GOT-V1-3a-2

The Discipline of the Unknown

Residual, Live Formulation, and Action Under Incomplete Closure

---

Abstract

The unknown is not mere absence. For a finite observer-system, the unknown is the ordinary field of relation beyond possession, relation outside present availability, relation available before recognition, relation misclosed by pattern, relation preserved as residual, relation held among live formulations, relation weighted without final settlement, and relation requiring action before closure.

Paper 1 established the horizon grammar: no horizon can be analysed without specifying the relation, the locus, the availability condition, and the closure demand. The present paper asks what follows once that grammar is applied to the unknown. A finite observer-system does not merely fail to know. It encounters different forms of unresolved relation, and those forms require different disciplines.

The unknown must therefore be separated from weaker or misleading substitutes. It is not subjective doubt, because unresolved relation may persist where confidence is high. It is not arbitrary possibility, because not every imagined formulation remains live. It is not mystical surplus, because residual must be localised and made updateable. It is not failure, because a properly stated residual may be the strongest available output under finite conditions. It is not action paralysis, because finite systems often must act before closure. It is not closure by another name, because action under residual does not convert uncertainty into warrant.

This paper distinguishes unavailable unknowns, available-but-unrecognised unknowns, misclosed unknowns, formulated residuals, live formulation fields, weighted residuals, and actionable residuals. It then develops the disciplines by which finite observer-systems carry them: access humility, closure discipline, failure-mode diagnosis, residual localisation, positive exclusion, live formulation, weighting without false precision, Best Available Formulation, weighted residual, action threshold, safeguard, monitoring, stop condition, and update trigger.

The governing rule is simple: the unknown must not be converted into closure, explanation, equal possibility, action-certainty, or unmarked provisional closure.

A disciplined observer-system becomes stronger not by pretending to escape the unknown, but by making the unknown visible, structured, weighted, revisable, and accountable to update inside and around the Soft Horizons through which finite systems must continue.

---

1. Introduction: unknownness, residual, and action under incomplete closure

This paper concerns the unknown as the structured condition that appears around relational availability, incomplete closure, misclosure, residual, and action under finite conditions.

Paper 1 established the foundation. A finite observer-system does not encounter reality as total state. It encounters relation through boundary, interface, scale, horizon, discrimination regime, abstraction, registration, retention, redeployment, comparator pressure, closure demand, and residual limit. It operates through Soft Horizons: provisional closures that make relation usable for recognition, modelling, classification, communication, judgement, or action while stronger closure remains unavailable. A horizon cannot be analysed without specifying the relation, the locus, the availability condition, and the closure demand.

The discipline of the unknown begins from that grammar but asks a different question. Paper 1 asked how relation becomes available, unclosed, provisionally closed, residual, or horizon-bound for a locus. This paper asks how finite observer-systems should preserve, order, weight, act under, and update the unknowns produced when relation is unavailable, available but unrecognised, misclosed, residual, provisionally closed, or action-demanding.

What is unknown is therefore not merely what has not yet been named, measured, proven, explained, or recorded. It may be relation outside availability. It may be relation available without recognition. It may be relation closed under a failing pattern. It may be relation provisionally closed for one use while remaining residual for another. It may be relation held among live formulations. It may be relation weighted without final settlement. It may be relation under which action must proceed before closure.

The unknown is not peripheral to inquiry. It is the ordinary field within which finite observer-systems operate. Most of reality, relative to any finite locus, is not presently available. Of what becomes available, only some becomes registered. Of what is registered, only some is retained. Of what is retained, only some is redeployed. Of what is redeployed, only some is warranted. Of what is warranted, only some is closed for the purpose being claimed. Even what is provisionally closed remains residual relative to stronger demands.

A Soft Horizon does not remove the unknown. It localises and operationalises a bounded part of it. This is why finite systems can proceed without possessing total closure. They recognise enough to move. They model enough to predict. They classify enough to decide. They communicate enough to coordinate. They act enough to survive. But none of these operations abolishes the unknown that remains outside, beneath, around, or inside the closure being used.

The term observer-system is used broadly. It does not mean a human subject by default. It may refer to an organism, cell, instrument, institution, computational system, formal mathematical community, scientific practice, ecological arrangement, legal procedure, social organisation, or coupled observer-network. The unknown is not defined anthropocentrically. It is indexed to finite loci capable of boundary-mediated relation, registration, retention, redeployment, correction, or action. Human knowing is one important instance. It is not the root form.

The discipline of the unknown has one governing rule:

The unknown must not be converted into closure, explanation, equal possibility, action-certainty, or unmarked provisional closure.

This rule protects five boundaries.

First, it protects closure from being claimed where the discriminator conditions for closure have not been met. A relation may be available, registered, named, or acted upon without being closed under the demand being made.

Second, it protects residual from becoming a mystical or rhetorical substitute for explanation. Residual is not a word for ineffable depth. It is a localised remainder: what remains unclosed, what the remainder prevents us from claiming, and what would force update.

Third, it protects live alternatives from being flattened into equal status merely because final closure is unavailable. Non-closure does not make all formulations equally disciplined. Some alternatives are positively excluded. Some are speculative. Some are weak but live. Some are strong but not closed.

Fourth, it protects action under uncertainty from presenting itself as if it had removed the uncertainty under which it acts. A finite system may need to act before closure. But action threshold is not truth threshold. Action does not become certainty because it was necessary.

Fifth, it protects provisional closure from hardening into inherited certainty after its residual, scope, tolerance, or update condition has disappeared. A Soft Horizon may be legitimate for one purpose and overstrong for another. The unknown is often lost when provisional closure travels without the conditions that made it provisional.

The task of the paper is to state how the unknown can be carried without being erased, worshipped, flattened, or used as false warrant for action. It develops a discipline for preserving unknown-statuses, identifying failure modes, localising residual, excluding failed formulations, structuring live alternatives, weighting without false precision, selecting a Best Available Formulation, carrying weighted residual, acting under threshold, preserving safeguards, monitoring for failure, and updating when the horizon moves.

The unknown is not a defect at the edge of method.

It is the condition that makes method necessary.

---

2. The unknown is not absence

The simplest mistake is to treat the unknown as a blank.

Something is either known or unknown. What is known has content; what is unknown lacks content. This binary is useful in ordinary speech, but it is too flat for disciplined inquiry. It collapses many different relations into a single negative term.

A relation may be unknown because it has not entered the observer-system's access conditions at all. It may be unknown because it has entered access conditions but has not yet been identified. It may be unknown because it has been identified under a pattern that later fails correction. It may be unknown because it has been provisionally closed for one use while remaining residual for another. It may be unknown because it has been formulated but not decided. It may be unknown because more than one formulation remains live. It may be unknown because action must proceed while closure remains unavailable.

These are not the same condition. Treating them as the same condition creates conceptual and practical failure.

They are unknown-statuses, not replacements for relational availability. Relational availability asks what status a relation has for a locus under stated conditions. The discipline of the unknown asks what must be preserved when that availability status does not satisfy the closure demand being made. The same target relation may therefore be unavailable for one locus, available but unrecognised for another, closed for action under a Soft Horizon, and residual for explanation.

The first form is the unavailable unknown. This is relation outside the present access regime of the observer-system. It may lie beyond the system's boundary, beyond its interface, beyond its

scale, beyond its temporal horizon, below its sensitivity, outside its discrimination regime, or beyond the abstraction by which it can register difference. The unavailable unknown is not necessarily unreal. It is not even necessarily weak. It is simply not available as relation for this observer-system under these conditions.

The proper discipline for unavailable unknowns is access humility. The observer-system should not convert non-availability into non-existence. It should not treat absence from its present regime as evidence that the target relation is absent absolutely. But it also should not convert non-availability into open speculation without constraint. The correct output may be non-availability, weak or indirect availability, constrained speculation, or access gap, depending on the available path.

The second form is the available-but-unrecognised unknown. Here something has crossed into the observer-system's access field, but it has not been stabilised as identity. A signal-bearing contrast may reach the locus without being classed. A perturbation may be sensed without being explained. A record may show anomaly without theory. A bodily state may become salient before diagnosis. A legal institution may receive testimony or register an anomalous case before it can classify the harm under its governing categories. An organism may respond to disturbance before any centralised representation exists. The relation is no longer outside availability, but it has not been closed.

The discipline for available-but-unrecognised unknowns is suspension with pattern search. The observer-system should preserve the relation as available but unclosed. It should resist the nearest available pattern where that pattern has not earned closure. It should state what has entered access, what recognition has not yet occurred, and what would be required for classification, explanation, or action.

The third form is the misclosed unknown. In this case the observer-system does not merely fail to identify. It identifies incorrectly. An available relation is forced into a pattern that does not preserve the target relation under correction. Optical illusion gives the primitive perceptual case. A line appears longer, a surface appears moving, a figure appears stable, but measurement or repeated correction shows that the perceptual closure does not preserve the relevant relation. Misclosure is not absence of registration. It is registration followed by an overstrong claim about what the registered relation establishes.

The discipline for misclosed unknowns is correction and failure-mode diagnosis. The system must ask which pattern was applied, what relation it failed to preserve, what comparator exposed the failure, what residual had been hidden, and whether the failure belongs to scale, abstraction, category, evidence path, record path, instrument, procedure, model, or action threshold.

The fourth form is formulated residual. Here the observer-system has enough discipline to state what remains unresolved. The unknown is no longer vague darkness. It has been localised. A claim, model, measurement, category, attribution, or explanation has reached a boundary where some discriminator condition remains unmet, or a Soft Horizon has been adopted for one purpose while leaving stronger demands unclosed.

The correct output is not silence and not closure. It is residual: a retained statement of the missing condition, the current limit, what the provisional closure is sufficient for, what it is not sufficient for, and the update route by which the residual could be reduced, transformed, or removed.

The fifth form is the live formulation field. Sometimes the unknown does not present as one missing condition around a single formulation. Several candidate formulations remain admissible. They have not all been proven. They have not all been excluded. They are not equal

merely because they remain open. The field must be structured. Which formulations remain live? Which have been positively excluded? Which are weak but not excluded? Which are strong but not closed? The unknown becomes a disciplined field of admissible alternatives.

The discipline for live formulation fields is admissibility control. Not every imagined possibility belongs in the live field. A formulation is live only where it preserves relation well enough to remain admissible under the current reference conditions and has not failed the relevant discriminator or comparator. Live formulation protects inquiry from premature closure, but it also protects inquiry from unlimited possibility.

The sixth form is weighted residual. Once live formulations are distinguished, the observer-system may be able to weight them. Weighting is not closure. It is proportional discipline under non-closure. A formulation may be better supported by available evidence, stronger comparator survival, better scale fit, superior explanatory compression, greater predictive resilience, or lower residual burden, while still falling short of closure. Weighted residual prevents the observer-system from treating unresolved possibilities as equal merely because final settlement is unavailable.

The discipline for weighted residual is proportional support without false precision. The system should state why one formulation is stronger than another, what kind of support grounds that weighting, what uncertainty remains, and what would alter the weighting. Weighting is a way to carry the unknown forward without either flattening it or prematurely closing it.

The seventh form is actionable residual. Finite systems often must act before closure. Biological systems must respond to threat before full identification. Institutions must make decisions before complete evidence. Scientific practice may proceed with a model while its ontology remains unsettled. Medical judgement may require intervention before all causal pathways are known. Engineering systems must operate under uncertainty. Actionable residual is the condition in which action is justified, required, or unavoidable while epistemic closure remains incomplete.

The discipline for actionable residual is threshold, safeguard, monitoring, stop condition, and update. The system must state why action is permitted or required, what consequence or urgency sets the threshold, what safeguards limit the action, what monitoring will detect failure, what condition would stop or reverse the action, and what would force revision. The danger is to convert action into certainty. The discipline is to act by threshold while preserving residual.

These forms show why the unknown cannot be treated as simple absence. Absence of closure does not mean absence of relation. Absence of knowledge does not mean absence of structure. Absence of final warrant does not mean equal possibility. Absence of certainty does not mean action paralysis.

The unknown is a family of relations between target, observer-system, abstraction regime, and closure demand.

This distinction also prevents a common inflation of knowledge. If the unknown is treated as blank, then any available formulation may appear to defeat it. A name seems better than no name. A model seems better than no model. A measurement seems better than no measurement. A category seems better than ambiguity. A decision seems better than suspension.

But an overstrong closure can be worse than an explicit residual. A misclassification can be worse than disciplined non-closure. A false explanation can suppress the conditions that would have allowed correction. A premature action can destroy the possibility of later update.

The unknown becomes disciplined only when the observer-system preserves the difference between non-availability, availability, recognition, provisional closure, misclosure, residual,

live formulation, weighting, action threshold, and update. Without those distinctions, the unknown is either erased or worshipped. It is erased when closure is claimed too soon or when a Soft Horizon hardens into final possession. It is worshipped when residual is treated as ineffable depth beyond analysis.[1]

Both failures avoid discipline.

The discipline of the unknown begins with a more difficult claim: the unknown has forms, and those forms matter.

---

3. Indexed unknownness: relation, locus, and constraint

The unknown is indexed. This is unavoidable.

A relation is unavailable, available, registered, retained, deployed, residual, or closed only relative to some locus under some regime. What lies outside one observer-system's boundary may lie inside another's. What is below one scale may be operative at another. What is noise for one discrimination regime may be signal for another. What is unformulable in one abstraction language may become expressible in another. What is unclosed under one method may close under a stronger comparator, instrument, proof system, model, or category.

Indexed unknownness does not make the unknown subjective. It does not mean that the observer invents the target relation. It does not mean that unknownness is merely a mental state. It does not mean that whatever the observer cannot close is therefore indeterminate in itself. It means that unknownness must be stated as a relation-locus-availability-closure claim: what relation is at stake, for which locus, under which availability condition, and under which closure demand.

A target relation can constrain an observer-system without being fully available to it. Reality does not need to be possessed as total state in order to constrain finite access. A chemical gradient can shape a cell's behaviour without becoming a human concept. A gravitational relation can constrain a measuring apparatus without becoming directly visible. A mathematical relation can constrain formal work before it is recognised as theorem, undecidable statement, or contradiction. A pathogen can alter an organism before diagnosis. An institutional category can fail a lived relation before reform names the failure.

The target constrains even when closure is absent.

The observer-system also constricts the target. It does not receive the target whole. It receives relation through boundary, interface, scale, threshold, noise regime, abstraction, registration pathway, retention capacity, and redeployment structure. This constriction is not optional. A finite system must select, filter, stabilise, and compress in order to encounter anything at all. A Soft Horizon is the disciplined name for the usable closure produced by such constriction where residual remains. An observer-system without boundary or discrimination would not possess the real more fully. It would fail to register relation.

The discipline therefore rejects two symmetrical errors.

The first is total-state realism. This error treats a record, perception, measurement, model, category, or other Soft Horizon as if it were possession of the target whole. The observer-system mistakes its access product for unbounded reality. It forgets boundary, scale, abstraction, comparator, retention, and residual. It treats the available and provisionally closed relation as if availability had already become total closure.

The second error is empty constructivism. This error treats observer-relativity as if it dissolved target constraint. Because every claim is indexed to a locus and regime, it concludes that the target is merely constructed by that regime. This fails in the opposite direction. Observer-systems do not create the consequences that correct them. They do not choose which measurements survive calibration, which classifications fail, which models break under prediction, which organisms survive, which proof steps follow, or which residuals become damaging.

A Soft Horizon is therefore neither total possession nor arbitrary construction. It is finite cut under real constraint.

The unknown lives between those errors. It is not the observer's private ignorance, but neither is it an object waiting in completed form for direct possession. It is a relational condition. It arises where the target constrains but the observer-system has not closed the relation under the closure demand being made, or where the observer-system has closed enough for one purpose while residual remains for another.

This is why the word observer must be handled carefully. If observer is read as human mind, the discipline becomes anthropocentric. The unknown becomes human uncertainty, human ignorance, human doubt, or human knowledge deficit. That is too narrow. A bacterium has an access regime. A plant has an access regime. A measuring instrument has an access regime. A model has an operational envelope. A court has procedural categories. A language has classification affordances. A scientific community has instruments, standards, records, and correction pathways. A computational system has tokenisation, architecture, training distribution, and output constraints.

This broader use does not attribute human consciousness to all such systems. It does not say that every finite locus has beliefs, intentions, concepts, or experiences. It says that finite loci encounter only through relation. Where such relation is bounded, classed, registered, retained, deployed, corrected, or acted upon, the unknown can be indexed to that locus without importing human subjectivity.

The point is not to humanise all observers.

It is to dehumanise epistemic structure enough to make human knowing one case within a wider architecture.

The unknown is therefore neither subjective nor absolute. It is not subjective because the target constrains the relation. It is not absolute because unknownness must be stated relative to a locus and regime. A relation unknown to one observer-system may be available to another. A relation unavailable at one scale may be tractable at another. A relation unclosed under one abstraction may close under another. A relation closed for action may remain unclosed for explanation. A relation closed locally may remain residual globally.

The discipline of the unknown requires this exact indexing. Without it, the observer-system either overclaims or dissolves. It overclaims when it treats local closure as total possession. It dissolves when it treats local mediation as arbitrary construction.

The disciplined position is harder but more stable: the relation is indexed; the constraint is real; closure is earned only under stated conditions.

4. Vast unknown and horizon unknown

The most common form of the unknown is not anomaly. It is not paradox. It is not undecidability. It is not a difficult problem already visible to inquiry.

The most common form of the unknown is non-availability relative to a finite observer-system.

This is the vast unknown. In the vast unknown, the governing condition is that relation has not entered the present access regime strongly enough for the relevant closure demand to be attempted. The proper output is therefore not proof-failure, explanation-failure, or residual by default. It is non-availability, weak or indirect availability, constrained speculation, or an explicitly stated access gap.

Most of the universe does not enter the access regime of any given locus. It does not cross the relevant boundary. It does not survive the relevant interface. It does not appear at the operative scale. It does not persist over the required horizon. It does not exceed the threshold. It does not become discriminable against noise. It does not fall into the available abstraction. It is not registered, retained, or redeployed by the locus under the present regime. It may be real, structured, consequential elsewhere, and still remain outside the observer-system's event horizon.

The vast unknown should not be treated as a set of hidden facts already formatted for the observer. That image smuggles in human-scale epistemology. It imagines the unknown as a library with unread pages, a map with unexplored regions, or a set of propositions awaiting discovery. Sometimes that metaphor is useful. Often it is misleading. Much of what lies outside an observer-system may not be formatted as object, fact, proposition, category, measurement, or explanation for that observer at all. It becomes such only if a pathway, abstraction, instrument, interface, or coupled locus makes the relation available in a form the system can register.

The discipline appropriate to the vast unknown is access humility.

Access humility is not sceptical collapse. It does not say nothing can be known. It does not flatten all claims. It does not deny target constraint. It states the current access condition without pretending that non-availability has become either non-existence or mystery. It permits the observer-system to say: this relation is outside current access; only weak or indirect availability has been established; speculation is constrained by the evidence and limits presently available; this claim cannot yet be attempted; this path would be required for stronger availability.

Access humility matters because unavailable relation is easily abused. It can be erased. A system may act as if what does not enter its current categories does not exist. Institutions often do this when harms are not recordable. Scientific practices may do it when an effect lies outside instrument range. Computational systems may do it when target relations are absent from training data. Legal systems may do it when lived relation cannot enter admissible evidence. The relation may remain consequential even where the system cannot receive it.

Unavailable relation can also be inflated. A system may treat non-availability as license for unconstrained speculation. Because the relation cannot be closed, anything is imagined as possible. The unknown becomes a rhetorical space for projection. This is not discipline. The fact that a relation is unavailable does not make every formulation admissible. Constrained

speculation must remain tied to available indication or evidence, known limits, reference conditions, or physically, formally, institutionally, or procedurally plausible paths.

The vast unknown therefore demands a narrow but important form of strength. The observer-system must be able to say what it cannot presently access without converting that lack into either denial or fantasy. It must preserve the access gap.

The horizon unknown is different.

In the horizon unknown, relation has entered some condition of availability but remains unclosed under the demand being made. The relation is not wholly outside access. Something has crossed the boundary. It may be sensed, registered, formulated, named, measured, recorded, modelled, procedurally received, or acted upon. But availability has not become closure.

The horizon unknown is the field of available-but-unrecognised relation, misclosed relation, anomaly, paradox, formulated residual, provisional closure, live formulation, and action under residual. It is more structured than the vast unknown because something has entered the observer-system's regime. It is also more dangerous, because availability invites closure.

A sound heard in darkness may be available without identity. A symptom may be registered without diagnosis. A scientific anomaly may be measured without explanation. A mathematical statement may be formulable without proof. A legal institution may receive testimony without a category able to preserve the harm. A machine system may output a classification without warranting the target relation. An institution may record a case while losing the relation that made the case consequential. These are not blank absences. They are relations that have entered the system enough to create a closure demand.

The first danger in the horizon unknown is premature recognition. The observer-system forces relation into the nearest available pattern because available relation creates pressure for closure. The unknown sound becomes threat. The unusual symptom becomes a familiar disease. The anomaly becomes error. The testimony becomes an existing legal category. The model output becomes fact. Premature recognition reduces anxiety but may lose the target relation.

The second danger is misclosure. Here the system not only recognises too quickly; it carries the closure forward as if it were warranted. The relation becomes record, category, model, decision, proof-claim, or action premise. The residual disappears. The system may then become harder to correct because the misclosure has acquired downstream force.

The third danger is residual laundering. The system acknowledges uncertainty but hides its force. It says "there is residual" while treating the claim as if the residual does no work. It records uncertainty in a footnote, caveat, risk register, model note, legal qualification, or institutional file while permitting the main closure to travel as settled. The unknown remains verbally acknowledged and operationally erased.

The fourth danger is action-certainty conversion. The system acts because action is necessary, urgent, proportionate, precautionary, or legally required. The action may be justified. But later the system narrates the action as if it had closed the uncertainty under which it acted. The threshold becomes proof. The decision becomes warrant. The fact of continuation becomes evidence that the relation was settled.

The horizon unknown therefore requires closure discipline.

Closure discipline does not forbid closure. It asks what kind of closure has been earned. A relation may be closed for immediate action and open for explanation. It may be closed for classification and residual for ontology. It may be closed for legal procedure and residual for social repair. It may be closed for local prediction and residual for mechanism. It may be closed under a Soft Horizon and open under a stronger demand.

The key distinction between the vast unknown and the horizon unknown is therefore not certainty. Neither is certain. The distinction is availability. In the vast unknown, the relation has not entered the present access regime strongly enough for the closure demand to be attempted. In the horizon unknown, the relation has entered availability but closure remains incomplete, failed, provisional, contested, residual, or action-bound.

This distinction changes the proper output.

For the vast unknown, the disciplined output may be non-availability, weak or indirect availability, constrained speculation, or access gap.

For the horizon unknown, the disciplined output may be suspension, misclosure diagnosis, formulated residual, live formulation set, weighted residual, Best Available Formulation, action threshold, safeguard, monitoring condition, stop condition, or update trigger.

The two fields interact. A relation may move from the vast unknown into the horizon unknown when an instrument, category, abstraction, procedure, model, record path, or coupled locus makes it available. A relation may also retreat into non-availability if a record is lost, an instrument fails, a category is withdrawn, a witness disappears, a model no longer applies, or a path cannot be reconstructed. Horizons move because availability moves.

The unknown is therefore dynamic. It is not a fixed region surrounding knowledge. It is the changing structure of relation under finite access, closure demand, abstraction, residual, and update.

A disciplined observer-system must know which kind of unknown it is handling.

Where relation is outside availability, it must preserve access humility.

Where relation is available but unclosed, it must preserve closure discipline.

Both are necessary. Without access humility, the system denies or invents what it cannot access. Without closure discipline, the system overclaims what it can access. Together, they prevent the unknown from becoming blank absence, mystical depth, speculative excess, premature explanation, equal possibility, or action-certainty.

The unknown has entered discipline only when the observer-system can say not merely “we do not know,” but what kind of unknown is present, what relation is at stake, for which locus, under what access condition, under what closure demand, and what would move the horizon.

---

5. Failure modes of the unknown

The master failure-mode is availability/closure collapse.

It occurs when a relation is treated as closed because it is available, registered, named, operationally useful, institutionally stabilised, emotionally compelling, technically expressed, or acted upon. Something has entered the observer-system's field, and the system mistakes that entry for settlement. The available relation is converted into identity, explanation, warrant, category, action-certainty, or finality before the relevant closure demand has been satisfied.

Every failure mode in the discipline of the unknown is a variant of this collapse. The observer-system either denies available relation, overcloses relation that remains underdetermined, hides residual, flattens live alternatives, treats uncertainty as mystical depth, or converts action under threshold into retrospective certainty.

The unknown becomes dangerous when its structure disappears. A finite observer-system does not merely confront blank absence. It filters, compresses, recognises, misrecognises, records, remembers, models, acts, and updates under limits. It also depends on Soft Horizons: provisional closures that make continuation possible. Failure does not arise because such closures are used. Failure arises when their provisional status, residual, scope, tolerance, or update condition is lost.

The first failure is **premature closure**. A relation becomes available, but instead of carrying it as unrecognised, anomalous, residual, live, weighted, or provisionally closed, the observer-system converts it into a settled identity. A useful first cut becomes final identification. A necessary Soft Horizon becomes reified. The unknown sound becomes intruder. The anomaly becomes error. The pain becomes diagnosis. The model becomes explanation. The category becomes ontology. The action threshold becomes proof.

Premature closure is not always irrational in origin. Biological systems often must respond before classification is complete. Institutions often must act before evidence is exhaustive. Scientific practice often needs working models before ontology is settled. A first cut may be necessary, and a Soft Horizon may be the correct finite output. The failure is not provisional closure. The failure is allowing provisional closure to travel as final identity after the support, horizon, residual, and update condition have disappeared.

The second failure is **mystification**. Here the observer-system avoids premature closure, but only by turning the unknown into a special kind of explanatory object. Residual becomes depth. Obscurity becomes insight. Unresolved relation is protected from localisation. A missing discriminator becomes a mystery that explains itself.

Mystification is the inverse of premature closure. Premature closure gives too much determinate form to the unknown. Mystification gives the unknown immunity from form. Both failures prevent discipline. The unknown must neither be forced into the nearest available category nor elevated beyond audit.

Mystification often appears when a real limit is encountered. The observer-system reaches the edge of a pattern regime. Something resists classification, measurement, explanation, prediction, or proof. Instead of stating the missing condition, the system treats resistance itself as explanation. The residual becomes a noun. The noun becomes authority. This can happen in metaphysical, institutional, scientific, cultural, or personal settings. A phenomenon may be called ineffable when the actual issue is an underdeveloped abstraction regime. A social condition may be called complex when the system has not localised its evidential gaps. A scientific problem may be described as mysterious when the missing comparator has not been identified. A legal or institutional harm may be declared intangible when the available categories are failing to preserve the affected relation.

The third failure is **sceptical collapse**. The observer-system recognises that closure is incomplete and concludes that no disciplined claim can be made. Because total closure is

unavailable, every formulation is weakened equally. Because all observation is finite, no observation is treated as relation-tracking. Because every model carries residual, no model is allowed to be stronger than another. Because every decision occurs under uncertainty, decision becomes arbitrary or merely political.

Sceptical collapse is not the opposite of premature closure. It shares the same refusal of status discipline. Premature closure treats availability as closure. Sceptical collapse treats non-total closure as if no local closure, weighting, warrant, or action threshold can be earned. Both fail to preserve the actual condition of finite inquiry. Some claims are stronger than others. Some formulations survive comparator pressure better than others. Some residuals are local and tractable. Some Soft Horizons are legitimate for their purpose. Some action thresholds are justified. The fact that no finite observer possesses total state does not erase these differences.

The fourth failure is **false equivalence**. This is the more local form of sceptical collapse. Where closure is unavailable, all remaining possibilities are treated as equally live. The observer-system refuses premature closure but loses discrimination among alternatives. A well-supported formulation, a weak but admissible formulation, an excluded formulation, an unconstrained speculation, and a rhetorical possibility are all allowed to occupy the same field.

False equivalence is especially tempting in contested cases. If no final proof has closed the matter, one can say that alternatives remain. But alternatives do not remain live merely because they can be imagined. A formulation must preserve relation under stated conditions and must not have failed the relevant comparator. Some possibilities are excluded. Some are speculative. Some are live but weak. Some are live and strong. Some are operationally sufficient under threshold while remaining epistemically residual. Non-closure does not abolish these differences.

The fifth failure is **rhetorical capture**. The unknown is used as a tool of persuasion rather than preserved as a structured status. Uncertainty may be amplified to block justified action. It may be minimised to force closure. It may be invoked selectively against one formulation and ignored in another. It may be used to protect institutional authority, evade responsibility, preserve a preferred category, discredit a challenger, or justify emergency action.

Rhetorical capture is dangerous because it can speak the language of discipline while defeating discipline. It may demand impossible certainty from an opposing claim while treating its own Soft Horizon as settled. It may claim humility while refusing update. It may invoke residual while hiding the specific missing condition. It may say “more evidence is needed” without identifying what evidence would change the claim. It may say “we cannot know” where the system does know enough to act under threshold. It may say “the matter is settled” where only a procedural closure has occurred.

The correction is to force the rhetoric back into unknown-status. What relation remains unresolved? For which locus? Under what closure demand? Which alternatives remain live? Which have been positively excluded? What residual remains? What would update the claim? What action threshold, if any, is being invoked? If those questions cannot be answered, the invocation of uncertainty is not yet disciplined.

The sixth failure is **fixation**. Here the observer-system preserves a formulation beyond the point where its residual, comparator failure, or update condition should have reopened it. Fixation may attach to a first explanation, a dominant model, a legal category, an institutional record, a machine classification, a professional convention, or a personal expectation. The system continues to use the formulation because it has become familiar, convenient, authoritative, or embedded.

Fixation differs from premature closure in timing. Premature closure happens when a relation is closed too early. Fixation happens when a provisional closure is not reopened after its conditions have changed. The Soft Horizon may once have been legitimate. It may have been the Best Available Formulation at an earlier point. But later evidence, comparator pressure, affected-locus challenge, scale shift, deployment failure, or changed consequence may make its residual newly active. If the formulation is not reopened, the unknown is no longer merely residual. It has been buried under inherited closure.

The seventh failure is **residual laundering**. This occurs when unresolved remainder is verbally acknowledged but operationally erased. The system says that uncertainty remains, but allows the claim, category, model, record, decision, or action to travel as if it were settled. Residual appears in a caveat, appendix, footnote, policy note, risk register, model card, legal qualification, institutional file, or review statement. It is present enough to protect the system from accusation of overconfidence, but not present enough to limit claim-strength, alter threshold, trigger safeguard, structure monitoring, or define update.

Residual laundering is one of the most important institutional failure modes. It allows a system to appear disciplined while acting as if no residual matters. It can say that limitations were noted. It can say that uncertainty was considered. It can say that risk was recorded. But the question is not whether uncertainty was mentioned. The question is whether the residual constrained the claim or action in proportion to its force.

The eighth failure is **action-certainty conversion**. A system acts under uncertainty because action is necessary, urgent, precautionary, proportionate, legally required, morally demanded, or practically unavoidable. The action may be justified. But once action occurs, the system later narrates it as if the underlying uncertainty had been closed. The intervention becomes proof that intervention was warranted. The decision becomes evidence that the matter was settled. The continuation of a policy becomes confirmation that the policy was correct. The fact that a threshold was crossed is retold as if a truth condition had been satisfied.

This failure matters because finite systems must act under residual. A medical intervention may be justified before full diagnosis. A legal precaution may be justified before final determination. An engineering shutdown may be justified before cause analysis. A biological retreat response may be justified before identification. In each case, the action can be warranted by consequence and threshold without claiming that the underlying relation is closed.

The discipline is to keep action threshold distinct from epistemic closure. Action may be justified where closure is incomplete. But the residual under which the action occurred must remain visible after action. The system must preserve what it did not know, why it acted anyway, what safeguards constrained the action, what monitoring followed, what would stop or reverse the action, and what would update the underlying formulation.

These failure modes show why the unknown requires more than tolerance. A vague appreciation of uncertainty is not enough. The observer-system must know when it is carrying unrecognised availability, when it is using a Soft Horizon, when closure has failed, when residual is localised, when alternatives remain live, when weighting is possible, when action is justified, and when update is required.

Without that discipline, the unknown is closed too soon, left too vague, flattened into equal possibility, captured rhetorically, fixed beyond its support, laundered into confidence, or converted into authority by the very closures that made finite action possible.

6. Residual: localised remainder under closure demand

Residual is the proper form of unclosed relation once that relation has become relevant to a claim, model, observation, classification, explanation, action, or Soft Horizon.

Residual is not everything unknown. It is not the vast unknown. It is not mere non-availability. It is not diffuse uncertainty. It is not a mood of caution. It is not a decorative caveat. It is the relevant unclosed remainder that bears on the claim being made. To function as disciplined output, that remainder must remain visible in the claim path because closure has not been earned at the strength being asserted.

A relation below the threshold for an active closure demand should not be marked as residual by default. The disciplined output may instead be non-availability, weak or indirect availability, constrained speculation, or access gap. Residual begins when relation has entered availability strongly enough for a closure demand to be attempted and some condition still prevents the claimed strength from being earned.

Residual therefore occupies a disciplined middle position. The observer-system does not abandon the claim entirely, but neither does it close the claim falsely. It carries the remainder as part of the claim's honest form. A model may remain usable while residual blocks ontology. A legal category may remain procedurally necessary while residual blocks full relation-preservation. A scientific explanation may remain strongest available while residual blocks final mechanism. A machine output may remain operationally useful while residual blocks target warrant. A medical judgement may remain actionable while residual blocks diagnostic closure.

Residual differs from ordinary uncertainty because it is localised. Ordinary uncertainty may be diffuse. The observer-system does not know what is missing. It merely lacks confidence. Residual is more disciplined. It states where the claim fails to close. It identifies the missing discriminator, weak comparator, scale gap, abstraction gap, retention gap, deployment gap, attribution gap, or action gap. Residual is uncertainty made local.

The first task of residual is **localisation**. The observer-system must identify where the relation remains unclosed. A claim may fail at the boundary because the candidate frame is wrong. It may fail at the interface because no pathway has been shown. It may fail at the scale because the operative scale differs from the reference scale. It may fail at the abstraction because the cut excludes the relation it claims to preserve. It may fail at registration because no event of selection into state is established. It may fail at retention because the state does not persist as re-identifiable. It may fail at deployment because possible use is mistaken for realised redeployment. It may fail at comparator because no adequate contrast has been constructed. It may fail at action because the threshold is not matched to consequence. A residual that does not locate itself remains too vague.

The second task of residual is **limitation**. A residual must state what cannot yet be claimed. This is what prevents residual from becoming a harmless caveat. If the residual does no limiting work, it has not been carried. A model residual may prevent ontological claim while permitting prediction. A category residual may prevent final classification while permitting provisional handling. An evidential residual may prevent legal finding while permitting investigation. A machine residual may prevent automated decision while permitting human-assisted review. An action residual may prevent retrospective certainty while permitting precautionary intervention.

Limitation is what gives residual force. The system must say not merely “uncertainty remains,” but “because this condition remains unresolved, this stronger claim is not warranted.” Without limitation, residual becomes rhetorical insurance.

The third task of residual is **blocking**. Residual blocks specific migrations of claim-strength. It blocks availability from becoming closure. It blocks registration from becoming warrant. It blocks action from becoming proof. It blocks category from becoming ontology. It blocks model from becoming target. It blocks institutional record from becoming truth by default. It blocks confidence from becoming relation-preservation. It blocks Soft Horizon from becoming final possession.

This blocking function is often uncomfortable because it interrupts usable closure. A system wants to proceed. It has a name, category, record, measurement, model, decision, or action path. Residual says: proceed only under the strength actually earned. It does not always forbid continuation. But it forbids overstatement.

The fourth task of residual is **update orientation**. A residual must state what would change the claim. The missing condition should not be left as inert uncertainty. It should point toward the route by which the horizon might move. What evidence would reduce the residual? What comparator would test the formulation? What scale shift would matter? What affected-locus challenge would reopen the category? What measurement would separate live alternatives? What monitoring would detect failure? What event would force withdrawal?

Residual without update orientation is incomplete. It may preserve uncertainty, but it does not make the claim answerable. A disciplined residual carries forward the conditions under which it can be reduced, transformed, intensified, or converted into positive exclusion.

Residual therefore has a distinctive grammar. It should state the relation under review, the closure demand being attempted, the current formulation, the missing condition, the claim blocked by that missing condition, the provisional use still permitted, and the update condition.[2]

This does not require a rigid template inside the prose of every paper, but it does require conceptual recoverability. A reader should be able to identify what remains unresolved, why it matters, and what would alter the claim. If that cannot be recovered, residual has become too vague.

Residual is also not defeat. A properly stated residual may be the strongest available output under finite conditions. It may be more disciplined than a premature conclusion. It may preserve relation better than a false explanation. It may support better action than an overconfident category. It may allow inquiry to continue without dissolving into equal possibility. It may carry uncertainty in a form that future evidence can actually answer.

This is why residual is not a lesser form of knowledge in the ordinary sense. It is a condition of honest finite claim-making. It states the boundary of the claim. It preserves the target relation from overclosure. It keeps the observer-system exposed to correction.

Residual also prevents the false binary between closure and silence. Without residual, a system may think it must either decide or say nothing. But finite systems often need a middle form. They need to say: this is the strongest formulation available, this is what it supports, this is what it does not support, this is what remains unresolved, and this is what would force revision. Residual is that middle form.

Residual must not be multiplied without discipline. Not every unresolved possibility belongs in residual. A residual should bear on the closure demand. It should not become a storage place

for every imaginable doubt. Overbroad residual weakens inquiry because it makes every claim appear equally burdened. The residual should be local to the relation and demand.

A residual should also not be softened into uncertainty where positive exclusion is warranted. If a candidate has failed under the relevant comparator, it should not remain in residual merely because someone might imagine it otherwise. Positive exclusion and residual are different. Residual carries unresolved relation. Positive exclusion records failed relation under stated conditions. Confusing them allows failed formulations to remain artificially alive, and allows genuine residual to be misread as indecision.

Residual is therefore the hinge between the unknown and method. It is how unclosed relation remains available to disciplined inquiry. It prevents false closure, but it also prevents vague suspension. It carries the unknown in a form that can limit, guide, weight, act, and update.

A finite observer-system becomes disciplined not by eliminating residual, but by learning where residual belongs and what it forbids.

---

7. Live fields: positive exclusion, admissibility, and formulation sets

Residual is not the same as positive exclusion.

A residual remains where a relation has not closed under the demand being made. Positive exclusion occurs where a formulation, candidate, category, explanation, hypothesis, model, attribution, or action path has failed under stated reference conditions. Residual preserves what remains unresolved. Positive exclusion preserves what has been ruled out.

This distinction is necessary because finite observer-systems often confuse unresolved relation with failed relation. They may exclude too soon, treating a weakly supported alternative as if it has failed. They may also preserve too much, treating an excluded alternative as if it remains live merely because final closure of the whole field has not been reached.

Positive exclusion is a disciplined negative result. It does not mean “I dislike this formulation,” “this formulation is inconvenient,” “this formulation lacks current popularity,” or “this formulation has not yet been proven.” It means that the formulation fails a relevant discriminator, comparator, evidence path, scale condition, abstraction requirement, or closure demand under stated reference conditions.

A scientific hypothesis may be positively excluded by evidence that contradicts its prediction under controlled conditions. A legal claim may be positively excluded by failure to meet a required burden. A classification may be positively excluded because it does not preserve the target relation. A machine proxy may be positively excluded because it fails under deployment comparator. A medical diagnosis may be positively excluded by differential evidence. A formal path may be positively excluded by proof failure under the rules of the system.

Positive exclusion strengthens the live field by reducing false openness. It prevents the observer-system from treating failed alternatives as if they still carry equal status. It also preserves the reason for exclusion. This matters because exclusion without a retained reason becomes dogma, while refusal to exclude becomes false equivalence.

A disciplined exclusion should therefore retain its conditions. It should be possible to say: this formulation is excluded under these reference conditions because it failed this discriminator or

comparator. If those conditions change, the exclusion may need review. The exclusion is not arbitrary. It is indexed.

Residual and exclusion can coexist. A field may remain unresolved while some candidates within it have been excluded. A diagnosis may remain open while certain causes have been ruled out. A scientific problem may remain unsolved while a class of explanations has failed. A legal question may remain contested while one argument is unavailable. A machine classification problem may remain residual while one proxy metric has been excluded. The existence of residual does not keep every candidate alive.

This prepares the concept of the live formulation field.

A live formulation field is the structured set of formulations that remain admissible under the current relation, locus, availability condition, closure demand, comparator pressure, and residual. It is not a list of all possibilities. It is not a landscape of imagination. It is not a rhetorical balance of every position someone might assert. It is the field of formulations that have not been positively excluded and that preserve relation well enough to remain under disciplined consideration.

A formulation may be live in different strengths. Some formulations are weakly live. They are not excluded, but their support is thin. They may preserve some relation but lack sufficient comparator survival. They may require conditions not yet established. They may depend on a path not yet available. They remain admissible, but fragile.

Some formulations are strongly live. They preserve more relation, survive more comparator pressure, fit scale and abstraction better, reduce residual, or support better prediction, explanation, classification, or action. They are still not closed, but they are better carried by the current field.

Some formulations are speculative. They may be imaginable, internally coherent, or rhetorically attractive, but they lack sufficient relation to the available field. Speculation is not always useless. It may guide future search. It may expose hidden assumptions. It may help formulate tests. But speculation should not be treated as live unless it enters the disciplined conditions of the field.

Some formulations are excluded. They have failed under stated conditions. They should not be revived without showing that the exclusion conditions have changed, that the prior comparator was inadequate, or that the formulation has been materially revised.

This vocabulary prevents two opposite errors. The first is premature narrowing, where the observer-system selects one formulation too soon and suppresses remaining live alternatives. The second is unlimited expansion, where the system refuses selection by keeping all possibilities equally open. A live field is neither closure nor chaos. It is structured admissibility.

Live fields are especially important in domains where action must occur before closure. Medicine, law, engineering, governance, science, and machine-mediated classification frequently require decisions among live formulations. A doctor may not know the final diagnosis but may know which possibilities remain live, which are excluded, which are dangerous, and which must be treated first under threshold. A court may not possess total truth but may know which claims meet procedural burden, which fail, and which remain residual. An engineer may not know the root cause of a failure but may know which failure modes remain live and which require immediate safeguard. A machine system may not warrant a final target relation but may expose which classifications remain uncertain and where human review is required.

Live formulation also preserves future correction. If only the chosen formulation is retained, later update becomes weaker. The system forgets what alternatives were considered, why they were excluded, why one was weighted more strongly, and what residual remained. This produces retrospective overclosure. The chosen formulation appears inevitable because the live formulation field and the history of its contraction have been erased.

A disciplined live field retains branch-memory without becoming paralysed by it. It records the relevant alternatives, their status, their support, their exclusion conditions, and their residual. It does not preserve every fantasy. It preserves what mattered to the claim.

The live formulation field also protects affected loci. Institutions often close fields too soon from the standpoint of those affected by their categories. A system may treat a person, event, risk, harm, or claim as settled because its internal category has closed. But affected loci may expose a residual the institution has not preserved. A live field must therefore be open to challenges that show the abstraction has failed to carry the relevant relation. This does not mean every challenge becomes equally live. It means the field must remain answerable to relation, not merely to institutional convenience.

A live field should therefore be governed by admissibility, exclusion, weighting, and update.

Admissibility asks whether the formulation preserves enough relation to remain in the field.

Exclusion asks whether the formulation has failed under stated conditions.

Weighting asks how strongly each admissible formulation is currently supported.

Update asks what would alter the field by strengthening, weakening, excluding, reviving, or replacing a formulation.

This structure matters because the unknown is not simply one thing hidden behind the known. It often appears as a field of possible closures, some stronger than others, some failed, some unresolved, some actionable, some speculative. The discipline of the unknown must be able to carry this field without flattening it.

A finite observer-system that cannot preserve live fields will either close too soon or remain open too vaguely. Both failures damage inquiry. Premature closure loses relation. Vague openness loses discrimination.

The live field is the disciplined middle.[3]

---

8. Weighting without closure

Weighting is the discipline of proportional support under non-closure.

It becomes necessary when more than one formulation remains live, or when a single formulation is usable but residual prevents final closure. The observer-system must continue to distinguish stronger from weaker claims without pretending that the strongest claim has become settled. Weighting allows finite inquiry to proceed between false certainty and false equality.

Weighting is not closure. A weighted formulation has not become true merely because it is strongest. It has not become final because alternatives are weaker. It has not become proof because action favours it. Weighting states relative support under current conditions. It says: given the relation, locus, availability condition, closure demand, abstraction, comparator field,

evidence path, residual, and update condition now available, this formulation currently carries more or less strength than another.

Weighting is also not confidence alone. Confidence may be psychological, institutional, social, rhetorical, statistical, or practical. It may track support, but it may also mislead. A system may feel confident because a formulation is familiar, authoritative, repeated, institutionally preferred, machine-generated, or easier to act upon. Weighting must return confidence to support conditions. Why is the formulation stronger? What relation does it preserve? Which comparator has it survived? Which residual remains? What would weaken it?

The simplest weighting distinction is qualitative. A formulation may be weak, moderate, strong, or dominant under current conditions. These terms should not be used decoratively. Weak means admissible but thinly supported. Moderate means supported but carrying significant unresolved conditions. Strong means well supported under present comparators while residual remains. Dominant means substantially stronger than alternatives while still not closed under the demand being made.

Qualitative weighting is often sufficient. Not every field requires numbers. In many domains, numerical weights create false precision. A legal category, institutional assessment, scientific interpretation, medical judgement, or ethical risk may be harmed by arbitrary quantification. A number can make a residual-bearing judgement appear more settled than it is. It can turn proportional support into pseudo-measurement. Weighting should be as precise as the support conditions permit, not more.

Ordinal weighting may be appropriate where formulations can be ranked without assigning stable numerical probability. One formulation may be better supported than another because it explains more of the available relation, survives stronger comparator pressure, fits scale better, carries less residual, or requires fewer unsupported assumptions. Ordinal weighting helps avoid false equality while avoiding spurious quantification.

Numerical weighting may be appropriate where probability, frequency, error, confidence interval, risk estimate, likelihood, model score, or statistical support is legitimately constructed under known conditions. But numerical weighting should not be mistaken for closure. It remains dependent on data quality, model assumptions, reference class, sampling, measurement, abstraction, comparator, and deployment relation. A number can carry residual. It can even hide residual more effectively than prose because it appears exact.

Weighting also remains bounded by what the supporting path actually preserved, as the selective-preservation account in Part 2a, Paper 1c makes explicit. A later probability, ranking, or confidence judgement cannot restore a discriminating relation that the retained path lost through transformation or never carried. Independent evidence, justified context, or a defensible probabilistic model may still weight the alternatives, but the additional support must be attributed to those conditions rather than presented as a distinction recovered from the original record.

Weighting therefore requires proportionality between the form of weight and the strength of the support path. A vague field should not be forced into numerical precision. A measurable field should not be left in vague language merely to avoid responsibility. A strongly supported formulation should not be flattened into “one possibility among many.” A weak speculation should not be inflated into a live alternative by rhetorical emphasis.

The main enemies of weighting are false precision and false equality.

False precision occurs when the observer-system gives a formulation a degree of support that the underlying relation cannot sustain. This can happen through numbers, categories, rankings,

scores, labels, traffic-light systems, machine outputs, risk bands, or institutional grades. A residual-bearing judgement becomes more exact in presentation than in grounding. The output looks usable, but the precision belongs to the interface rather than to the relation.

False equality occurs when the observer-system refuses to distinguish support because closure is incomplete. All formulations are described as uncertain, contested, possible, or unresolved. This may appear humble, but it can be as misleading as false precision. If one formulation has survived strong comparator pressure and another has not, treating them equally hides knowledge already earned. If one has failed, keeping it live hides positive exclusion. If one is actionable under threshold and another is speculative, treating both as equal possibilities distorts action.

Disciplined weighting asks what kind of support is doing the work.

Evidence support concerns what has been observed, recorded, measured, or otherwise made available.

Comparator support concerns what alternatives the formulation has survived.

Abstraction support concerns whether the category or model cut preserves the target relation.

Scale support concerns whether the formulation holds at the relevant scale and horizon.

Explanatory support concerns whether the formulation reduces residual in the relation being explained.

Predictive support concerns whether the formulation survives future or out-of-sample branch pressure.

Practical support concerns whether the formulation justifies action under threshold, safeguard, and consequence.

These supports may diverge. A model may have predictive support without explanatory closure. A legal formulation may have procedural support without full causal explanation. A medical judgement may have practical action support before diagnostic closure. A machine output may have benchmark support without deployment support. A scientific hypothesis may have explanatory elegance without sufficient evidential support. Weighting must show which support is being claimed.

This prevents a strong formulation in one register from overclaiming in another. A formulation may be strongest for action and weak for explanation. It may be strong for local prediction and weak for ontology. It may be strong under institutional procedure and residual for affected-locus relation. It may be strong under model score and weak under target preservation. A disciplined weight is indexed to the closure demand.

Weighting also prepares Best Available Formulation. A Best Available Formulation is not merely the highest-confidence claim. It is the strongest currently admissible formulation under residual, with its residual retained. Weighting is what allows a system to identify such a formulation without pretending it has closed the field.

The relation between weighting and residual is therefore direct. Residual identifies what prevents closure. Weighting assesses how much support remains despite that residual, and how the residual compares across live formulations. One formulation may be stronger because it leaves less residual. Another may be stronger because its residual is better localised. Another may be weaker because its residual blocks the central closure demand. Another may be actionable despite heavy residual because the consequence threshold is different from the epistemic closure demand.

Weighting must remain updateable. A weight is not a final rank. It is a current status under stated conditions. New evidence may strengthen a formulation. Comparator failure may weaken it. A scale shift may expose hidden residual. A new abstraction may revive a previously weak formulation. A positive exclusion may remove an alternative. An action outcome may feed back into the field. A monitoring result may trigger revision. Weighting is disciplined only when it remains answerable to such movement.

A finite observer-system needs weighting because it cannot wait for final closure before every continuation. But it also cannot treat every possible continuation as equal. It must select, prioritise, investigate, model, decide, and act under incomplete closure. Weighting is the structure that lets it do so without misdescribing its own strength.

Weighting does not defeat the unknown.

It prevents the unknown from becoming undifferentiated.

---

9. Best Available Formulation and Weighted Residual

A finite observer-system often reaches a point where the field is no longer flat, but closure is still unavailable.

The relation has entered availability. Residual has been localised. Some candidates have been positively excluded. Some formulations remain live. Weighting has distinguished stronger from weaker support. Yet the closure demand has not been fully satisfied. The observer-system must now decide what, if anything, can be carried forward.

The answer is the Best Available Formulation.

A Best Available Formulation is the strongest currently admissible formulation of a relation under stated conditions, selected from a live formulation field, weighted by available support, and carried with its residual intact. It is not final truth. It is not certainty. It is not consensus. It is not convenience. It is not the most confident statement. It is not the formulation that best supports action already preferred. It is the strongest current Soft Horizon of formulation under incomplete closure.[4]

The Best Available Formulation is necessary because finite systems cannot always remain in undifferentiated suspension. Suspension may be correct where the field is too weak, too malformed, or too under-available to carry any formulation responsibly. But where support has concentrated, where alternatives have been structured, and where residual has been localised, the observer-system may need to carry a formulation forward without pretending that final closure has been reached.

A scientific practice may carry a model as best available while mechanism remains residual. A medical system may carry a working diagnosis while causal closure remains incomplete. A legal institution may carry an interim category while final adjudication remains unavailable. A machine system may carry a classification for triage while interpretation remains residual. An organism may carry an action-form under partial detection. An institution may carry a procedural decision under evidence limits. In each case, the formulation is not strengthened by hiding its incompleteness. It is strengthened by preserving it.

The Best Available Formulation must therefore be typed to the closure demand it serves.

A formulation may be best available for action but not for explanation. It may be best available for local prediction but not for ontology. It may be best available for classification but not for mechanism. It may be best available for provisional coordination but not for final judgement. It may be best available for monitoring but not for sanction. It may be best available for temporary precaution but not for permanent policy. The same formulation cannot silently migrate across these demands.

This typing prevents one of the most common overclosures. A system selects a formulation because it is adequate for a limited purpose. The formulation then travels downstream as if it were adequate for a stronger purpose. A treatment decision becomes diagnostic certainty. A predictive model becomes explanation. A legal interim category becomes moral description. A classification used for routing becomes identity. A precautionary judgement becomes settled risk. The formulation's original bestness was local; downstream use makes it overstrong.

A Best Available Formulation should therefore carry its own limits. It should make recoverable the relation under review, the locus for which the formulation is being made, the availability condition, the closure demand, the live field from which it was selected, the candidates excluded, the support that weights it, the residual that prevents closure, the use it permits, the use it blocks, and the condition under which it would change.

This does not require a rigid template in every ordinary case. It requires that the structure remains recoverable where the formulation becomes consequential. The stronger the claim, the higher the consequence, the more affected the locus, the less reversible the action, or the more authority the formulation carries, the more explicit the Best Available Formulation must become.

The Best Available Formulation is not weakened by acknowledging uncertainty. It is weakened when uncertainty is hidden.

A finite observer-system that states its best current formulation, its basis, its residual, its weakening conditions, and its replacement conditions has produced a stronger epistemic object than one that declares closure prematurely. Strength under finite conditions is not the absence of residual. It is the disciplined relation between formulation, support, residual, and update.

This makes the Best Available Formulation a hinge concept. It turns unknown discipline toward continuation without surrendering to action-certainty conversion. It allows the observer-system to proceed while keeping the unclosed relation alive. It answers the practical question that appears across finite domains: what can be responsibly carried forward when the relation is not closed but the field is not equal?

The answer is not simply the formulation the system wants to use. It is not the formulation supported by institutional habit. It is not the formulation most familiar to an expert community. It is not the formulation produced with the highest machine confidence. It is not the formulation that reduces anxiety or preserves procedure. It is the formulation that currently best preserves the relation under the strongest available discipline while visibly carrying the conditions under which it would change.

Weighted residual is the companion of Best Available Formulation.

Residual by itself states what remains unclosed. Weighted residual states how the remaining unresolved field is structured after live formulations have been identified, exclusions have been made, support has been weighted, and a Best Available Formulation has been selected. It is the form in which the unknown remains attached to disciplined continuation.

A weighted residual records what must not be forgotten while the current formulation is carried forward.

It records which alternatives remain live, which were weakly live, which were excluded, which were speculative, which residuals block stronger claims, which residuals do not block the present use, what support differentiates the field, and what would alter the weighting. It prevents the chosen formulation from erasing the field that made it provisional.

Weighted residual is especially important because selection changes memory. Once a formulation is selected, it becomes easier to treat it as natural. The field narrows around it. Records begin to accumulate under it. Institutional procedures adapt to it. Model outputs reinforce it. Language forms around it. Later observers may inherit the selected formulation without inheriting the uncertainty under which it was selected.

Weighted residual resists that loss.

It says: this is the formulation we are carrying, but this is the unresolved structure that still matters. This alternative remains weakly live. This candidate was excluded under these conditions. This residual blocks explanation but not prediction. This residual blocks deployment but not controlled trial. This risk was known but treated as non-blocking at the time. This support would change if a comparator failed. This update would reopen the field.

A weighted residual may be qualitative. It may state that one formulation is strongly supported, another weakly live, another speculative but not excluded, and another excluded under current conditions. It may be ordinal, ranking formulations without assigning numerical probability. It may be probabilistic where the evidence structure genuinely supports probability. It may be scenario-based where consequence dominates. It may be comparator-relative where the key issue is which formulation survives which contrast. The format is not universal. The burden is proportional representation of the unresolved field.

Weighted residual also prevents decorative uncertainty.

It is common to append a vague limitation statement to a strong claim. The system says uncertainty remains, but the main formulation continues to travel as if closed. This is not weighted residual. A disciplined weighted residual changes the force of the formulation. It states which uncertainty matters, what it blocks, what it permits, how alternatives compare, and what would alter the field.

The blocking relation is central.

Some residuals block all downstream use. Others block only stronger use. A residual may block mechanism but not correlation. It may block ontology but not prediction. It may block individual classification but not aggregate monitoring. It may block automation but not human review. It may block coercive action but not voluntary support. It may block final judgement but not investigation. Without this distinction, residual becomes either paralysing or irrelevant.

If every residual blocks every use, finite systems cannot continue. If no residual blocks use, residual becomes ornamental. Weighted residual gives residual its correct force by asking: what does this unresolved condition prevent us from claiming, doing, generalising, automating, enforcing, or forgetting?

The Best Available Formulation and weighted residual therefore form one unit.

The Best Available Formulation says what may currently be carried forward.

The weighted residual says what must not be forgotten while carrying it.

Without the Best Available Formulation, weighted residual risks becoming unstructured uncertainty. Without weighted residual, the Best Available Formulation risks becoming overclosure. Together they allow finite systems to continue under incomplete closure without falsifying the unknown.

---

10. Action under unknown conditions

Finite observer-systems often act before closure.

This is not an exceptional defect. It is one of the ordinary conditions of finite existence. Organisms act under partial detection. Instruments trigger thresholds before interpretation is complete. Institutions make interim decisions before all evidence is available. Scientific practices proceed with models while ontology remains unsettled. Medical systems intervene before causal pathways are fully known. Computational systems classify under uncertainty. Human beings decide under incomplete information. Refusal to act may itself become action.

The discipline of the unknown cannot therefore be a discipline of endless suspension.

Suspension is sometimes necessary. If the field is malformed, if residual is blocking, if consequence is too high, if support is too weak, or if no action threshold has been reached, suspension may be the only disciplined output. But suspension is not always available and not always responsible. Where relation becomes action-relevant before closure, the observer-system must decide whether action threshold has been reached.

An action threshold is not a truth threshold.[5]

It is the point at which the cost, risk, urgency, reversibility, consequence, and available support justify action despite remaining residual. A formulation may be insufficient for final explanation but sufficient for precaution. It may be insufficient for diagnostic closure but sufficient for treatment. It may be insufficient for conviction but sufficient for investigation. It may be insufficient for deployment but sufficient for controlled trial. It may be insufficient for general theory but sufficient for local engineering use.

The threshold depends on the action being taken.

This point must be held tightly. Action closure is local to the action named. It does not close recognition, explanation, attribution, ontology, warrant, or future judgement unless those closure demands are separately satisfied. A system may act responsibly under residual and still lack epistemic closure. The danger is to narrate action as if action had solved the relation that made action difficult.

Action thresholds must be typed by consequence.

Low-consequence actions may proceed under lighter support. High-consequence actions require stronger support, stronger residual visibility, stronger safeguard, stronger monitoring, and stronger reversal conditions. Irreversible actions require more discipline than reversible actions. Coercive actions require more discipline than voluntary actions. Safety-critical actions require more discipline than exploratory actions. Actions that expose others require more discipline than actions whose consequences remain local to the actor.

The unknown is therefore not ethically neutral. Residual becomes more demanding as consequence rises.

Urgency complicates the threshold. A high-urgency situation may justify action under heavier residual than would ordinarily be acceptable. A biological system may flee before it identifies the predator. A doctor may treat before every test returns. An engineer may shut down a system before cause analysis. An institution may issue temporary protection before final adjudication. Urgency can lower the action threshold for some purposes, but it does not remove residual. It makes residual more important because the system is acting before closure.

Reversibility also matters. A reversible action can often proceed under lower support because correction remains possible. An irreversible action demands stronger justification because update cannot fully undo the branch. Some actions are partially reversible but leave persistent records, costs, harms, institutional effects, or path dependencies. A temporary record may become permanent in practice. A provisional classification may shape later decisions. A model trial may alter the system it studies. A protective measure may create reputational harm. A reversible-looking action may not be fully reversible for the affected locus.

Action under residual must therefore carry its reversibility profile.

Safeguard is the constraint placed on action while residual remains. It prevents residual from licensing unlimited action. A safeguard may limit scope, duration, affected-locus burden, automation, discretion, harm exposure, data use, escalation, or institutional authority. It may require human review, second measurement, appeal, consent, notice, transparency, independent oversight, monitoring, proportionality, staged intervention, error correction, domain restriction, sunset, or withdrawal condition.

Safeguard is not decorative reassurance. It is part of the warrant for acting under residual. If an action proceeds because closure is incomplete, the incompleteness cannot become permission for broad uncontrolled intervention. A risk that is not fully known cannot justify unlimited intervention merely because it is uncertain. A model that is not fully understood cannot command deference merely because it is technical. A category that remains unstable cannot become an enforcement trigger without constraint. Safeguard is the discipline that prevents action under uncertainty from becoming authority without correction.

Monitoring is the bridge between action and update.

A residual-bearing action should specify what will be watched, what counts as expected outcome, what counts as warning, what counts as failure, and what would force revision. Without monitoring, action under residual becomes uncontrolled commitment. The observer-system acts, but the action does not remain answerable. Monitoring turns residual-bearing action into an updateable process.

Stop conditions are equally important.

A residual-bearing action should state when it must stop, pause, reverse, narrow, escalate, or return to review. If the action has no stop condition, the residual has not truly been preserved. The action is behaving as though closure had been achieved. Stop conditions are the practical form of epistemic humility. They acknowledge that action is proceeding under incomplete closure and that the world may answer against it.

Action under unknown conditions also requires separation between action success and claim closure.

An action may work for reasons other than those assumed. A treatment may help by a different mechanism. A policy may succeed because of external conditions. A model may guide useful behaviour while misdescribing the target. A defensive response may avoid harm without proving that the perceived threat was real. Action success updates the field, but it does not

automatically close the explanation. The relation between action and outcome must itself be audited.

Action failure must also be interpreted carefully. Failure may exclude the formulation that justified the action. It may instead show that the action did not properly implement the formulation. It may reveal a scale gap, timing issue, boundary failure, wrong bearer, insufficient dose, missing interface, weak safeguard, changed environment, or unanticipated dependency. Failure is not automatically disproof, but it is always update pressure. Residual-bearing action must preserve enough structure to interpret failure without laundering it.

Action also changes the unknown.

Before action, the observer-system faces a field of possible relations. After action, one branch has been selected. Other branches may become unavailable. The action may create records, alter the target, produce feedback, change incentives, affect trust, transform the observer-system, or stabilise one formulation over others. The unknown after action is not the same as the unknown before action.

This is why action under residual must be explicitly marked. A disciplined action statement says: action is being taken under these conditions, for these reasons, with this residual, at this threshold, under these safeguards, with these monitoring paths, with these stop conditions, and with these update triggers. These elements keep the action Soft Horizon from becoming retrospective closure.

The strongest action under unknown conditions is not the action that denies uncertainty. It is the action that remains corrigible while acting. Corrigibility keeps an action Soft Horizon soft. The action may proceed, but it remains answerable to residual, monitoring, stop condition, affected-locus response, comparator pressure, and update.[6]

The discipline of the unknown therefore includes action as a central case. Most observer-systems do not merely contemplate the unknown. They live through it. They must continue, avoid harm, seek opportunity, allocate attention, protect boundaries, coordinate, and revise under feedback. The unknown is not only an epistemic condition.

It is an operating condition.

---

11. Stabilisation bias, live field contraction, and update

Once a branch stabilises, it can appear more inevitable than it was.

This is one of the deepest distortions in the management of the unknown. The formulation that survives becomes the formulation later observers inherit. The alternatives that failed, were excluded, were never tested, remained weakly live, lacked institutional support, or disappeared from records become harder to see. A Soft Horizon that stabilises may therefore be mistaken for the natural cut of the target rather than a provisional closure that became historically dominant.

Stabilisation bias occurs when the current structure of records overrepresents the winning path.

This can happen in perception, science, law, culture, institutions, personal memory, model development, formal systems, and governance. Once a provisional closure gains records, examples, terminology, procedures, incentives, institutional backing, machine implementation,

and retrospective explanation, it can appear target-complete. The residuals, alternatives, and branch conditions that made it provisional become less visible.

This matters because live formulation sets are often reconstructed after the fact. Later observers ask what alternatives existed, but the archive is already shaped by the branch that stabilised. Failed alternatives may look foolish because they failed. Unselected alternatives may look marginal because they were not preserved. Suppressed alternatives may look unsupported because the systems that would have generated their support were never built. Weakly live alternatives may look impossible because the language that could formulate them disappeared. The live field becomes distorted by survival.[7]

Live Field Contraction is the process by which alternatives cease to remain live within the field under review.

In the Part's counterfactual architecture, Live Field Contraction extends the Counterfactual Comparison Discipline established in Paper 1 into formulation history: the question is no longer only which alternatives test a closure, but which formulations remained admissible, how they ceased to remain live, and whether that contraction was earned by comparator pressure or produced by historical disappearance. Its object remains the live formulation field of the observer-system. Presence in, or removal from, that field does not establish that an alternative was physically available as a causal branch for the evaluated locus.

Some exclusions are earned. A formulation fails comparator pressure. A prediction breaks. A mechanism cannot localise. A scale bridge collapses. A category cannot preserve the relation. A model fails deployment. A legal argument fails a required burden. These exclusions should remain exclusions unless the conditions that produced them change.

Other exclusions are historical rather than epistemic. A formulation may be abandoned because it lacked resources, because a dominant institution preferred another path, because an instrument was not yet available, because a language lacked a category, because action selected one branch, because a procedure could not receive the evidence, or because preserving alternatives was too costly. These are not equivalent to earned exclusion.

The discipline of the unknown requires asking how alternatives left the field.

Did they fail? Were they never properly tested? Were they excluded by a closure demand they were not designed to satisfy? Were they displaced by action? Were they suppressed by institutional form? Were they invisible because the abstraction regime could not formulate them? Were they weak in themselves, or weak relative to the access conditions of the time? Did they lack evidence because the evidence path was unavailable? Did they appear implausible because the target relation was being cut incorrectly?

These questions do not revive every abandoned alternative. That would be false equivalence. Many alternatives are properly excluded. Many are malformed. Many are too weak to carry forward. The point is not to repopulate the field indiscriminately. The point is to distinguish epistemic exclusion from historical disappearance. Stabilisation is evidence, but it is not always proof of superiority. Persistence may reflect target constraint, but it may also reflect path dependence.

Stabilisation bias also affects categories.

A category that becomes standard can appear to have always been the natural cut. Later users forget the alternatives the category suppressed. The category then carries authority from stability rather than from continuing relation-preservation. Unknown relations at the category boundary disappear because the category has become familiar. A disciplined observer-system should periodically ask what a category still fails to preserve, which affected loci challenge it,

what residual remains at its edges, and whether its original abstraction still matches its current use.

Scientific models are vulnerable in the same way. A model that works across many domains gains legitimate strength. But legitimate strength can turn into reification if the model's residual, scale limits, comparator field, and excluded alternatives disappear. A model may become the grammar of a discipline. Questions that fit it become natural. Questions that do not fit it become marginal. Over time, the model may shape the availability of the field it claims to describe.

Institutional records show the same structure. Once a person, event, risk, harm, or status is recorded under a category, later systems inherit the category more easily than the uncertainty that accompanied it. The record becomes clean. The residual becomes invisible. The chosen branch becomes administratively real. This is why weighted residual and action records matter: they preserve what was unresolved at the time the branch stabilised.

Update is the discipline that keeps stabilisation from becoming final possession.

An update trigger is a condition that forces the observer-system to reopen, revise, narrow, replace, suspend, or reweight a formulation, residual, live field, category, action, or Soft Horizon. It is not a general willingness to learn. It is a specified route by which the world, comparator, affected locus, record, outcome, or abstraction can answer back.

Update triggers include new evidence, failed prediction, comparator failure, affected-locus challenge, model drift, instrument recalibration, changed scale, category failure, action harm, monitoring warning, stop-condition activation, record correction, legal or procedural change, formal contradiction, previously unavailable evidence, residual resolution, or the appearance of a stronger formulation. The exact trigger depends on the relation and closure demand.

Without update triggers, residual becomes inert. The system admits uncertainty but has no condition under which uncertainty can change the claim. This is temporal laundering. A formulation is provisional when adopted, but permanent in practice. A decision is precautionary when made, but maintained without renewed threshold. A model is experimental when deployed, but normalised without review. A category is unstable when introduced, but archived as settled. A safeguard is promised, but no one can trigger it. A residual is noted, but cannot move anything.

Update keeps the unknown active.

A disciplined update does not treat revision as embarrassment. It treats revision as the expected consequence of finite relation. If new availability changes the field, the field should change. If a comparator fails, the weighting should change. If an affected locus exposes a category failure, the abstraction should be reopened. If monitoring shows harm, the action should stop, narrow, reverse, or escalate. If a residual becomes blocking, the Soft Horizon should lose strength. If an excluded formulation becomes newly admissible because conditions changed, the exclusion should be reviewed.

This does not mean every challenge forces revision. A system can distinguish noise from update. A claim does not become weak merely because someone contests it. An excluded formulation does not become live merely because it is reasserted. A Best Available Formulation does not lose strength whenever a speculative alternative appears. Update discipline requires comparator pressure. The trigger must bear on the relation under review.

The horizon moves when availability, abstraction, comparator, scale, residual, or closure demand changes.

An instrument may make previously unavailable relation available. A new abstraction may make a formerly vague relation formulable. A stronger comparator may exclude a candidate. A failed action may reveal a hidden dependency. A new affected-locus challenge may show that the category does not preserve the relation it governs. A model drift signal may show that past performance no longer applies. A formal result may show that a closure demand cannot be met from within the present system. A record correction may alter the evidential field.

When the horizon moves, the observer-system must not pretend that the old Soft Horizon remains unchanged. It may remain usable, but only under renewed conditions. It may narrow. It may weaken. It may be replaced. It may become positively excluded. It may become a Best Available Formulation of lower strength. It may become an historical formulation retained for branch-memory but no longer carried forward as current.

Update therefore completes the discipline of the unknown.

Access humility handles relation outside availability. Closure discipline handles relation available but unclosed. Failure-mode discipline prevents unknown-status collapse. Residual localises what remains unresolved. Positive exclusion removes failed formulations. Live fields structure admissible alternatives. Weighting orders non-closure. Best Available Formulation permits continuation. Weighted residual preserves the unresolved field. Action threshold permits action without action-certainty. Safeguard, monitoring, and stop condition keep action corrigible. Stabilisation-bias discipline preserves branch-memory across Live Field Contraction. Update keeps the system answerable when the horizon moves.

The unknown is disciplined only when it can change the future claim.

The live formulation field developed here remains a field of finite observation, formulation, residual, and action under incomplete closure. Its counterfactual function is epistemic and historical: it preserves which formulations remained live, how the field contracted, and why alternatives were excluded, weakened, abandoned, or lost. It must not be identified with the physically supported possibility field of an Information Processing Locus. A live formulation may remain epistemically available without being an operative physical alternative for the locus whose causal potential is under review; conversely, a locus may possess physically available branch-changing alternatives without representing them as live formulations. The two relations therefore require separate discriminator conditions.

The next paper begins at that boundary. Once physically supported distinction-conditioned organisation makes counterfactually different downstream branch-realisation available to a specified Information Processing Locus, counterfactual structure becomes causal rather than merely epistemic. The question is no longer only how a finite observer-system carries alternative formulations, but what branch-changing potential the locus physically possesses and what follows when some part of that potential is realised. That is the primitive relation developed as Power.

---

12. Conclusion: the unknown as the main field of finite Observation

The unknown is not a peripheral residue around knowledge.

It is the main field within which finite observer-systems operate.

A finite observer-system does not possess reality as total state. It encounters relation through boundary, interface, scale, horizon, discrimination regime, abstraction, registration, retention, redeployment, comparator pressure, closure demand, and residual limit. It lives through Soft Horizons: provisional closures that make relation usable for recognition, classification, modelling, communication, judgement, and action while stronger closure remains unavailable.

The unknown is therefore not blank absence. It appears as relation outside availability, relation available before recognition, relation misclosed by pattern, relation provisionally closed for one use while residual for another, relation localised as residual, relation held among live formulations, relation weighted without final settlement, relation carried as Best Available Formulation, and relation acted under before closure.

Each form requires its own discipline.

Unavailable relation requires access humility. Available-but-unclosed relation requires closure discipline. Misclosure requires correction. Residual requires localisation, limitation, blocking, and update orientation. Live formulation requires admissible alternatives. Positive exclusion requires retained reasons for failed candidates. Weighting requires proportional support without false precision or false equality. Best Available Formulation requires the strongest current Soft Horizon of formulation with residual intact. Weighted residual requires memory of what must not be forgotten. Action requires threshold, safeguard, monitoring, stop condition, and update. Stabilisation requires branch-memory. Update requires conditions under which the horizon can move.

This structure is not anthropocentric. Human observers are one class of finite observer-system. The same architecture applies, in different forms, wherever a finite locus encounters relation through boundary, interface, scale, discrimination, registration, retention, deployment, and correction. A cell, organism, instrument, computational system, legal institution, scientific practice, social system, formal community, or coupled observer-network can encounter unknowns relative to its regime. The discipline does not project human consciousness onto all such systems. It removes human consciousness from the foundation of unknownness.

The unknown is also not anti-realist. The target constrains. Observer-systems do not invent the consequences that correct them. Models fail. Measurements drift. Categories break. Organisms die. Instruments saturate. Predictions miss. Actions harm. Proofs do not follow. Residuals become consequential. The relation is indexed, but the constraint is real. The discipline of the unknown preserves realism by refusing to replace the target with an overclosed formulation.

The unknown is not a failure of method. It is the condition that requires method.

A method that can only handle closure is not sufficient for finite observers. A method that can only suspend but cannot weight will paralyse. A method that can only act by pretending certainty will become dangerous. A method that can only preserve residual but cannot formulate Best Available Formulation will fail continuation. A method adequate to finite relation must be able to carry the unknown without dissolving it.

The proper strength of a finite observer-system is therefore not total possession. It is disciplined exposure to correction.

Such a system can say what is available, what is unavailable, what is registered, what is retained, what is deployed, what is misclosed, what is excluded, what remains live, what is weighted, what is best available, what remains residual, what action threshold has been reached, what safeguard constrains action, what monitoring is required, what stop condition applies, and what would force update.

It does not escape the unknown.

It becomes accountable to it.

The discipline of the unknown can be compressed into one final rule:

No unknown-status is disciplined until relation, locus, availability condition, and closure demand are visible. Do not harden what is only provisionally closed. Do not close what remains residual. Do not mystify what can be localised. Do not flatten what can be weighted. Do not act as if action were certainty. Do not forget what would force update.

A finite observer does not become disciplined by knowing everything.

It becomes disciplined by preserving the conditions under which the unknown can answer back.

Endnotes

[1] Unknown-status comparators: Michael Smithson, Ignorance and Uncertainty: Emerging Paradigms (1989), and Matthias Gross, 'The Unknown in Process' (2007), both resist treating ignorance or non-knowledge as one undifferentiated absence and develop differentiated or dynamic forms of not-knowing. The present paper makes a different technical move: unknown-status is indexed to relation, locus, availability condition, and closure demand, and is generalised across finite observer-systems rather than confined to human judgment or sociology. No technical inheritance is claimed.

[2] Historical provenance: Robert K. Merton, 'Three Fragments From a Sociologist's Notebooks' (1987), develops specified ignorance as a way of making what is not known explicit and inquiry-directing. Residual here is narrower and more operational: it localises the unmet condition under a stated closure demand, states what stronger claim is blocked, and carries an update route. The relation is historical provenance rather than dependency.

[3] Historical methodological precursor: T. C. Chamberlin, 'The Method of Multiple Working Hypotheses' (1890), argues for carrying multiple explanatory hypotheses rather than allowing one ruling hypothesis to dominate inquiry. The present live-formulation architecture adds admissibility, positive exclusion, heterogeneous weighting, residual, and explicit branch-memory; Chamberlin's method is not adopted as a technical dependency.

[4] Comparator: Peter Lipton, Inference to the Best Explanation (2nd ed., 2004), develops inference by selecting the explanation that would, if true, best explain the evidence. Best Available Formulation is deliberately broader and weaker in closure: it may be typed to action, prediction, classification, monitoring, or explanation; it is selected across heterogeneous support relations; and it must carry residual and update conditions. BAF therefore does not inherit truth by explanatory bestness.

[5] Inductive-risk comparator: Richard Rudner, 'The Scientist Qua Scientist Makes Value Judgments' (1953), and Heather Douglas, 'Inductive Risk and Values in Science' (2000), make the consequences of error relevant to scientific acceptance or methodological judgment. The present paper preserves a sharper status boundary: consequence may alter the threshold for action under residual without thereby satisfying the stronger closure demand of truth, explanation, or ontology. The comparison is discriminatory, not inherited.

[6] Decision-under-deep-uncertainty comparator: Robert J. Lempert, Steven W. Popper, and Steven C. Bankes, Shaping the Next One Hundred Years (2003), develop robust and adaptive policy analysis across multiple plausible futures rather than relying on a single forecast. The present paper shares the requirement that action remain corrigible under unresolved futures, but does not inherit RAND's quantitative decision method; threshold, safeguard, monitoring, stop condition, and update are defined as general relation-preserving burdens.

[7] Stabilisation and live-field comparators: Baruch Fischhoff, 'Hindsight ≠ Foresight' (1975), shows that outcome knowledge can make the realised branch appear more predictable retrospectively and distort reconstruction of what was knowable beforehand. Thomas S. Kuhn, The Structure of Scientific Revolutions (2nd ed., 1970), provides a science-historical specialisation in which paradigm stabilisation and retrospective disciplinary reconstruction can make development appear more linear and inevitable than it was. P. Kyle Stanford, Exceeding Our Grasp (2006), supplies complementary pressure against treating elimination of the alternatives actually conceived as proof that every serious alternative has been

exhausted. Stabilisation Bias and Live Field Contraction generalise these relations across finite closure regimes, distinguish earned positive exclusion from forgetting, suppression, non-formulation, and branch loss, and preserve the possibility that the finite live field itself was incomplete. None of these external theories is inherited.

References

Chamberlin, T. C. (1890). The method of multiple working hypotheses. Science, 15(366), 92-96. https://doi.org/10.1126/science.ns-15.366.92

Douglas, H. E. (2000). Inductive risk and values in science. Philosophy of Science, 67(4), 559-579. https://doi.org/10.1086/392855

Fischhoff, B. (1975). Hindsight ≠ foresight: The effect of outcome knowledge on judgment under uncertainty. Journal of Experimental Psychology: Human Perception and Performance, 1(3), 288-299. https://doi.org/10.1037/0096-1523.1.3.288

Gross, M. (2007). The unknown in process: Dynamic connections of ignorance, non-knowledge and related concepts. Current Sociology, 55(5), 742-759. https://doi.org/10.1177/0011392107079928

Kuhn, T. S. (1970). The structure of scientific revolutions (2nd ed., enlarged). University of Chicago Press.

Lempert, R. J., Popper, S. W., & Bankes, S. C. (2003). Shaping the next one hundred years: New methods for quantitative, long-term policy analysis. RAND Corporation. https://doi.org/10.7249/MR1626

Lipton, P. (2004). Inference to the best explanation (2nd ed.). Routledge.

Merton, R. K. (1987). Three fragments from a sociologist's notebooks: Establishing the phenomenon, specified ignorance, and strategic research materials. Annual Review of Sociology, 13, 1-29. https://doi.org/10.1146/annurev.so.13.080187.000245

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

Smithson, M. (1989). Ignorance and uncertainty: Emerging paradigms. Springer-Verlag. https://doi.org/10.1007/978-1-4612-3628-3

Stanford, P. K. (2006). Exceeding our grasp: Science, history, and the problem of unconceived alternatives. Oxford University Press.