems a way to carry claim-strength without false possession.

The method is successful when it makes a claim more answerable, not when it makes the claim more confident.

It is successful when it shows that a claim is strong, weak, residual, actionable, non-auditable, overextended, or updateable under stated conditions.

It is successful when it preserves the path by which a Soft Horizon can be used without becoming an idol.

---

11. Conclusion: claim strength without false possession

The Epistemic Method begins from a simple condition: finite observer-systems must claim, classify, model, decide, rely, and act before they possess total closure.

This condition cannot be escaped. It can only be disciplined.

The method is that discipline. It audits claim-strength and, where claims are relied upon, reliance-strength. It asks whether a claim preserves the relation it purports to carry under stated conditions, and whether any trusted carrier satisfies the relying locus's closure demand strongly enough for the reliance being placed upon it.

Its governing rule remains:

No claim earns strength beyond the recoverable conditions under which it preserves the relation it purports to carry. Where the claim is relied upon, no reliance closure earns strength beyond the recoverable conditions under which the trusted relation, carrier, or closure-process satisfies the relying locus's closure demand.

Everything else follows from that rule. A claim must state its relation. A relation must be indexed to locus, availability condition, closure demand, boundary, interface, scale, and horizon. Abstraction must be exposed. Evidence, record, and carrier path must not become warrant by default. Strong claims require comparator discipline. Residual must not be hidden inside verdict. Best Available Formulation must not become closure. Action under residual requires threshold, safeguard, and update. Method-compliance must not become truth.

Where a claim concerns power, its strength remains tied to the Counterfactual Inclusion Field physically supported by the distinction-conditioned organisation of the specified Information Processing Locus under the conditions being claimed. Where a claim concerns responsibility, attribution remains tied to the recoverable consequential branch relation attributable to the relevant locus: Counterfactual Exclusion where application is claimed, or a comparator-stable difference to the relevant Counterfactual Inclusion Field where responsibility arises through non-application, retention, maintenance, or material field-conditioning without application. Responsibility Warrant remains distinct from Responsibility Provenance. Where epistemic alternatives are reconstructed, the live formulation field and its contraction history remain distinct from physical causal inclusion. Where the method itself acquires authority, its stewardship, technical form, application, dissent, forks, branch-memory, and amendment path remain answerable to demonstrated relation-preservation. These are not additions to every audit. They are the conditional burdens activated by the relation being claimed.

The method is strict against false strength. It refuses to let epistemic authority attach to name, record, category, procedure, institution, model, certificate, release, schema, machine output, public confidence, or action already taken unless the relation-preserving conditions are recoverable.

The method is also permissive toward finite action. It does not demand total certainty before claim or decision. It permits provisional closure, best available formulation, trust carriers,

institutional judgement, machine assistance, and action under residual. It understands that a system unable to act until final closure would not become more rational. It would become unusable.

The discipline lies in carrying the incompleteness honestly.

A Soft Horizon is not a failure. It is the ordinary condition of finite use. The failure begins when the Soft Horizon hardens into false possession: when the category becomes the target, the record becomes warrant, the model becomes ontology, the threshold becomes proof, the carrier becomes ground, the decision becomes truth, or the method becomes closure.

The Epistemic Method does not abolish Soft Horizons.

It keeps them answerable.

It gives finite systems a way to claim without pretending to possess, to act without pretending to know finally, to rely without laundering trust, to formalise without worshipping form, and to revise without losing the relation under review.

Its final demand is modest and severe:

where authority travels, grounding must travel with it.

Endnotes

[1] Toulmin (1958) provides a major structural neighbour in practical argumentation: claims are assessed through recoverable grounds, warrants, backing, qualification, and rebuttal rather than by a single universal formal schema. The present claim path is not the Toulmin model; it additionally carries indexed relation, locus, availability, closure demand, abstraction, record and carrier paths, comparator, residual, reliance, threshold, safeguard, and update. No argumentation framework is inherited.

[2] The phrase relational uncertainty is established in interpersonal-communication research. Knobloch and Solomon (1999) use it for uncertainty concerning the self, partner, and relationship in dating relationships. The present paper uses the term differently: uncertainty about whether a claim, closure, reliance relation, or Soft Horizon preserves the relation it purports to carry under indexed conditions. This is terminological provenance and disambiguation, not conceptual inheritance.

[3] Conditional Power and Responsibility comparators: Albantakis and Tononi (2015) formalise intrinsic cause-effect power in discrete dynamical systems through counterfactual causal repertoires. This is a strong formal specialisation of the requirement that a Power claim recover a physically supported counterfactual possibility field, but the Epistemic Method does not inherit Integrated Information Theory, its intrinsic-causal commitments, or its formal repertoire. Chockler and Halpern (2004) use structural causal models to formalise degrees of responsibility and distinguish responsibility from blame. This is a formal specialisation of counterfactual responsibility attribution, but the present Responsibility audit is built from the NTC/IPL/Power architecture, distinguishes Responsibility Provenance from Responsibility Warrant, and does not inherit a universal scalar or the Chockler-Halpern model. These sources therefore locate the conditional audit branches without becoming their technical ground.

[4] Ulrich (1987) makes boundary judgements and affected parties explicit in Critical Systems Heuristics. This is a specialised methodological neighbour to the present locus, affected-locus, boundary, and abstraction discipline. The Epistemic Method is not restricted to social planning or normative systems design and does not inherit Critical Systems Heuristics.

[5] Moreau et al. (2015), reflecting the W3C PROV architecture, formalise provenance through recoverable relations among entities, activities, and agents. This is a technical specialisation of path recoverability. Provenance may support assessment of quality, reliability, or trustworthiness, but in the present method provenance or record recovery does not become warrant by itself.

[6] Mayo (2018) makes evidential severity depend on testing capable of revealing flaws in a claim. This is a strong statistical specialisation of the present requirement that comparators exert real discriminatory pressure. The Epistemic Method does not inherit error-statistical formalism and does not restrict comparator discipline to statistical inference.

[7] Rudner (1953) argues that standards for accepting scientific hypotheses are affected by the seriousness of error. This is a historical specialisation of consequence-sensitive acceptance and action thresholds. The present method separates that practical threshold from truth or closure status: action may be justified under residual without the unresolved relation being narrated as closed.

[8] Method and compliance comparators: Power (1999) analyses how audit and verification practices can become organisational objects tied to accountability and control. This is a direct institutional comparator for the warning that method-compliance can harden into a compliance ritual. Cartwright, Hardie, Montuschi, Soleiman, and Thresher (2022) argue that scientific reliability is not secured by generic appeals to method, rigour, or objectivity alone, but depends on heterogeneous networks of supporting scientific products and practices. That supplies adversarial pressure against treating an Epistemic Method as a universal reliability guarantee. The present discriminator remains relation-preservation under recoverable conditions: method-use can expose and discipline a claim path, but it cannot replace domain judgement or create closure merely by completion.

References

Albantakis, L., & Tononi, G. (2015). The intrinsic cause-effect power of discrete dynamical systems—From elementary cellular automata to adapting animats. Entropy, 17(8), 5472-5502. https://doi.org/10.3390/e17085472

Cartwright, N., Hardie, J., Montuschi, E., Soleiman, M., & Thresher, A. C. (2022). The tangle of science: Reliability beyond method, rigour, and objectivity. Oxford University Press. https://doi.org/10.1093/oso/9780198866343.001.0001

Chockler, H., & Halpern, J. Y. (2004). Responsibility and blame: A structural-model approach. Journal of Artificial Intelligence Research, 22, 93-115. https://doi.org/10.1613/jair.1391

Knobloch, L. K., & Solomon, D. H. (1999). Measuring the sources and content of relational uncertainty. Communication Studies, 50(4), 261-278. https://doi.org/10.1080/10510979909388499

Mayo, D. G. (2018). Statistical inference as severe testing: How to get beyond the statistics wars. Cambridge University Press. https://doi.org/10.1017/9781107286184

Moreau, L., Groth, P., Cheney, J., Lebo, T., & Miles, S. (2015). The rationale of PROV. Web Semantics, 35, 235-257. https://doi.org/10.1016/j.websem.2015.04.001

Power, M. (1999). The audit society: Rituals of verification. Oxford University Press.

Rudner, R. (1953). The scientist qua scientist makes value judgments. Philosophy of Science, 20(1), 1-6. https://doi.org/10.1086/287231

Toulmin, S. E. (1958). The uses of argument. Cambridge University Press.

Ulrich, W. (1987). Critical heuristics of social systems design. European Journal of Operational Research, 31(3), 276-283. https://doi.org/10.1016/0377-2217(87)90036-1