0.0 Title: the persistent metaphysical trace of prior causation, part one: a foundational theory of information

Author: Luke Tredinnick, l.tredinnick@londonmet.ac.uk

0.1 Abstract: this paper develops a foundational theory of information as the persistent trace of prior causation, realised as a trajectory within transformation space. The theory is predicated on the minimal assumption that information exists. From that assumption it derives the conditions required for informational existence, and recursively redescribes them through relation, scale, transformation, subjectivity, and intersubjectivity. Although the paper is narrow in scope, its potential reach is broad. This is because information functions as an implicit or explicit foundational concept across contemporary theoretical discourse. This reach does not indicate the unrestricted breadth of the theory, but the significance of the concept it addresses and the degree to which that concept remains under-theorised. The paper uses a derivational approach from minimal assumptions under methodological, ontological, and epistemological constraint. Its major findings concern the nature of information as trace, the emergence of meaning through the actualisation of inferential potential within a subjective domain, and the significance of information as a critical and theoretical concept. The paper is situated within the information science tradition, but has relevance across disciplines in which information functions as a foundational concept.

0.2 Status of the manuscript: this paper is circulated as a preprint version of an ongoing theoretical project. The tracework framework and its methodological constraints are presented as fixed, but the exposition, terminology, internal cross-references, and articulation of specific derivations may be refined in later versions.

© Luke Tredinnick 2026. This work is licensed under a Creative Commons Attribution 4.0 International Licence (CC BY 4.0). Full text available: DOI: https://doi.org/10.5281/zenodo.21157999

0.3 Biographical note: Luke Tredinnick is a Reader in the School of Computing and Digital Media at London Metropolitan University. His research explores the theoretical foundations of information, with particular interests in meaning, digital culture, artificial intelligence, creativity, and the social organisation of knowledge. He has written on digital information environments, generative artificial intelligence, and the changing relationship between technology, culture, and professional information practice. **ORCID:** 0000-0001-6947-283X.

Preprint

0.4 Changelog: preprint version 1.0 July 2026; v. 1.1 July 2026 [copyright statement and DOI]; v. 1.2 August 2026 [memory added to s. 3.3; retrodictive failure mode revised s. 8.3; bridge structure added to s. 3.4; minor clarification s. 7.8; note on working notes added s. 0.5]; v. 1.3 [MONTH] 2026 [Shannon equivalence disambiguation]; v. 1.4 October 2026 [clarification of stabilisation domain s. 3.2; minor disambiguation of section and subsection references; other minor corrections].

0.5 Companion papers: a formal companion to this paper exists as a separate working document, developed independently of the derivation presented here. It develops the architecture of sections 2–7 as explicit set-theoretic and relational constructions, tested adversarially against constructed counterexamples. Its purpose is not to certify the theory but to test whether its architecture survives formal pressure.

Preprint

1. The problem of information

1.1 Introduction: quod stat, est. This paper presents a foundational theory of information that is intended to address the fragmentation of information as a theoretical and conceptual object. It is minimal in the strictest sense, assuming only that information exists. The derivational argument that follows shows that if information exists, then its existence can be accounted for through production, persistence, and transformation. Because production and persistence are not equivalent operations, transformation is asymmetrical. Given this, the paper situates information as a consequence of causal chains whose residues remain open to reconstruction when stabilised through a recursive process termed generative tracework. Information exists as the persistent metaphysical trace of prior causation (see Tredinnick, 2025) realised as a trajectory within transformation space. Like trails in a cloud chamber, it is a temporary structural formation that retains a glimmer of a prior causal event already absent at the moment the trace forms. What follows establishes that structure.

1.2 Information as a fragmented concept: the central problem addressed by this paper is the fragmentation of information as a theoretical and conceptual object, and the concomitant fragmentation of those discursive domains within which assumptions about the nature of information have become foundational. The nature of information has been called both "the fundamental problem of information science" (Brookes, 1975) and "the hardest and most central problem" in the philosophy of information (Floridi, 2013: 30). Yet despite its defining role in contemporary discourse, we lack a unified account of information's ontological status and affordances. Hjørland, for example, has described information as "a polysemic and vague word" (Hjørland, 2023). Krzanowski has observed that "we still have a rather vague understanding of what information is" (Krzanowski, 2023: 1), and Beavers that "we can scarcely say precisely what the term means" (Beavers, 2016: 16). The long tradition of theoretical debate has not resolved this problem; fragmentation is a consistent preoccupation in the theory of information (e.g. Fairthorne, 1967; Belkin, 1978; Qvortrup, 1993; Cornelius, 2002; Capurro and Hjørland, 2003; Zins, 2007; Logan, 2012). Consequently, information has been theorised in markedly different ways in different discursive domains, and the resulting descriptions have often appeared incommensurable (Capurro and Hjørland, 2003; Ellis, 1992; Raber, 2003; Tredinnick, 2006, 2015, 2025; Krzanowski, 2023). This implies that the problem of fragmentation may lie at a more

fundamental level than previously addressed, with significant consequences across a wide range of domains in which information is assumed to play a central role.

A foundational theory of information must do two things. It must specify the conditions under which anything can be considered informational, and it must be capable of accounting for those phenomena widely regarded as informational across the diverse contexts in which the concept is deployed (Tredinnick, 2025). Both parts of this definition are integral. Theories that cannot specify the conditions under which information exists lack precision, and those that cannot account for its diverse expressions lack sufficiency. Scepticism towards the possibility of such a theory has been commonplace, generally reflecting doubts about sufficiency. Shannon, for example, argued that "it is hardly to be expected that a single concept of information would satisfactorily account for the numerous possible applications of this general field" (Shannon, 1953: 105), and Floridi has questioned the viability of a "grand unified theory," citing the semantic richness, contextual variability, and disciplinary divergence that characterise informational concepts (Floridi, 2013: 30). Indeed, many of the attempts to unify the theory of information have also focused on sufficiency over precision, specifying what information does rather than what information is, as if the latter were already resolved. The result has been a proliferation of motivated definitions that outpaces their ontological grounding (Yovits, 1969: 6; Belkin, 1978: 55; van Rijsbergen, 1996: 385; Burgin, 2003: 148; Capurro and Hjørland, 2003: 349; cf. Tredinnick, 2025). This paper proposes a unifying ontological structure operating beneath prior accounts and from which both precision and sufficiency can be derived.

1.3 Theoretical framing: the analysis in this paper follows a prior thematic review (Tredinnick, 2025). That thematic review sought to map the boundaries of a foundational theory and provides the fuller theoretical context for this paper. It identified three areas that currently bind the concept of information: what is known or generally agreed; what is unknown or contested; and what appears intractable. It is known or generally agreed that: 1) information is a real phenomenon with material instantiation rather than wholly abstract or ideal; 2) information carries, transmits, or conveys meaning; 3) information is related to knowledge or knowing; 4) difference plays a decisive role in what information is; 5) information has subjective and objective characteristics; 6) its objective characteristics are independent of observation; and 7) Shannon's account of information has decisive reach. It is unknown or contested: 8) whether information is a thing, a structure, or a process; 9) whether information is intrinsically or situationally

meaningful; 10) whether data, information and knowledge are distinct, related, or identical phenomena; 11) what difference is and what role it plays; 12) how information's objective and subjective characteristics interact; 13) what role time plays in information; and 14) whether information's material instantiation matters on an epistemic level. Currently intractable problems include: 15) the precise nature of information; 16) how meaning arises in informational things; 17) why objective and subjective characteristics appear incommensurable; and 18) what role computational accounts have in the nature of information. The review also identified significant gaps in the tradition, particularly in respect of the role of time and the inability of prior theories to account for the structural reach of Shannon's framework. A fuller discussion of how these issues emerge, and how they have been addressed in the literature can be found in the prior thematic review (Tredinnick, 2025).

To these issues can be added a further set of boundary problems drawn from Floridi's open problems in the philosophy of information (Floridi, 2004a; cf. Crnkovic and Hofkirchner, 2011). Floridi identified eighteen open problems across five areas: the analysis of the concept of information, semantics, intelligence, the relation between information and nature, and values. Many of these are already implied by the thematic boundaries listed above, including the nature of information, the emergence of meaning, the relationship between information and knowledge, and the objective status of environmental information. The remaining problems most directly relevant to the present account include: 19) whether a grand unified theory of information is possible; 20) the problem of false information; 21) levels of abstraction; 22) the value of information; 23) whether natural intelligence can be implemented non-biologically; and 24) the philosophical grounding of information ethics. In combination, these issues provide a comprehensive mapping of the boundaries of a foundational theory of information. They specify what such a theory should be expected to resolve, or show to be irresolvable. This represents a genuinely high bar for any theory of information to surpass, drawn from the literature independently of this paper.

The unresolved issues in the nature of information mark thresholds where prior theories of information reach their explanatory limit (Tredinnick, 2025). In other words, they demarcate the boundaries where false assumptions may be doing additional work. No prior theories account for them fully. Shannon's framework, for example, is formally precise and operationally powerful, but it brackets semantic content and treats temporality primarily as a statistical parameter (Shannon, 1948). Floridi's account addresses semantic meaning but relies on truth as a foundational constraint (Floridi, 2003; 2004b; 2010; 2013). Existing attempts at unification have

generally also failed to reconcile these tensions. Burgin's general theory of information, for example, provides a classificatory unification across multiple conceptions and thus inherits rather than explains the assumptions embedded in the theories it systematises (Burgin, 2003; 2004; 2005; 2009). Brier's cybersemiotics integrates semiotics and information theory but remains primarily interpretive, lacking both formal constraints and material instantiation (Brier, 2008). Each prior theoretical account succeeds insofar as it strictly delimits the problem space, bracketing aspects of information that fall outside its explanatory frame. This perhaps explains their apparent incommensurability, implied by Capurro's trilemma (Capurro et al., 1997) and Ellis' irreducible duality of research paradigms (Ellis, 1992). Different theories reveal different facets of a single theoretical object that in its totality remains occluded. A fuller account is provided in the prior thematic review (Tredinnick, 2025).

Tracework approaches the problem at a different level of explanation. It proceeds from the single minimal assumption that information exists. This assumption does not specify what information is, what form it takes, or what properties it has. It requires only that, if information exists, there must be conditions under which it is informational rather than not informational. From this the derivation specifies the conditions under which that assumption can hold [cf. s. 2]. That derivation yields a number of named principles, each following from what precedes it.

The first of these is the principle of informational grounding. Traditionally, information has been constituted in already-given elements: symbols, data, and discrete states. But if information exists, something must account for its existence. To treat it as simply given presupposes the very conditions that make it identifiable as information in the first place. That something exists which is informational in kind already requires persistence, differentiation, and structured consequence. These cannot be assumed without also assuming the conditions under which they come to be. To assume already-stabilised primitives is to assume an unexamined prior process. The minimal assumption that information exists therefore entails that the conditions under which anything can exist informationally must be specified. This shifts the foundational question from what information is to how information comes to be. Rather than starting from already-given elements and analysing their properties, tracework starts from the conditions that must be met for persistent structured consequences capable of functioning informationally to be produced. In tracework, information is produced when causal processes leave sufficiently stabilised structure for generative tracework to begin.

Preprint

Production is a process that unfolds in time. Informational grounding therefore requires that information is inherently temporal: it must exist for an interval sufficient for it to function informationally with respect to its grounding, or it cannot exist as information nor have been produced as such. It must also continue to exist over time because existence outside time, or existence for no time, cannot satisfy the minimum criteria of persistence, differentiation, and structured consequence [cf. s. 2.4]. This is the principle of temporal extension [s. 2.6]. Standard accounts generally obscure this by beginning from already-stabilised informational primitives. Once these are assumed, time can no longer be constitutive of their formation, but can only appear externally, as the interval between discrete states or their sequence. The result is a kind of quantised temporality in which information is locked into sequential fixed states. Information is treated as if it were first available as a fixed state or structure, and only afterwards placed into temporal relation, and temporal ordering is therefore incidental rather than constitutive. Even theories with explicit temporal constraints tend to retain this fundamentally synchronic picture. Shannon entropy characterises a probability distribution as a time-slice (Shannon, 1948); the General Definition of Information characterises a data structure synchronically (Floridi, 2013); and Dretske's causal account treats the signal as a carrier of information about its source at the point of transmission (Dretske, 1981; 1983). Tracework inverts this implicit temporal framing. Information is not first given and then moved through time; it is stabilised through temporally extended causal interaction. A static description of information is always a cross-section of something fundamentally extended in time, a point sampled from its trajectory, and therefore incomplete because it abstracts away the process by which the trace was produced, stabilised, and made inferentially available. Time is not the medium through which information moves but the dimension within which it exists, and within which its full extension must be defined.

Because production and persistence are not equivalent operations, transformation is directionally ordered. A trace stands in a different relation to what precedes it than to what follows from it. What is fixed at formation cannot be recovered by running formation in reverse: the forward and backward relations are not two descriptions of a single operation but are distinct. This distinction specifies the productive and inferential orientations in the temporal extension of information. Under the productive orientation, a trace constrains what can follow from it; under the inferential orientation it constrains what could have preceded it. Orientation therefore determines which operations are available rather than merely selecting among descriptions of an invariant object. This is the principle of orientational non-equivalence, given here in restricted form; its full generality is derived later [s. 2.6] and developed across the scaling, transformation, and subjective architectures [cf. s. 3-7].

The epistemic expression of non-equivalence is the principle of epistemic asymmetry. If information exists only as the extended consequence of prior causal processes, then it is produced and therefore not eternal: the state before production differs from the state after. The causal processes that produce information are complete at the moment information exists, so inference about those causes cannot follow the direction of their production; it must proceed backwards, from information towards the conditions that produced it. Because production unfolds forwards while inference proceeds backwards, any non-stipulative account of information grounded in causation must derive its structure from this asymmetry in the relationship between information and time. This is the principle of epistemic asymmetry [s. 2.6].

Epistemic asymmetry and orientational non-equivalence jointly entail another consequence. If inference proceeds from trace to cause while production proceeds from cause to trace, then the epistemic starting point is the trace rather than the source. And because orientation is operationally determinative rather than merely a matter of perspective, the productive and inferential orientations cannot be treated as equivalent operations in reverse: what can be recovered through backward inference is not what was fixed at the point of forward production. Standard accounts of information treat the source as prior both ontically and epistemically: the source precedes the trace causally and is taken to ground meaning. Tracework by contrast differentiates between these two operations: while on an ontic level cause precedes trace, on an epistemic level trace precedes cause. The source is not the origin from which information flows but a causal condition that remains, in principle, reconstructable in the trace. Information is not encoded by a sender and decoded by a receiver. It is the reconstruction of the causal conditions that produced the trace, of which communicative intention is one instance. The arrow of epistemic priority therefore runs from consequence to cause rather than from cause to consequence. This is the principle of epistemic priority inversion [s. 2.6], the symmetry break through which tracework expresses its structure.

If epistemic priority runs from trace to cause, then what is fixed in production cannot determine what is reconstructed in inference. This is the principle of informational intransitivity [s. 2.6]. Standard accounts treat information as transitive: encoded at a source, transmitted through a channel, and received at a destination. Tracework rejects this. Information is not transmitted as self-identical content. It propagates through transformation. What appears as transmission is the preservation of sufficient causal history across a sequence of transformations to maintain inferential potential at each stage. Physical carriers may move in space, but spatial movement is itself a transformation [cf. s. 3.4]. Information persists only where enough causal history is

preserved and re-stabilised under the altered conditions into which it moves. Transmission is therefore not a fundamental operation but an appearance generated by the preservation of inferential potential through transformation.

Finally, if information is always locally constituted through stabilisation rather than transmitted from a source, then absence is itself a stabilised condition capable of sustaining inferential potential. This is the principle of presence-independence [s. 2.4]. A universal assumption has been that information is a property of presence. Signals, data, marks, and the elementary primitives of prior theories are all presences, and absence is treated as information's limit rather than one of its modes. This reflects a metaphysics of presence that is arguably deeply and historically ingrained in the Western intellectual tradition (Derrida, 1967). Tracework specifies the conditions under which information forms, persists, and sustains inferential potential irrespective of presence. Wherever absences are the stabilised consequence of prior causal conditions, they become informational. Absences may also sustain inference about their causal conditions under the right constraints, as they remain materially determined, such as a hole in a punched card. Each absence remains a nexus within a relational network. The consequences of this principle extend across physical, legal, archival, and documentary domains, but are also fundamental to tracework: information is not inherently present; it may also be structurally absent [cf. s. 5.6].

Together, these principles articulate the theoretical architecture within which generative tracework operates. Each follows from the minimal conditions required for the assumption that information exists to hold [cf. s. 2.4; s. 2.5]. The fragmentation of information as a critical and theoretical object reflects not its implicit complexity, but the ways in which deeply embedded assumptions have shaped the field. Tracework shows these commitments to be downstream of information’s fundamental nature. Information is simple in nature, inexhaustibly complex only in its effects. Prior theories of information can therefore be understood as operating within distinct restricted regimes of description, their apparent incommensurability arising from differences in the level at which they articulate a broader structural process.

Tracework owes particular recognition to Dretske's account of information, particularly in respect of the role of causation (Dretske, 1981;1983). Tracework inherits its commitment to causal grounding and its rejection of Shannon's reductionism. However, its account also departs in significant ways. Tracework does not treat meaning and intentionality as the primary target of a foundational theory of information. Prior to matters of semantic interpretation lies the problem

of the conditions under which information exists at all. Whereas Dretske treats source as prior to signal, tracework treats signal as epistemically prior to source. Tracework remains agnostic about which theory of productive causation is correct, whereas Dretske presupposes a more determinate causal relation. The disjunction problem that arises within Dretske's account is shown to be a consequence of these presuppositions rather than a feature of causal information, and is addressed directly [s. 2.7].

Tracework is governed by three constitutive methodological commitments that form a triple lock on methodological closure addressing ontological, epistemological and methodological limits respectively. First, it makes only those metaphysical commitments required for the coherent description of tracework itself, introducing no ontological categories beyond those required to account for stabilisation, persistence, and inferential legibility under temporal asymmetry. Thus it is the minimal possible theory of information. Second, it makes claims only about phenomena defined within tracework, declining extension to domains not already characterised in terms of traces, trajectories, and admissibility. Thus it is only a theory of information. Third, every claim advanced within the framework is constrained to follow from the assumption that information exists and the minimal stabilising conditions for its possibility derived in section 2. No additional principles are introduced to secure particular results. Under this triple lock of methodological closure, the framework admits no extension by conjecture. Claims either follow as consequences of minimal constraint structure or fall outside its scope.

1.4 Aims and structure: this paper develops a foundational theory of information as a trajectory in transformation space grounded in the persistent metaphysical trace of prior causation, where a trajectory denotes an ordered, constraint-governed sequence of transformations. The architecture of the theory is developed across six levels of recursive organisation. Each of these redescribes the same structure under progressively more complex conditions of its expression. Section 2 introduces the persistent trace of prior causation as a minimal ontological structure and derives from it the architecture required to sustain inferential potential about prior causal conditions. Section 3 addresses relational organisation, deriving the conditions under which traces become organised into trace-structures through coupling and recursive plasticity. Section 4 outlines the scaling dynamics of tracework, addressing how trace-structures are organised across viability domains through coarse-graining and re-stabilisation, and deriving the syntagmatic and paradigmatic dimensions of organised tracework. Section 5 outlines

transformation, deriving the geometry of transformation space, the admissible forms of trace transformation and trajectory organisation, and the thresholds and decay modes that define the limits of admissible movement. Section 6 redescribes the prior sections under reflexive self-regulation, deriving meaning and subjectivity as structural consequences of tracework operating over its own trace-structure at sufficient recursive depth. Section 7 derives the intersubjective domain, and explores its role in regulating tracework. Section 8 closes the theoretical account by deriving the principle of non-self-grounding, specifying the four irreducible failure modes of tracework, and testing tracework against the boundaries of a foundational theory established above. The paper concludes by situating information science within the landscape the theory defines [s. 9]. It is supported by a formal derivation of the two-dimensionality of transformation space (Appendix A); a glossary of original theoretical terms (Appendix B); and a named principles index (Appendix C). Part two, forthcoming as a separate preprint, addresses demonstrations.

Preprint

2. The persistent trace

2.1 Introduction: this section establishes the conceptual core of tracework: the conditions under which information exists in kind. That core consists in a single ontological structure, the trace, from which the minimal architecture required to describe the temporal extension of information through stabilisation is derived. A trace can form and persist only under temporal asymmetry. Tracework is therefore dynamic in character, describing processes that unfold under temporally asymmetric constraint through stabilisation, transformation, and failure.

2.2 From causal interactions to stabilised difference: Bateson's aphorism that information is "a difference that makes a difference" (Bateson, 1972: 460) has become one of the most widely cited definitions in information science (Duguid, 2021: x). Whereas Shannon defined the bit with mathematical precision (Shannon, 1948), Bateson reached for something more indeterminate: a distinction sitting prior to form, magnitude, or encoding. Difference, in this sense, was consequential prior to its definition as an informational object. Yet the meanings implied by difference across the contexts in which it has been mobilised vary, and there has been a relative lack of specific ontological and epistemological grounding. Even Bateson himself uses difference in three distinct senses in the aphorism (Tredinnick, 2025). As a consequence, a seductively economical concept breaks down under ontological heterogeneity that masks divergent assumptions.

In the tradition that followed, difference generally came to function as a minimal ontological primitive: a definition of what information *is* in fundamental terms rather than a principle from which that definition should follow. Theories of information that address pattern (Bates, 2005; 2006; Burgin, 2009; Floridi, 2003), structure (Wersig and Neveling, 1975; Stonier, 1990; 1991; 1992; Mingers, 1999; 2004; Burgin, 2009; Hofkirchner, 2009; 2012; 2013), or form (Bateson, 1972; Ruyer, 2024) tended to emphasise how defined differences combine into ontologically significant higher-order structures. What was lost in this move was the understanding that pure difference in Bateson's terms was not reducible to a stabilised distinction, but was circling an anti-reductionist constraint on definition without quite being able to articulate what followed. Although it remained nebulous, Bateson's difference displaced a central theoretical problem. An important insight was lost in its reification in already-stabilised primitives.

Preprint

Theories that begin from stabilised differences, whether that prior stabilisation is explicit or only implied, presuppose the process they need to explain. Already-stabilised primitives are commonplace in the discourse on the nature of information. Examples include the hartley (Hartley, 1928), ban (Turing, 1941), bit (Shannon, 1948), negative entropy (Wiener, 1961; Brillouin, 1962), natural sign (Dretske, 1981), sign (Brier, 2008), and data (Floridi, 2013). Although their derivations differ, each assumes an unarticulated prior process by which those primitives come to be. That assumption implies an underlying structural correspondence between them, described here as the principle of Shannon availability equivalence. It states that any theory of information grounded in already-stabilised elementary units inherits the formal structure of the bit at the level of informational availability, regardless of the theoretical vocabulary in which it is expressed, the means of derivation, or the semantic conditions subsequently added. Once treated as already-stabilised, any such informational concept assumes the formal structure of a bit: a discriminable difference between alternatives. It may subsequently be elaborated as sign, datum, natural sign, semantic content, negative entropy, or meaningful form, but its informational availability already depends on its discernibility in binary terms. The fundamental units in use can always be redescribed at the level of formal informational availability as Shannon bits. The apparent incommensurability of theories of information therefore arises not from fundamental ontological differences at the level of the unit, but from post hoc extensions that conflict with the inherited assumptions.

Shannon equivalence commits theories of information to a cluster of assumptions that follow from the assumption of a stabilised difference. Time can only enter as an external parameter, something imposed on information rather than constitutive of it. Presence is given in the stabilised nature of difference, and absence can appear only as the comparative term of a binary distinction, never as an independent informational condition. Where no account is given of how the difference came to be, the originating condition must be treated as prior both ontically and epistemically, since there is no other basis on which to ground the difference. A stabilised difference is also fixed and self-identical, and from this transitivity follows: information can be treated as travelling from source to destination unchanged, because nothing in the account specifies why it should transform. The conditions of persistence, organisation, and higher-order stabilisation therefore remain unexamined.

Tracework rejects these assumptions. Its account of information breaks with the principle of Shannon equivalence. It situates information not as primitive but as product, arising in the stabilisation of variation within a field of continuous flux. This is expressed by the principle of

informational grounding [s. 1.3]. If information exists in kind, then there must be a prior process through which it comes to be; unless that process is specified, the nature of information cannot be known but must be assumed. Physical systems exhibit continuous causal variation without presupposed entities or fixed partitions. The nature and substance of that variation are incidental to the continuous flux. For variation to become informationally operative and consequential, it must be organised such that its dynamics exhibit sufficient persistence under transformation to exist in kind as information. This is the minimal precondition for information to exist. Without introducing additional priors, this organisation must arise through the constraint structures generated by causal interaction itself, under which some patterns of variation acquire stability relative to the surrounding field, while others dissipate without a trace. The conditions under which this occurs are derived in the subsection below [s. 2.4]. Information is the persistent trace of prior causation.

2.3 Definitions: the following terms constitute the minimal technical vocabulary required for the derivation that follows. They are given here as working definitions of distinctions that the subsequent derivation will show to be structurally required. The purpose in defining them here is to clarify the argument before its full derivation has been established and to ensure continuity with the later account of transformation space, viability domains, stabilisation domains, difference, and constraint. No additional ontological or epistemological commitments are presupposed beyond the foundational assumption itself.

Transformation space designates the space of possible configurations and their transformations [cf. s. 4.2]. It carries no commitment to any particular metric, topology, or physical substrate: it names only the space within which configurations can be related by transformation. A viability domain is a bounded region of transformation space within which the conditions required for trace formation can in principle be realised [cf. s. 2.8]. The bounds of a viability domain are set by whatever constraints govern the persistence and discriminability of differences within that region. A stabilisation domain is the narrower operative regime, relative to a given scale or resolution, within which traces and trace-structures remain discriminable and persistent [cf. s. 2.8]. A stabilisation domain is always nested within a viability domain, further constrained by the specific conditions under which differences can be sustained and carried forwards at a given level of description.

Difference designates non-identity under transformation: two configurations differ if they can be distinguished within a stabilisation domain. No metric, observer, or threshold of significance is presupposed beyond what the stabilisation domain itself makes available. This is the weakest possible criterion of difference consistent with the derivation that follows.

A constraint is a restriction on admissible configurations or transformations. Under the productive orientation (below), constraint restricts what successor configurations can follow from a given configuration. Under the inferential orientation (below), it restricts what prior configurations remain admissible as the causal conditions of a present trace. Constraints may be realised physically, formally, statistically, or relationally, but for information to exist they must be operative as restrictions within some viability domain.

Inferential tracework describes the backward temporal orientation, productive tracework the forward temporal orientation, and generative tracework describes their joint operation as a single system. Tracework, used without qualification, refers to the theoretical framework.

2.4 The trace and inferential trace stabilisation: The trace is the minimal stabilised difference capable of preserving both inferential potential and productive continuation. Inferential potential is the relation by which traces sustain reconstruction of their prior causal interactions, whether or not that reconstruction is ever actualised. This relation remains viable only where the trace can continue to participate in generative tracework. To participate in generative tracework is to provide the basis for further trace formation within some relevant viability domain. For example, tree rings preserve the history of their formation independently of observation, but retain inferential consistency only insofar as they can still participate in further trace formation: measurement, comparison, dating, climatic reconstruction, or other continuations in generative tracework. Reconstruction of prior causal interactions should therefore be understood as continuation through transformation [cf. s. 2.5; s. 2.6].

Because traces preserve the structured consequences of prior causal interactions under directional temporal asymmetry, the order in which those interactions stabilise matters. Trace formation is path-dependent. Just as the order of steps in an algorithm determines its outcome, so the order of trace stabilisation determines what is produced: each operation is conditioned by the state its predecessors have left behind. Tracework is not commutative. This extends

recursively through generative tracework [cf. s. 3.2]. While convergence is possible, divergence is the general case.

Difference can function informationally only if it is discriminable [s. 2.3; cf. s. 2.5]. This requires a domain within which difference can be resolved. But every stabilisation domain operates at a given scale, with thresholds of persistence, granularity, and admissible transformation [s. 2.3; cf. s. 4]. Therefore no difference is distinguishable absolutely; difference is distinguishable only relative to the resolution at which the domain can stabilise it. At finer scales than the domain can resolve, variation dissipates into background fluctuation [cf. s. 2.8]; at coarser scales, multiple lower-level differences are compressed into a single trace [cf. s. 3.2]. Differentiation is therefore subject to scale-dependent limits of resolution and discrimination. It is an outcome of stabilisation rather than something given in advance.

Traces are locally stabilised within a material substrate. However, their ontological persistence is realised through inferential consistency: the continuation of the same inferential potential across transformation regardless of material instantiation. Inferential consistency is specified when the trace jointly satisfies three conditions:

1. Differentiation: the difference is distinguishable relative to a given stabilisation domain, within the scale-dependent limits of resolution and discrimination that domain can sustain.

2. Persistence: the difference endures over a temporal interval sufficient for stabilisation, relative to the relevant stabilisation domain.

3. Causal history: the difference retains the structural effects of prior causal interaction in a form that can sustain inferential potential.

These inferential stabilising conditions are not attributes of a trace itself but constraints arising from stabilisation within a given stabilisation domain. They establish the trace as the minimal structure capable of performing inferential work. Without discriminable states there is nothing to reconstruct; without those states persisting there is nothing from which reconstruction could proceed; and without some preserved connection to the prior causes of difference there is no grounding for the claim that reconstruction is of anything at all [cf. s. 2.5]. The trace does not encode causation but bears its imprint in the way a landscape is changed by glaciation, structurally transformed by what has passed. Causal history is therefore a structured sequence of constraint-forming transformations over which a trace's inferential consistency is established.

Each step restricts the space of admissible continuations under productive tracework, and those restrictions become available under inferential tracework as relations of admissibility. However, the sequence itself is not necessarily recoverable from the trace. The intrinsic relation to its causal history means the trace cannot be fully detached from what produced it. It is inherently temporal, carrying the past into the present as the structural consequences that make potential inference about those causal conditions possible.

Inferential consistency does not presuppose an identity relation, but arises from the preservation of structured consequence across transformation. Traces are subject to continual change, persisting only through ongoing transformation under which inferential consistency is preserved. Persistence as a continuing trace is therefore conferred by the partial preservation of its capacity to sustain reconstruction of its prior causal conditions within a stabilisation domain. Although traces are stabilised in a material substrate, inferential consistency and that material substrate can come apart. In some cases, material structure may survive while a specific form of inferential consistency fails, although other inferential potential may persist. For example, a document written in a lost language may lose semantic meaning but retain broader historical significance. Conversely, a trace may retain inferential consistency while its physical substrate changes beyond recognition, for example in the material decay of ancient documents [cf. s. 2.8] or in their re-stabilisation through digitisation [cf. s. 3.2]. Inferential consistency and physical persistence are therefore distinguishable and separable, rather than being parallel descriptions of a single persistence relation. What survives in the trace is an echo of its prior causes.

These conditions for inferential consistency are presence-independent. Differentiation requires that a difference be distinguishable within a stabilisation domain. Persistence requires that the distinguishable state endure across transformation. Causal history requires that the structural consequences of prior interactions be preserved. Wherever an absence satisfies these conditions it must function informationally. However, presence-independence does not imply that traces are substrate independent; a difference, whether a presence or an absence, can only be realised within a material substrate.

Relative to a given scale and viability domain, two traces are equivalent insofar as inferential tracework yields compatible reconstructions of prior causal conditions, or productive tracework yields compatible spaces of admissible continuation [cf. s. 3.2]. Path-dependence means that equivalence across all relevant scales between traces with different histories is a limiting case. For example, two copies of the same document may be equivalent within a retrieval regime,

Preprint

sustaining compatible inferential potential for the purposes of information access, while retaining different production histories that become relevant in archival context. The bibliographic record treats them as equivalent. The collection and item-level descriptions treat them as unique. Both are right at their respective scales of resolution. Equivalence is therefore always domain-relative and scale-dependent [cf. s. 2.6; s. 4.3]. Its full significance becomes clear once trace-structures are derived: the same configuration can function productively as a trace available for further organisation and inferentially as a trace-structure whose consistency is reconstructed from constituent traces [cf. s. 3.2].

2.5 Productive trace stabilisation: traces stabilise under both inferential and productive orientations. These are not independent in principle but arise from the same underlying condition: that traces preserve the structured consequences of prior interactions while remaining capable of participating in subsequent transformation. The productive stabilising mechanisms outlined below therefore parallel the three inferential conditions, expressing the same constraints under the opposite temporal orientation [s. 2.4]. Productive continuation is specified when the trace jointly satisfies three conditions:

1. Variation: the counterpart of differentiation. Variation determines whether causal production yields distinguishable configurations rather than undifferentiated flux.

2. Propagation: the counterpart of persistence. Propagation determines whether causal effects extend forwards sufficiently to sustain further trace formation.

3. Generation: the counterpart of causal history, determining whether prior causal conditions, under the relevant viability constraints, yield configurations capable of grounding further trace formation.

Both productive and inferential stabilising conditions are jointly required. Information can only sustain inferential potential when it admits continuation, because any actualisation of the trace's inferential potential is itself a further trace relation arising out of productive continuation [cf. s. 6.3]. Because inferential potential persists only where productive continuation remains possible, traces exist only through ongoing transformation, and are therefore subject to continual change.

Both requirements rest on the grounding of differentiation in production. This follows from the minimal assumption of tracework. For information to exist it must be distinguishable as information rather than as arbitrary variation. It must be *about* something in at least the minimal sense of a relation that differentiates, and must therefore stand in a relation of distinction to what it is not. To function informationally, differentiation must arise from, and remain constrained by, the conditions producing it. Differentiation that is not constrained is arbitrary, and cannot function informationally. Any candidate for the grounding of constraint other than prior causation, whether reference to independently specified future states, abstract relations, or semantic content, introduces additional commitments that the framework has not derived and is not entitled to assume. Any such alternative conflicts with the minimal account and violates Ockham's razor: differentiation grounded in production is sufficient, and what is sufficient requires no supplement. Therefore, within the minimal framework, the only admissible grounding for such constraint is prior causation as productive formation. This constraint, though derived here, can be translated into more concrete terms. A letter in an archive carries informational value not because of the marks on the page in isolation, but because those marks are the structured consequence of the conditions in which they were produced: physical, circumstantial, and purposive. They are, in this minimal sense, *about* those conditions. Stripped of that history, the letter would be informationally inert. Tracework shows that all information derives from differentiation constrained by the conditions of its production.

Causal history names the structured consequence of production as it is preserved across transformation. It is not the ground of informational differentiation itself, but the form in which grounding in production is retained within the trace. For differentiation to be informational rather than arbitrary, it must arise from and remain constrained by the conditions under which it was produced. Causal history is the condition that secures this: it is what distinguishes a trace from mere variation. Floridi's veridical principle can be read as an attempt to preserve this requirement within a semantic framework (Floridi, 2013), but because it begins from already-stabilised data it must express the requirement as truth rather than derive it as causal history. Tracework does not introduce a truth condition. The grounding function that truth performs in Floridi's semantic account is discharged at the informational level by causal history, which is preserved through persistence rather than stipulated as a semantic requirement. The three inferential conditions are therefore consequences of what it means for differentiation to be informational. Differentiation supplies the bare distinction; persistence sustains it across transformation; causal history structures it by the conditions of its production and retains those conditions within

Preprint

it. All three are aspects of a single requirement: that differentiation be grounded in production under continuation.

The same principle applies to the productive conditions, because the requirement that differentiation be grounded in production constrains differentiation as such, under either orientation. A differentiation that cannot be carried forwards ceases to be structured by its conditions of formation and collapses into undifferentiated variation. The productive conditions are therefore required for informational differentiation to be sustained. Variation is required if productive continuation is to yield distinguishable configurations. Propagation is required if the effects of prior configurations are to extend sufficiently to sustain further trace formation. Generation is required if prior states are to yield new configurations through which structured differentiation is carried forwards across transformation. These conditions describe the same constraint structure encountered under different temporal orientations, although in an individual trace the corresponding inferential and productive capacities can come apart. Any alternative either fails to constitute informational differentiation, fails to sustain it, or introduces commitments the framework cannot justify. Thus all six stabilising conditions are derived from discriminable difference, which itself must be an outcome of the assumption that information exists.

The derivation maps an implicit reading of Bateson's aphorism that information is “a difference that makes a difference” (Bateson, 1972: 460) [s. 2.2]. His first “difference” is the bare distinction, which tracework specifies under both orientations as differentiation and its productive counterpart, variation. His “makes a difference” is the requirement that the distinction be consequential, which tracework specifies as grounding in production carried forwards under continuation: inferentially as causal history sustained through persistence, and productively as propagation and generation. Bateson's formulation captures the requirement of differentiation, but does not specify the conditions under which differentiation is non-arbitrary, persists, or is sustained. Difference in Bateson's minimal sense is therefore necessary but not sufficient. Only when specified as differentiation structured by its conditions of formation, sustained across transformation, and maintained as discriminable does it become sufficient. Bateson's definition of information is therefore a compressed expression of a structure that tracework explicates.

Neither orientation subsumes the other. Traces may lose inferential potential or productive operability in one viability domain while retaining them in another. For example, when a tree is felled, the biological viability domain that sustained the formation of new rings ceases [cf. s. 4.5].

Yet the ring sequence may remain both inferentially viable and productively generative in other viability domains, such as dendrochronological research. The productive and inferential conditions of generative tracework are therefore not only independent in their realisation but may be sustained simultaneously within entirely different viability domains. Continuation may take place in multiple domains at once. What counts as a viable domain for inferential tracework need not coincide with what sustained productive tracework, and the two may diverge indefinitely across time. However, productive and inferential stabilising conditions are inseparable in principle: inferential viability persists only insofar as some productive continuation remains possible somewhere in generative tracework. That possibility does not have to be realised. Where the possibility of productive continuation is exhausted across all relevant viability domains, inferential potential is no longer viable: nothing remains that can provide the basis for further trace formation. At that point the trace fails as an informational structure. Because the two orientations address distinct temporal requirements, breakdown of either undermines generative tracework in different ways. These modes of failure are not contingent errors but structural consequences of tracework, examined in detail later [cf. ss. 5.5–5.6].

2.6 Consequences of temporal asymmetry: the account developed above specifies the conditions under which traces can form, persist, and sustain inferential potential. It also establishes that traces are temporally extended. Production is a process and unfolds across an interval [s. 2.5]; persistence requires that a difference endure long enough to be stabilised [s. 2.4]. A configuration that exists for no interval could neither be produced nor persist, and so could not satisfy the conditions for tracehood. This is the principle of temporal extension [s. 1.3]: time cannot be excluded from the description of a trace, because a trace exists only where the causal history of prior interactions is preserved and carried forwards through productive continuation. That preservation requires transformation to be directionally ordered such that prior states constrain their successors: without this no causal history would be retained, no inferential consistency could be sustained, and no structured consequence could accumulate. The directional constraint required for trace stabilisation is therefore temporal asymmetry. The physical world satisfies these conditions [cf. s. 2.8].

With the trace stabilisation architecture now in place, the principles introduced earlier [s. 1.3] that have not already been derived [s. 2.2–2.5] can now be shown to follow from temporal asymmetry. The first of these is the principle of orientational non-equivalence, previously given

in restricted form [s. 1.3]. Because trace relations are constituted under directional constraint, orientation within tracework is never a neutral perspective on an invariant object. Under the productive orientation, a trace participates in further causal generation, constraining what configurations can follow. Under the inferential orientation, the same trace sustains reconstruction of prior conditions, constraining what configurations could have preceded. These are not two descriptions of the same operation: what is fixed under the productive orientation at the point of trace formation is not recoverable under the inferential orientation as the same operation run backwards. The difference is not merely epistemic but functional. Orientation determines what can be produced, reconstructed, or carried forwards. As is shown later, this applies wherever orientational distinctions are generated by the architecture of tracework: internal versus external continuation, scalar versus relational description, aspectual versus structural specification. Throughout tracework the orientational difference is operationally consequential rather than merely a matter of perspective. The consequences of this are developed across sections 4, 5, and 6, where it underlies the incommensurability of first-person and third-person descriptions of the subjective domain [s. 4; s. 5; s. 6.3; s. 6.4].

Because the causal processes that produce a trace are complete at the moment that trace exists, inference about those processes cannot follow the direction of their production: there is no longer a forward process to follow. It must instead proceed backwards, from the trace towards the conditions that produced it. This is the principle of epistemic asymmetry, the epistemic expression of temporal asymmetry. It is a necessary but not sufficient condition, constraining admissible reconstruction without by itself determining it. Production proceeds forwards through transformation, from prior condition to stabilised consequence; reconstruction proceeds in the opposite epistemic direction, from structured residue to the prior condition whose effects it preserves. Productive priority and epistemic priority therefore do not coincide. The epistemic starting point is the trace itself, and epistemic priority runs from trace to condition rather than from condition to trace. This is the principle of epistemic priority inversion.

If epistemic priority runs from trace to condition, information cannot be transmitted as self-identical content from source to destination, since such transmission would require informational determination to be carried forwards independently of the local conditions under which reconstruction occurs. Instead, inferential potential is constituted wherever sufficient causal history is preserved under transformation, and what appears as transmission is the preservation of sufficient inferential potential across successive transformations to sustain

reconstruction at each stage. Information is therefore not transmitted as self-identical content, but propagates only through transformation. This is the principle of informational intransitivity.

Informational intransitivity does not imply that information is immobile. Spatial movement is itself a transformation. When a trace moves from one location, substrate, or domain to another, its coupling relations and conditions of accessibility change [cf. s. 3]. Informational continuity is preserved only where sufficient inferential consistency survives those changes to sustain inferential potential. For example, an alien artefact arriving on Earth might testify to the existence of its makers while no longer sustaining the reconstructions its inferential potential could yield in the context of its production. What survives the passage is reconstructed in the new domain rather than transmitted intact from its origin. Yet where inferential consistency is preserved, the trace remains objectively the same trace; what changes is the reconstruction available within the new domain. The domain-relativity of reconstruction, described in the information science tradition as domain analysis (Hjørland and Albrechtsen, 1995), is a structural consequence of informational intransitivity. However, tracework locates this relativity in reconstruction rather than in information itself: where inferential consistency is preserved, the trace remains invariant across domains, even though the reconstructions it sustains do not. Spatial movement, copying, and communication are transformations in which the reconstructive reach of its inferential potential may be preserved, degraded, redistributed, or re-stabilised. What continues is not an informational substance moving unchanged through a channel, but inferential potential sustained by the partial preservation of causal history across changing conditions.

Another consequence of informational intransitivity concerns retrocausation. Any hypothetical retrocausation would have to operate below the level at which informational structure stabilises, since informational structure is fixed only at the point of its formation. Below that threshold there is no informational structure to constrain; above it, backward influence cannot operate as informational constraint without already presupposing the trace it would need to alter. Reconstruction therefore operates under local conditions and cannot be bound by the productive orientation that preceded it. Informationally, the future cannot bind the past. This parallels the chronology protection conjecture: just as that conjecture forecloses causal loops by preventing backward influence on prior conditions, informational intransitivity forecloses transmissional loops by preventing the inferential orientation from being bound by the productive orientation (Hawking, 1992).

Preprint

Temporal asymmetry changes the relative size of the space available to a trace in each orientation. In productive tracework, a trace can become the common origin of many independent continuations. When a book is printed, it may become a thousand individual traces, each pointing back towards the same act of composition, and each then free to diverge from all the others, through damage, marginalia, or the accumulated significance of ownership. As trajectories split and re-converge across a shared coupling network they produce an intricate web of relationships. In inferential tracework, by contrast, a trace's candidate origins do not multiply in the same way. A trace may leave its causal history underdetermined among several admissible candidates [cf. s. 2.7], but this underdetermination does not compound in the same manner as forward continuation. The space of what a trace can become is not, in general, matched by an equally rich space of what it could have come from. This is not an absolute asymmetry: as noted, convergence is possible, but is not the general case [s. 2.4]. However, as a tendency it is consistent with the thermodynamic irreversibility discussed further below [cf. s. 2.7].

This asymmetry is visible in how the information professions divide their labour. Bibliographic description, in treating multiple copies as instances of the same work, operationalises the convergent, backward-pointing relation reaching towards a putative idealised original; archival description, in treating each copy's subsequent history, damage, provenance, and ownership as significant in its own right, operationalises the divergent, forward-pointing one specifying individuality [cf. s. 2.4]. Neither is the more authentic description, because each tracks a different part of the same asymmetric structure.

2.7 Causation and complexity: in its most general sense, information is that which testifies to its own causation [s. 2.4; s. 2.5]. What we regard as informational is precisely this relationship between a trace and the inferential potential it sustains: the capacity of a trace to structurally echo the prior causes of its formation. Whatever fails to meet the conditions of trace formation does not qualify as information. That this relationship runs from effect to cause is the asymmetry on which tracework rests. This does not determine what causation is in principle, but how causal structure presents within tracework. Tracework remains agnostic between competing accounts of productive causation, whether interventionist (Woodward, 2003), dispositional (Bird, 2007), process-based (Salmon, 1994), or counterfactual (Lewis, 1973). This marks a significant departure from Dretske's account, which is rooted in counterfactual causation (1981; 1983).

Preprint

Tracework requires only that prior conditions produce subsequent states in ways that leave inferentially legible residues. What metaphysically grounds causation remains open.

This agnosticism is not unprincipled. Causal accounts often implicitly depend on informational conditions. For example, in Woodward's account, causation is defined as that which would be revealed by intervention (Woodward, 2003), but specifying what counts as an intervention presupposes informational access to causal structure. Similarly, Salmon's account relies on the transmission of marks (Salmon, 1994), which presupposes that a mark can be identified, persist through transmission, and be distinguished from background variation. Lewis's counterfactual account of causation depends on comparisons between actual and possible worlds (Lewis, 1973), and therefore requires determinate access to the counterfactual structure through which causal dependence is assessed. In each case, causal specification depends upon conditions of discriminability, persistence, and reconstruction that are informational in the present sense. These accounts do not merely require information as an epistemic supplement to causation; they presuppose assumptions about the nature of information in order to make causal specification possible. This introduces a potential circularity in causal-based theories of information.

Tracework therefore cannot simultaneously derive a theory of causation and a theory of information from a single orientational specification. Fully specifying causation requires following the productive orientation, while fully specifying information requires following the inferential orientation. Because these orientations are not equivalent, any framework that attempts to specify both must remain partially agnostic about one in order to say anything determinate about the other. Committing to an explicit theory of causation would already imply a prior account of the informational conditions under which causal structure can be identified, compared, or reconstructed, thereby presupposing what tracework is intended to derive. Causal agnosticism is a required condition of any account of information rooted in prior causal history [s. 2.6]. For information to exist it must have been productively caused, but no other assumptions about causation can legitimately be made.

The same methodological restriction marks tracework's departure from complexity theory. Tracework has real affinities with complexity theory, particularly around emergence and non-linearity. But these affinities are descriptive rather than foundational. Complexity theory describes dynamic systems through rich interaction among already individuated elements, from which coherent global patterns emerge without being reducible to any element in isolation (Cilliers, 1998; 2005; Manson, 2000; Urry, 2005). In tracework terms, such descriptions already

presuppose discriminable states, persistent relations, and constraints under which interaction can be described. Hofkirchner's emergent account of information, for example, treats information as emerging from complex interaction, but the interacting entities and their relational structure are already given (Hofkirchner, 2013). Complexity theory can therefore describe emergent organisation once differences and relations are in play, but it cannot ground their informational status without presupposing what tracework derives. It is not simply that more is different (Anderson, 1972); difference at higher scales becomes informational only where complex prior causes leave stabilised traces through which it can be reconstructed.

These two examples apply the same condition: a single orientational specification cannot simultaneously ground both the productive and inferential dimensions of information. This shows that causation operates in two analytically distinct but related ways in tracework that can become conflated. As a generative constraint structure, it governs how variations in the field propagate and stabilise into differences. As the selectively preserved residue of prior interactions, it also appears within stabilised traces as the partial causal structure that the backward inferential arrow can reach. These functions are not independent. What makes causal relationships legible in generative tracework is not their forward productive character alone, but the consistency of inferential potential they sustain. Nevertheless the distinction between them can be overlooked. Causation is what inferential tracework consistently finds at the other end of its reconstruction of stable traces.

This distinction reframes the disjunction problem that appears in causal theories of information such as Dretske's (1981; 1983). Standard causal accounts face difficulty differentiating between multiple prior causes that may have produced the same informational effect. This problem arises because causation is treated as prior to information: the ground from which informational content is read (Adams and Aizawa, 2021). Causal theories of content assume that causal events have determinate meanings. In tracework, by contrast, causation is not prior to the trace as a determinate semantic ground, but reconstructed from it. The trace is the structured consequence of whatever causal conditions produced it. Those conditions are available only through its inferential potential, and only subject to stabilisation. Causal history is itself scale-relative: the prior cause reconstructed at the end of inferential tracework is not the total microphysical history of the trace, but the causally relevant profile stabilised at the operative scale. The frog's neural response, for example, does not carry information about flies as such, but about whatever causal conditions produce that response: flies, small moving dark objects, or any other stimulus that generates an equivalent trace at the relevant level of description (cf. Dretske,

1981: 35). Whatever cause is reconstructed at the end of inferential tracework is the cause that matters within that stabilisation domain. Where different causes produce inferentially indistinguishable traces, they are not informationally distinct at that level, whether or not they differ under some other description. Where they produce distinct traces, those traces discriminate between them without requiring a normative criterion external to the causal account. What the frog understands of this is another matter. The disjunction problem is thus reframed as a level-of-description misalignment [cf. s. 3.2; s. 4.3]: the apparent explanatory problem arises from representational assumptions that tracework does not share.

Scale-relativity applies generally in tracework. The cause of a book, reconstructed from the book, is not the microphysical history of ink molecules settling onto paper but the act of writing under which the text was composed; the cause of a painting is the act of painting, not the motion of pigment particles; the cause of a footprint is the passage of a foot, not the individual displacements of soil grains. The trace testifies to its causes stabilised at the scale at which the trace is itself a trace. The finer-grained history is not false, and remains available to differently scaled reconstruction, but it is not what the trace, at its operative scale, is a trace of. Inferential tracework recovers the cause commensurate with the resolution at which the trace stabilises, because that is the scale at which its inferential potential was fixed as a consequence of coarse-graining [cf. s. 4.2]. To reach past it towards microphysical detail is to change the scale of description, and so to ask after a different trace. This is addressed in more detail in the scale architecture [s. 4].

2.8 Limits of trace stabilisation: the limits of trace stabilisation describe the boundedness of tracework. They follow directly from the joint requirement of the inferential and productive stabilising conditions [s. 2.4; s. 2.5]. Both orientations are required because neither subsumes the other [s. 2.5], and the failure of either is sufficient to terminate trace persistence. Both orientations are also materially instantiated, because traces are themselves materially instantiated, and their formation and continuation are therefore subject to physical constraint. Any materially instantiated process under physical constraint is bounded in two directions: there is a lower bound below which the conditions for that process cannot be satisfied at all, and an upper bound beyond which physical constraint prevents further realisation. Two orientations, each bounded in two directions, yield four intrinsic limits. These are the structural expressions of

a single requirement: the joint necessity of inferential and productive conditions under physical constraint.

At the inferential lower bound, difference fails to achieve sufficient persistence, differentiation, and causal history to constitute a trace. What exists dissipates without leaving stabilised structure. At the inferential upper bound, the density of distinguishable states within a bounded region of physical space reaches its physical limit: further differentiation is no longer possible within the constraints governing that region. At the productive lower bound, the rate of productive continuation falls below that required to sustain further generative tracework: the trace ceases to participate in forward continuation and, where no alternative viability domain sustains it, fails as an informational structure. At the productive upper bound, the finite speed at which physical constraints can propagate through a substrate limits the scale at which coherent stabilisation can be maintained: the system exceeds the rate at which its own constraints can bind it.

These limits are not arbitrary, but derived from the stabilisation conditions. However, they are also corroborated by the physical constraints governing material reality. Bekenstein's (1981) upper bound on the information content of physical systems establishes that informational structure is constrained by physical law. Landauer's principle situates information processing within thermodynamic cost, tying informational change to physical constraint (Landauer, 1991). He writes that:

> Information is inevitably inscribed in a physical medium. It is not an abstract entity. It can be denoted by a hole in a punched card, by the orientation of a nuclear spin, or by the pulses transmitted by a neuron (Landauer, 1991: 29).

The finite speed of propagation ensures that informational processes are temporally extended rather than instantaneous, and thermodynamic processes exhibit effective irreversibility at the scales at which stable differences can form and persist (Wiener, 1961). These considerations do not generate the limits: they demonstrate that the universe satisfies the structural conditions the limits express.

The limits of trace stabilisation constrain an interior viability domain: the region within transformation space where the lower and upper bounds of both inferential and productive orientations are jointly satisfiable, and where generative tracework is therefore possible. This region is not uniform. Generative capacity varies across it, declining as the bounds are

approached. Its existence as a bounded finite space is presupposed by the scaling architecture [cf. s. 4] and defines the space within which the transformation architecture operates [cf. s. 5].

2.9 Why structure appears: the architecture derived above follows from the single assumption that information exists. The fact that this minimal assumption generates structure is a consequence of selection under temporal asymmetry and orientational non-equivalence. Stabilisation is inherently selective: only variations that satisfy the inferential and productive stabilising conditions can persist as traces. The population of persisting traces is therefore, by definition, a structured subset selected by constraint. Apparent organisation is the surviving residue of that constraint operating over time. That which does not satisfy the conditions of informational persistence is absent from the informational landscape. It exists as difference but not as information.

The same selective logic recurs across levels of organisation. Trace-structures, viability domains, trajectories, subjective domains, and the intersubjective domain are not independent kinds of thing to which separate explanatory principles must be applied. They are successive reorganisations of what survives under asymmetrical constraint. At each level, only those configurations that satisfy the relevant conditions of persistence become available to further organisation at that level. The structure that appears at each level is therefore the consequence of the same selective pressure operating recursively over what has already survived prior rounds of stabilisation. This does not mean that the same structure performs the same operation at every level. Orientational non-equivalence prevents recurrence from collapsing into repetition. The same constraint structure functions differently according to the orientation under which it is encountered: productively as a condition on what can follow, inferentially as a condition on what can be reconstructed, scalarly as a condition on what can be resolved, relationally as a condition on what can be coupled, and, later, subjectively as a condition on what can be inhabited and carried forwards under commitment. The formal analogues that recur across tracework are therefore consequences of the same architecture encountered under different orientations.

But selection and orientational non-equivalence do not fully explain informational complexity. Selection explains why structure appears and persists; orientational non-equivalence explains why recurrent structure acquires different operational roles; but neither explains why structure becomes recursively elaborated and richly significant. Tree rings are one thing; cave paintings, decorative murals, data, and scientific theories belong to another order of organisational

complexity. To explain how difference becomes significant at a human level, self-regulation is required [cf. s. 6]. Where tracework begins to operate over its own products, selecting among admissible continuations, binding itself to some rather than others, and reorganising the conditions of its own persistence, it becomes capable of deeper recursive elaboration. What selective persistence produces as surviving structure, self-regulation recursively deepens into organised complexity. This fuller extension of the architecture is developed later [cf. s. 6], where the account properly belongs; however, it is introduced here to signal that the structural richness of tracework given in the following sections is not derived from trace stabilisation alone. Reflexive self-regulation is already latent within the architecture.

Preprint

3. Tracework recursivity

3.1 Introduction: traces persist through transformation. Like Heraclitus' river, the trace is never the same twice. However, because inferential and productive stabilisation express different aspects of epistemic asymmetry, looking upstream and downstream are not equivalent operations in tracework. Persistence therefore cannot consist in the absence of change but only in the maintenance of structure compatible with constraint across change. A trace persists as a trajectory through transformation: an equilibrium between production and decay. This section develops the relational structure of generative tracework, through which stabilised traces become organised into trace-structures. That organisation is governed by coupling constraints, which determine how traces both combine and retain inferential consistency across continued transformation.

3.2 Trace-structure ontology: traces retain the structural effects of the causal interactions through which they were formed [s. 2.4]. Where traces arise through mutual causal interaction, the inferential potential of each is partially constituted by the inferential potential the others sustain, and their continuation is jointly constrained. For example, a chess move alters what subsequent moves are available to both players. Configurations of this kind are called trace-structures: stabilised organisations in which inferential consistency is preserved relationally rather than locally, and in which finer-grained causal history is partially compressed while coarser-grained relational organisation remains operative [cf. s. 4.2]. Trace-structures are outcomes of the stabilisation conditions [s. 2.4; s. 2.5], made possible by the temporal asymmetry under which those conditions operate.

The principle of trace equivalence states that the same configuration may be described as a trace or as a trace-structure depending on context. Productively, the configuration functions as a trace, a stabilised unit conditioning further continuation and capable of entering higher-order organisation. Inferentially, it functions as a trace-structure, an organisation whose consistency is reconstructed from the constituent traces from which it was formed. The two descriptions are equivalent in the sense that they concern the same configuration, not in the sense that the orientations under which it is accessed are themselves equivalent; orientational non-equivalence prevents the productive and inferential readings from collapsing into one another. This principle does not collapse orientation into scale. The two orientations remain analytically distinct from scale, but in the trace/trace-structure relation they traverse scale in opposite

directions: production treating the configuration as a unit available for further organisation, inference as an organised outcome whose conditions lie in prior constituent relations.

Because trace-structures are nothing other than traces under relational organisation, they inherit the six stabilising conditions already derived at trace level [s. 2.4; s. 2.5]. But inherited conditions do not reappear unchanged under the principle of orientational non-equivalence. Read forwards, trace-structures stabilise under the same six conditions as any trace, but read backwards, those same conditions are expressed not as attributes of an individual trace but as characteristics of a relational configuration. Under relational organisation, the six trace-level conditions reorganise into a two-plus-one pattern. Two families appear as paired requirements because relational persistence must be maintained both inferentially and productively: the structure must remain recoverable as the same organisation, and it must remain capable of further continuation. The third family is asymmetric because the trace-structure cannot be treated as already given. It must first have been constituted through prior generative tracework before it can itself become generative.

In trace-structures, inferential potential is distributed across configurations rather than lodged in any single constituent trace. Therefore the structure must be capable of maintaining consistency under substitution of constituent traces so long as the governing relational constraints remain intact. This is relational stability, the first stabilising condition. Productively, the same requirement appears as resilience: the capacity of the configuration to continue participating in generative tracework under perturbation, local loss, or partial failure, provided sufficient redundancy remains within the coupling network that sustains it. What persists is the distributed capacity of the organisation to continue sustaining inferential potential across multiple relations.

Relational stability and resilience are not sufficient, however. A configuration may remain consistent under substitution while losing the pattern that makes it recognisable as the same organisation across transformation. What must endure inferentially is therefore the governing pattern of constituent traces. This is pattern persistence. Productively, the same requirement appears as extension: the capacity of the structure to carry its relational order forwards through new trace formation rather than merely preserving what it already has. Pattern persistence and extension express the same structural requirement under opposite temporal orientations: that a relational organisation remain both retrospectively recoverable and prospectively productive.

These two paired families describe the conditions under which a relational configuration persists once constituted. But trace-structures cannot be presupposed as already constituted, because

that would reproduce at a higher level the same violation of informational grounding that tracework rejects at the lower: it would assume the product while leaving the process of its constitution unexplained. They must therefore be understood as recursively stabilised outcomes of prior tracework. This is compositionality: trace-structures are constituted through recursive stabilisation and as a result inherit path-dependence and non-commutativity at the level of organisation. Productively, this asymmetry appears as generativity: the capacity of a trace-structure to participate in the production of new configurations through recursive combination under coupling, channelling and constraining how further traces are generated, combined, and transformed. Unlike relational stability and pattern persistence, compositionality and generativity are asymmetric: they concern the fact that trace-structures are themselves products of generative tracework and in turn become conditions of further production.

Stabilising domains are therefore a kind of trace structure. The principle of reflexive articulation states that a trace can only articulate with other forms of tracework. Stabilising domains are therefore constituted within tracwework itself. A trace-structure functioning productively as the organised context of further trace formation constitutes a stabilisation domain at the relevant scale. It is not an external container or independently given field. It is a trace-structure considered in its constraint-forming operation: prior tracework whose relational organisation determines which subsequent differences remain discriminable, which traces can couple, and which transformations remain admissible. Trace-structure names what the configuration is; stabilisation domain names what it does. Because traces articulate only with traces, the constraints constituting a stabilisation domain must themselves be informationally stabilised.

These dynamics can be illustrated by documents (cf. Buckland, 1991; 1997; 1998). A document is not given but constituted through the recursive stabilisation of linguistic, typographic, generic, documentary, and institutional traces, which expresses the document's compositionality. Because what persists in a document is its relational organisation rather than any single instantiation, it can maintain inferential consistency across changes of substrate, format, or encoding provided the governing structure is preserved. This is pattern persistence: the governing pattern of relations endures across the transformation even as the document is re-instantiated. It may also maintain that organisation as its constituent traces are substituted, preserving its inferential consistency across variation in wording, presentation, and argument so long as the constraints defining that role remain operative. This is the document's relational stability. Productively, a document is generative insofar as it participates in the formation of new documents or other traces. It is resilient insofar as it continues to function despite loss of

individual copies or institutional disruption, provided sufficient redundancy remains within the network sustaining it. It exhibits extension insofar as its causal history is carried forwards through citation, translation, and digitisation. A document is therefore a trace-structure whose persistence and productivity are distributed across the relations that continue to stabilise it.

3.3 Trace-structure persistence: the existence of trace-structures raises the question of how relational configurations remain stabilised when their constituent traces are subject to continuous transformation. A relational configuration can persist in only two ways: through continuity of form, or through continuity of relational positioning. For a trace-structure, these correspond to compatibility with the generative conditions from which it arose, such that backward reconstruction of its causal history remains possible, and compatibility with the wider network of traces and trace-structures within which it is embedded, such that inferential potential is sustained by relational organisation even where finer-grained causal recoverability has weakened. No third option is available once stabilisation has become relationally distributed: any trace-structure that persists must do so by remaining compatible with its generative history, its sustaining network, or both. These two axes of stabilisation are called correspondence and coherence.

Correspondence stabilisation follows from the causal history condition derived at trace level [s. 2.4]. A trace-structure is a stabilised consequence of prior generative tracework; therefore, the organisation of its relations must preserve, in some degree, the structural effects of the causal interactions through which it was formed. Without this it would lack inferential grounding and could no longer function informationally. Correspondence is the condition under which those inherited constraints remain recoverable across subsequent transformation. A trace-structure stabilises through correspondence when it remains embedded within a relevant viability domain such that its causal history continues to constrain what it can become, and such that backward inference can still reach those prior conditions through the structure itself. This implies neither representation nor a grounding in external reality. Correspondence is continued constraint-compatibility across transformation. A deposited sedimentary sequence, for example, persists as a trace-structure not because it stands for the conditions of its formation, but because its organisation remains structurally compatible with those conditions across later transformation.

Once more than one trace or trace-structure exists within a shared stabilisation domain, each conditions what the others can sustain either inferentially or productively. A trace-structure must

therefore persist not only by remaining compatible with the conditions of its formation, but by remaining compatible with the traces and trace-structures within which it is embedded. Coherence stabilisation follows from the relational character of persistence at this level. Where persistence is distributed across a configuration, the configuration remains viable only insofar as its constituent relations continue to fit within the same operative constraint structure. This implies neither abstract consistency nor mere formal non-contradiction. It is the continued capacity of a trace-structure to occupy a compatible position within an actual trace network. A document, for example, persists not only through correspondence with the conditions of its production, but through its embedding in the document network: through citation, commentary, reinterpretation, institutional recognition, and documentary adjacency.

Most actual trace-structures stabilise through both correspondence and coherence in combination expressed to different degrees in different configurations. High correspondence does not preclude high coherence, and strong coherence may locally compensate for weakened correspondence, though never indefinitely, because coherence cannot float entirely free of generative grounding. Somewhere within the network, some compatibility with prior causal conditions must remain. For example, a well-preserved medieval manuscript embedded within an active scholarly and institutional tradition remains structurally compatible with the conditions of its production while also remaining relationally compatible with the network of texts, commentaries, catalogues, and institutions that sustain its significance. Pure correspondence and pure coherence are therefore limiting cases. In practice, trace-structures stabilise through a combination of both, the relative dominance of each varying with scale, substrate, and domain.

The fact that trace-structures persist through their embedding in a network continuously traversed by prior tracework means that network cannot remain unaffected. Prior interaction between traces is itself part of their causal history, and under coherence that history is distributed across the relations through which continuation is sustained. This produces network residue: the accumulated deformation of connectivity resulting from prior tracework within the network. Where correspondence stabilisation leaves history in the causal organisation of individual traces, coherence stabilisation leaves history in the relational infrastructure that embeds them: concentrating constraint along frequently traversed pathways, generating asymmetries of accessibility, and producing scar-like deformations in coupling structure that persist independently of the traces that first produced them. Network residue is therefore an aggregate effect, recoverable as causal pressure even where the path identity of individual traces is not. Its effects remain active in the distribution of future accessibility even where the originating traces

have decayed, conditioning both correspondence and coherence going forwards by shaping which pathways remain open and which configurations are most likely to persist.

Network residue gives tracework an account of memory. As a property of tracework memory is the retained deformation of a network by prior tracework. A system 'remembers' where previous operations have altered the distribution of accessibility, constraint, expectation, or pathway-strength within the network through which later operations must pass. Recall is accordingly retrodictive movement through that residue [cf. I.5.4] rather than the neutral retrieval of a stored copy

3.4 Trace coupling constraints: once trace-structures are given, stabilisation is no longer determined solely by the conditions governing each constituent trace. Each trace conditions what the others can preserve inferentially and carry forwards productively. This co-dependency introduces a new level of structural constraint that governs how traces combine. Coupling describes the structured relation through which traces condition one another's continuation, so that the persistence of each becomes partly dependent on the constraints sustained by the others. The forms coupling can take follow from the stabilising conditions already derived [s. 2.4; s. 2.5]. Traces exist under temporal asymmetry, reaching backwards to their prior causal conditions and forwards towards further trace formation. Any coupling constraint must therefore be articulated under one or both of these orientations.

There are four possible relational dimensions along which traces can couple. First, traces occupy positions in temporal sequences, so coupling must govern what may have preceded and what can follow. Second, traces occupy positions within spaces of admissible alternatives, so coupling must govern substitutive compatibility across those spaces. Third, traces are embedded within relational networks whose internal organisation conditions their persistence. Therefore, coupling must govern relational compatibility and the distribution of support within those networks. Fourth, traces exist across scales and levels of description, so coupling must govern cross-scale admissibility under transformation. These four relational dimensions are exhaustive, following from temporal extension, differentiation, relational organisation, and scale within stabilisation domains. Each yields a paired constraint under the inferential and productive orientations, producing four paired classes. However, a fifth class sits outside the logic of relationality. Traces do not couple in the abstract, but co-stabilise within material substrates [s. 2.4]. Their material co-presence constrains what they can jointly sustain. Physical coupling

Preprint

operates through the causal consequences of material coexistence: proximity, contact, and substrate compatibility. It is the material ground that makes the four paired classes possible.

The coupling architecture therefore has a four-plus-one structure: four paired classes exhausting the relational dimensions along which traces condition one another, and one substrate-mediated class governing the material conditions under which coupling becomes possible. Trace-structural coupling requires physical coupling together with at least one operative relational class. Physical coupling alone is necessary but not sufficient: it provides the material ground under which relational coupling becomes possible, but without at least one relational class governing how traces condition one another's continuation, no distributed inferential consistency arises and no trace-structure is constituted.

The first pair of coupling constraints governs temporal ordering: sequence and succession. Under temporal asymmetry, inferential tracework can only reach what has been preserved of prior causal history, while productive tracework can only continue from what the present state sustains. For example, a river is carved into the landscape over centuries, and that prior shape conditions its future development. Sequence is the inferential constraint, governing what may have preceded a trace given the partial and compressive preservation of causal history. Succession is its productive counterpart, governing what can follow under productive continuation. Multiple distinct sequences may be compatible with the same inferential profile, particularly under coarse-graining or where intervening traces have been lost, just as multiple distinct continuations may be compatible with the same present state under succession. The asymmetry is not that one direction admits multiplicity and the other does not [s. 2.6], but that admissible continuations can themselves become the origin of further, independently branching continuations, while admissible prior sequences remain candidates for the same fixed event rather than seeds of a comparably growing space.

A second pair of constraints governs the substitutive dimension: paradigmatic and combinatorial coupling. Paradigmatic coupling is the inferential constraint, governing which traces are admissible substitutes at a given position within the surrounding structure. Combinatorial coupling is its productive counterpart, governing whether traces may be jointly stabilised so as to participate in generative tracework. For example, in Ranganathan's faceted classification each facet position admits a range of admissible terms, while the rules of facet combination govern which may be jointly sustained across positions (Ranganathan, 1933). Classification is paradigmatic tracework: the stabilisation of admissible relations among traces within a

coherence domain. Faceted classification makes this explicit. Each facet position admits a range of substitutable terms, while the rules governing facet combination determine which traces can be jointly stabilised across positions. Individually stable traces may nonetheless fail to combine if their joint stabilisation violates domain-specific constraints such as material incompatibility, organisational exclusion, or resolution mismatch.

Traces are also embedded within relational networks whose internal organisation conditions their persistence. This is governed by a third pair of coupling constraints: coherence and redundancy coupling. Where the first two pairs govern traces in relation to positions, coherence and redundancy govern traces in relation to the networks within which they are embedded. Coherence coupling is the inferential constraint, arising when traces are embedded within a shared relational network whose internal compatibility conditions constrain the admissible stabilisations of each constituent trace. Redundancy coupling is its productive counterpart, governing whether inferential potential and further trace formation are distributed across multiple relations such that no single trace or pathway alone sustains productive participation. Like spiders' webs, coherence and redundancy coupling jointly sustain relational organisation across both temporal directions: sticky but tensile, warping under tension, preserving both constraint on backward reconstruction and robustness of forward continuation. Under repeated tracework, such warping accumulates as the network residue derived previously [s. 3.3].

The three pairs given so far govern coupling within a stabilisation domain, but traces also exist across scales and levels of description [cf. s. 4]. A fourth pair governs the conditions under which coupling persists across those changes: resolution and domain coupling. Resolution coupling is the inferential constraint, governing how stabilised differences at one level constrain, and are constrained by, their finer- or coarser-grained decompositions. Domain coupling is its productive counterpart, governing whether generative tracework can continue when scale, context, or descriptive framework shifts. Like the infinite zoom of Powers of *Ten* (Eames and Eames, 1977), resolution and domain coupling jointly delimit the cross-scale horizon of generative tracework, specifying how backward reconstruction degrades under compression and how forward trace formation depends on the persistence of the viability and stabilisation conditions within which it remains admissible.

Physical coupling differs in kind from the paired classes. It governs the mutual constraint of traces that co-stabilise in spatial and causal proximity within the same material substrate or shared causal medium. In its productive aspect, it determines how the stabilisation of one trace

materially affects the stabilisation of others under forward continuation. In its inferential aspect, it constrains what can be reconstructed about prior material conditions from the joint configuration of co-present traces. Like mycelium spreading through dead wood, physical coupling operates through contact and embedding. It differs from sequence coupling because it does not order traces in time, and from coherence coupling because it does not operate through relational compatibility as such, but through substrate-mediated, proximity-dependent effects.

The four-plus-one coupling classes derived above govern the admissibility and continuation of a single trace or trace-structure taken on its own terms. But because trace-structures are subject to trace equivalence [s. 3.2] and are accessed productively as a trace, a further relation follows without introducing new coupling classes. Where the productive continuation licensed by one trace-structure reaches, via succession or combinatorial coupling, a configuration that falls within the inferential admissibility of a second, distinct trace-structure, the two are coupled across orientation rather than within it. This is is a consequence of applying the coupling architecture already derived to two structures at once rather than one, drawing on the same joint necessity of inferential and productive conditions established for the individual trace [s. 2.5], creating a cross-orientation bridge. Its existence is exhaustively dichotomous in the same sense as subsumption and complementarity are for admissibility [s. 2.4]: either the reaching configuration is already productively coherent within the structure it departs from, or the bridge is carried by some proper substructure of it, reduced under the same coupling constraints. A cross-orientation bridge does not weaken the exhaustiveness of the four-plus-one coupling architecture but is implied when two coupled structures are considered jointly.

3.5 Limits of trace coupling: coupling arises when the stabilisation of one trace conditions what another can sustain inferentially or productively, which is itself the stabilisation of a relation. This process depends on the same stabilising conditions [s. 2.4; s. 2.5], as required by trace equivalence [s. 3.2]. A coupling relation must be discriminable within a stabilisation domain, must persist across transformation, and must retain the structural consequences of the causal interactions through which it was established. Where those conditions fail at the relational level, coupling fails. The limits of coupling are therefore the limits of trace stabilisation [s. 2.8], reappearing at the level of relations.

At the inferential lower bound, relations fail to achieve sufficient differentiation and persistence to sustain coupling at all: there is nothing stable enough to enter into a conditioning relation with

anything else. At the inferential upper bound, the relational space becomes saturated: the density of coupling relations within a bounded stabilisation domain reaches its limit and no further relational differentiation is possible. At the productive lower bound, coupling relations can no longer sustain forward continuation: relations persist but cease to participate in further generative tracework within the operative stabilisation domain. At the productive upper bound, the finite speed at which constraint can propagate across a coupling network is exceeded: the network can no longer bind itself as a coherent relational domain.

To these four, the coupling architecture adds a fifth: fragmentation. The four intrinsic limits are limits on tracehood as such. But coupling introduces configurations whose inferential consistency is distributed across relations among multiple traces rather than lodged in any one of them. That distributed character introduces a new mode of failure that has no counterpart at the trace level. A trace-structure can fail as a configuration while its constituent traces remain individually viable because the relational binding has itself failed. Where coherence coupling can no longer bind a trace-structure as a single coupled domain, the structure decomposes into disconnected components that cannot be reintegrated under the operative coupling constraints. Individual traces may persist and retain their own inferential potential, but the distributed inferential consistency that depended on their relational binding is lost. Fragmentation is therefore a structural consequence of the distributed character of relational organisation. It completes the four-plus-one structure of the coupling architecture, mirroring the coupling classes [s. 3.4]. The full dynamic account of all five limits under transformation is developed later [cf. s. 5.5; s. 5.6].

3.6 Recursive plasticity: the trace is defined by its capacity to sustain inferential potential through transformation. Because reconstruction is itself a form of productive continuation, this capacity depends on the continuing possibility of further trace formation somewhere in generative tracework. Without such continuation, a trace cannot persist as an informational structure. That persistence comes at a cost. Inferential potential is progressively lost through two irreversible processes. The first is the asymmetry between causation and inference: stabilisation preserves only what remains sufficient for inferential consistency at the operative scale, not the full causal history of production. The second is coarse-graining: multiple lower-level differences are compressed into fewer coarser-grained structures, and differences that cannot be sustained

at the coarser resolution are lost [cf. s. 4.3]. Both forms of loss accumulate with scale. Traces persist by trading fine-grained causal history for inferential consistency at a coarser scale.

This produces structural drift in generative tracework towards coarser levels of organisation. The drift is not goal-directed, but follows from the conditions of stabilisation themselves. Gould's analogy of the drunkard's walk explains this process (Gould, 1996): a drunkard walking randomly from a wall will appear to move away from it not because they are trying to, but because the wall prevents movement in one direction. Similarly, traces that can sustain inferential potential through productive continuation under compression persist and couple into more organised configurations, while those that cannot either cease to participate in generative tracework or rigidify, depending on which stabilising conditions fail first. In addition, because transformation includes trajectories splitting into multiple forward continuations [s. 2.6], traces will tend to proliferate at coarser grains. More stable, coarser-grained traces will therefore tend to predominate not because stability is selected for, but because instability exits the viability domain, whereas stable traces persist and proliferate. The drift is not monotonic. Local fragmentation remains common, and scale transition occurs only where the density of inferential potential crosses the coherence threshold required for re-stabilisation [cf. s. 4.2]. What accumulates is not complete causal history but progressively compressed structure that continues to support partial, level-relative reconstruction. Retrodiction remains available as a corrective orientation within this drift, and is discussed fully later [cf. s. 5.4].

Recursive plasticity is the name given to the bounded capacity through which generative tracework remains viable within this drift. It is the capacity of a trace or trace-structure to undergo transformation while preserving inferential consistency, occupying the interval between dissolution at the lower bound and rigidity at the upper bound. It does not introduce additional constraints, but describes the extent to which the stabilisation and coupling conditions already derived remain jointly satisfiable under ongoing transformation. Recursive plasticity is not flexibility in the general sense. Unconstrained variability is a mode of stabilisation failure, while excessive constraint collapses into saturation. In practical terms, recursive plasticity is what allows a document to be reformatted, translated, indexed, or reinterpreted without losing its capacity to participate in inference.

The forms of recursive plasticity are derived from the architecture already established. Because inferential consistency is preserved through two stabilisation operations, correspondence and coherence, a trace has plasticity along two corresponding axes. It has correspondence plasticity

insofar as it can sustain inferential consistency across changes of physical instantiation while remaining compatible with its causal conditions of formation. Archival provenance operationalises this condition: the principle that records must remain recoverable to the conditions of their formation across changes of custody and arrangement is the institutional enforcement of correspondence plasticity. It has coherence plasticity insofar as it can reorganise its relational structure without losing inferential consistency within the network in which it is embedded. These two forms follow directly from the two stabilisation operations [s. 3.3].

However, a trace may retain both correspondence plasticity and coherence plasticity while remaining vulnerable to local failure within its coupling network: a single pathway failure that destroys productive participation even where compatibility is preserved along both stabilisation axes. The distribution of inferential potential across coupling relations is therefore an independent condition of plasticity that does not reduce to either correspondence or coherence. This is redundancy plasticity. Splitting [s. 2.6] is one of the mechanisms by which such distribution arises: where productive continuation generates multiple independent pathways, redundancy increases; where it does not, a trace-structure remains exposed to single pathway failure regardless of how well correspondence and coherence are otherwise sustained. Redundancy follows from the coupling architecture [s. 3.4] rather than from the stabilisation axes alone. Recursive plasticity therefore has three derived forms: correspondence plasticity, coherence plasticity, and redundancy plasticity. The first preserves causal grounding across transformation, the second preserves relational organisation, and the third preserves distributed support. No fourth form of recursive plasticity is possible without introducing either a further stabilisation axis or an additional independent coupling condition.

The forms of recursive plasticity define the space of admissible transformation for any given trace within its stabilisation domain. Without sufficient coherence plasticity, trace-structures become rigid: coherent but unable to reorganise under changing conditions. Without sufficient correspondence plasticity, they become substrate-dependent: unable to sustain inferential consistency across changes of instantiation. Without sufficient redundancy plasticity, they become brittle: unable to preserve inferential potential across multiple pathways or absorb perturbation without collapse of productive participation. At the lower bound, plasticity collapses into dissolution; at the upper bound, into rigidity. The full account of these limits is developed later [cf. s. 5.5].

Preprint

The distinction between information and meaning made later in this paper turns on this capacity. A perfectly rigid system may sustain inferential potential under fixed correspondence conditions, but it cannot sustain meaning, since meaning requires the capacity to revise, reorganise, and redeploy inferential potential across contexts [cf. s. 6].

Preprint

4. Tracework scaling

4.1 Introduction: trace-structures are not merely patterns of relational organisation. Their constitution is also a scaling event, in which tracework moves across determinate levels of resolution. In tracework, scale does not introduce new ontological or structural entities, but redescribes existing elements under altered conditions of stabilisation and viability. Scaling is therefore an active constituent of tracework, intrinsic to its generativity and decisive in shaping what can be preserved, reconstructed, and productively continued through transformation. Within tracework, everything scales: causes, traces, stabilisation constraints, coupling constraints, viability domains, transformations, and trajectories. This section derives the scaling architecture of tracework, including the processes of coarse-graining and re-stabilisation through which trace-structures are organised across viability domains. That architecture is grounded in the principle of trace equivalence [s. 3.2]: across levels and under both temporal orientations, the same stabilising conditions obtain uniformly, differing only in the means by which they are satisfied.

4.2 Scaling as re-stabilisation: coarse-graining is the process by which finer-grained detail is compressed into coarser structures as the scale of description increases. It can therefore be understood as the stabilisation of a trace-structure as a trace at another scale, in which finer-grained relational organisation becomes productively available as a coarser-grained aggregate trace in further continuation [s. 3.2]. This re-stabilisation shifts the invariance sustained across scale from the correspondence-dominant recoverability of formation conditions towards coherence-dominant stabilisation [cf. s. 3.6].

Coarse-graining compresses available inferential potential irreversibly at the new scale at which the trace is taken up. Because discriminability is scale-dependent [s. 2.4], and because discriminable differences at finer scales are rendered indistinguishable at coarser scales, anything that falls below the operative resolution of the coarser stabilisation domain cannot be recovered from within it. Coarse-graining therefore loses fine-grained correspondence grounding: the detailed causal history through which constituent traces sustain inferential reach towards the specific conditions of their formation. That structure remains present in finer-grained descriptions or in the physical substrate, but it no longer functions as the basis of stabilisation at the coarser scale. What survives is inferential consistency sustained by coherence: the relational

Preprint

organisation of the configuration as a whole, rather than the causal history of its constituents. Coarse-graining is therefore a structural shift in the dominant mode of stabilisation.

Scale transition follows from the conditions of generative tracework. More densely differentiated configurations support richer inferential potential and more productive pathways, expressing greater generativity at the coarser scale [s. 3.2], but they also multiply the points at which correspondence, coherence, and productive continuity may fail. Ongoing transformation intensifies this tension by undermining the integrity of constituent traces and increasing the divergence of their inferential potential. Under these conditions, only those trace-structures whose competing inferential potential admits compression into a dominant trace, and whose productive pathways admit consolidation into a coherent continuation, remain viable. Scale transition is therefore the structural consequence of selection under transformational tension: what persists at the coarser scale is what survives compression. This process is non-commutative [s. 2.4]: the order in which constituent traces are stabilised and compressed determines what configuration survives, and different sequences of stabilisation produce structurally divergent outcomes even from the same constituent traces.

The role of coarse-graining in tracework can be illustrated by human memory. Memory is a re-stabilised configuration in which gestures, sensations, and fine-grained situational details are compressed into a form that sustains inferential potential unavailable in the immediacy of the lived moment from which it arose. A dominant configuration of inferential consistency survives through selective forgetting. The same holds for trace-structures at coarser scales. They are integrations of what preceded them: stabilised consequences of how their constituent traces came to be coupled in that particular configuration. New inferential and productive orientations are realised by re-stabilised relational organisation at the cost of compressed causal history below that scale. Traces persist by trading fine-grained causal history for coherence-organised inferential consistency, preserving the possibility of productive continuation by reconstituting tracework at a coarser level. The recursive constitution of coarser structures from finer ones is compositionality operating across levels of organisation [s. 3.2]. The consequences of this shift for scale-dependent viability and for the structural dimensions that become visible within organised tracework are developed later [cf. s. 4.3; s. 4.4].

4.3 Scale-dependent viability: the drift towards coherence stabilisation [s. 4.2] has consequences for how traces and trace-structures persist across scales. Because re-

stabilisation attenuates correspondence grounding while consolidating coherence organisation, the conditions of viability are not uniform across scales. At finer-grained scales, stabilisation is highly sensitive to micro-level perturbations. Small disruptions to the causal conditions that produced a trace can directly undermine its inferential consistency. Correspondence must therefore be tightly maintained. At coarser-grained scales, by contrast, coherence predominates. Micro-level variation falls below the operative resolution of the stabilisation domain, and relational stability is maintained across a broader range of internal states [s. 3.2]: inferential consistency is sustained through relational compatibility within the coupling network rather than merely through correspondence with prior causal conditions. Coarse-graining thus functions as a ratchet. Each step towards coarser organisation increases tolerance to micro-level perturbations by compressing differences below operative resolution.

This ratchet operates within limits. Coarse-graining compresses differentiation in order to sustain inferential consistency at a coarser scale, but differentiation is itself the minimal condition of trace formation [s. 2.3]. Compression that proceeds far enough to eliminate the differentiation required for any distinguishable structure to remain does not produce a viable coarser trace; it approaches dissipation by over-compression rather than by insufficient formation [s. 2.8]. The drift towards coarser stabilisation is therefore self-limiting: it cannot consume the differentiation on which its own further continuation depends. In addition, as fine-grained correspondence grounding weakens under coarse-graining, the trace-structure becomes progressively more dependent on its coupling network for inferential viability. Robustness at the coarser scale is therefore not intrinsic but a function of redundancy plasticity [s. 3.6]: the distribution of inferential potential across multiple coupling relations such that no single pathway failure is sufficient to undermine inferential consistency at the operative scale. At finer scales, stabilisation is correspondence-dominant with respect to detailed causal history. As scale coarsens, that fine-grained correspondence is compressed, and inferential consistency depends increasingly on coherence among the relations that survive. This does not mean that coarser structures cannot be correspondence-dominant. Rather, correspondence is re-anchored at the coarser scale, where the trace testifies not to the full microhistorical sequence of its formation but to a compressed causal profile stabilised within that domain.

This has a direct consequence for trace equivalence [s. 3.2]. Two traces are equivalent at a given scale insofar as inferential tracework over either yields compatible reconstructions of prior causal conditions at that scale. Because coarse-graining compresses the finer-grained differences that would distinguish them, equivalence is more readily satisfied as scale coarsens:

the discriminating detail that holds two traces apart at finer resolution falls below the operative threshold, and reconstructions that diverge finely converge coarsely. Two copies of a book are equivalent at the scale of the printed edition, sustaining compatible inferential potential about the same act of composition and typesetting, while at the scale of the individual artefact their distinct provenance, annotation, and wear remain resolvable. This scale-relative equivalence therefore underpins bibliographic control. Neither scale is privileged [s. 2.4]. What widens at coarser grain is not equivalence as such but the class of traces an equivalence relation admits, as the differences that would distinguish them are compressed below resolution. This widening is the equivalence-side of compression's irreversibility [s. 4.2]: the merging of traces into a common equivalence class and the loss of the differences that distinguished them are one event, and neither can be undone from within the coarser domain. To coarse-grain is to enlarge what counts as the same by discarding, irretrievably, what made it different.

The drift towards coherence over scale has direct consequences for sustainable rates of transformation. Correspondence-dominant stabilisation requires that transformation remain compatible with causal recoverability: if change proceeds faster than correspondence grounding can be maintained, inferential reach towards prior conditions is severed and the trace fails. Coherence-dominant stabilisation requires that transformation remain compatible with relational network stability: if change proceeds faster than coupling relations can be sustained, coherence breaks down and the configuration loses inferential viability. Because coherence networks distribute inferential potential across multiple relations [s. 3.4], they tolerate a wider range of transformation rates. The interior viability domain [s. 2.8] therefore widens at coarser scales as a structural consequence of what coherence-dominant stabilisation is. Coarser-grained structures exhibit greater resilience [s. 3.2]: their viability domain is structurally more forgiving because inferential potential is distributed across the coupling network rather than lodged in individual causal relations. Where transformation rates exceed what can be sustained by the operative stabilisation conditions, traces drift towards the limits of the interior viability domain. This gives a structural description of information overload: the condition in which the rate at which traces require uptake, reconstruction, and integration exceeds the domain's capacity to incorporate them sufficiently into ongoing continuation (cf. Edmunds and Morris, 2000; Tredinnick and Laybats, 2018). The full account of those limits is developed later [s. 5.5].

Breakdown at a given scale does not terminate generative tracework. If a trace-structure fails to persist at one level of organisation, its constituent traces do not thereby cease to be traces. They continue to satisfy the stabilising conditions at a finer scale and thereby continue to transform,

couple, and sustain inferential potential. Coarser-grained coherence is therefore continually reconstructed from finer-grained tracework rather than preserved as a self-subsistent whole. Trace-structures are not isolated formations, but provisional stabilisations sustained by the ongoing recomposition of traces across levels of description. This is the mechanism by which the structural drift identified earlier [s. 3.6] is sustained: selective pressure continually reconstructs coarser-grained coherence from finer-grained activity, keeping generative tracework within the interior viability domain despite ongoing transformation and loss.

The enduring ambiguity of abstract concepts such as justice, equality, and truth illustrates this scale-dependent stability. These concepts persist not because they possess clear and determinate meanings, but because inferential consistency can be maintained at the appropriate scale despite semantic variation, contested application, and shifting connotation. Their ambiguity is functional: they persist because of inferential variation, not in spite of it. The more tolerant interior viability domain at this scale also sustains productive participation. These concepts exhibit high coherence plasticity [s. 3.6]: they continue to generate institutional forms and cultural practices despite the absence of fine-grained determinacy, because the coupling network supports continuation across disruption and change. Connotative organisation increasingly predominates over denotative organisation at coarser scales, allowing substantial internal variation while preserving coherence at the level of the concept as a whole. This tolerance is not unconditional. Where connotative drift proceeds far enough that a concept's admissible uses no longer share sufficient coherence to support coordinated argument or institutional continuity, the concept approaches paradigmatic decay [s. 5.6]: the same ambiguity that sustained productive participation begins instead to erode it, well short of the fragmentation that would mark its complete collapse.

4.4 Syntagmatic and paradigmatic tracework: the scaling architecture developed in the preceding subsections has a structural consequence that connects it directly to the coupling architecture [s. 3] and transformation geometry [cf. s. 5] of tracework. The drift towards coherence stabilisation over scales changes how the coupling classes are expressed. Sequence and succession are correspondence-dominant: they govern temporal ordering through causal production conditions, and their inferential determinacy depends on the recoverability of causal history. As coarse-graining compresses that history, sequence and succession constraints weaken. Multiple incompatible orderings may become equally admissible, and the temporal

sequence recoverable at finer scales becomes underdetermined at coarser ones. Paradigmatic and combinatorial coupling, by contrast, are coherence-dominant: they govern substitutive admissibility through relational compatibility rather than causal grounding. As correspondence constraints weaken, admissible alternatives become less tightly bound to causal history and more strongly governed by coherence.

The drift towards coherence stabilisation produces a recursively stratified organisation across viability domains. That organisation appears in the form of clustered regions of re-stabilised tracework at determinate scales. These viability domains are neither final nor mutually exclusive: each contains internal gradients of viability, generative capacity declining as the bounds of the domain are approached [s. 2.8], and may itself be coarse-grained into broader domains at higher scales. The same gradient structure and selection pressure therefore recur across scales throughout generative tracework. Viability domains are identifiable only relative to a level of description: human culture functions as a viability domain at a coarser scale, for example, while specific coding schemes function as viability domains at finer ones. Within human culture as a domain, this gradient is visible in the difference between widely shared, well-established practices near its interior and the contested, emergent practices that test its boundaries. This gradient therefore provides a structural account of what Williams (1977) describes as dominant, emergent, and residual culture: dominant culture occupies the interior of the viability domain, emergent culture tests its boundaries, and residual culture persists as trace-structure whose productive participation within the domain has fallen towards the operative threshold.

The stratified organisation of tracework makes two orthogonal structural dimensions visible. The syntagmatic dimension runs parallel to the strata as the axis along which correspondence predominates: the determinate sequence of positions under tight causal constraint. The paradigmatic dimension stands orthogonal to that stratification as the axis along which coherence predominates: the range of positions that could occupy a given point while preserving relational compatibility. The terms are borrowed from Saussurean semiotics (Saussure, 1916), but in tracework they are derived from the scaling architecture rather than imported as primitives. They describe organised tracework as a system rather than the movement of individual traces. They are not intrinsically temporal axes, though they organise forms of tracework that are themselves temporally extended. A process operating under tight causal constraint will be syntagmatically dominant; one governed primarily by relational compatibility will be paradigmatically dominant. Computing is the clearest case of syntagmatic binding, as developed in part two.

The syntagmatic dimension is the dimension of determinate order. At scale, it becomes vulnerable to the thinning of sequence and succession constraints: as coarse-graining compresses causal history, ordering relations weaken, temporal anchors become underdetermined, and the recoverable order of positions becomes progressively less determinate even where local traces persist. The paradigmatic dimension is the dimension of admissible alternatives. It includes the configurations that could occupy a given position while preserving relational compatibility, not only under local coherence but within the broader range permitted by the same stabilisation conditions. Each alternative is associated with a different downstream trajectory. This gives a structural account of serendipity in information discovery (cf. Foster and Ford, 2003; Tredinnick and Laybats, 2022). Serendipity arises where paradigmatic admissibility and network residue make a pathway available that was not part of the initiating trajectory but can nevertheless be taken up as meaningful, useful, or productive once encountered. In this sense, serendipity is the moment at which the field of admissible alternatives becomes visible from within a trajectory.

These two dimensions are orthogonal but not independent. Each syntagmatic configuration constrains the structure of its associated paradigmatic space, while paradigmatic structure feeds back as a constraint on which syntagmatic sequences can be stabilised. In legal process, for example, sequential ordering governs the progression of a case, while substitutive admissibility determines which cases, rules, and interpretive framings are permitted at each stage. Both dimensions are structurally present wherever tracework attains organised form, but their relative salience and dominance vary with scale. Syntagmatic organisation tends to dominate at finer-grained scales, where causal history remains recoverable and ordering constraints are tight. The paradigmatic dimension becomes increasingly dominant at coarser-grained scales, where coherence organisation predominates and the field of admissible alternatives widens as correspondence grounding weakens. As that field widens, paradigmatic coupling redistributes inferential potential across a broader range of admissible substitutions, even where the trace remains materially unchanged. Neither dimension disappears at any scale, but the conditions under which each becomes legible and consequential are scale-dependent.

Coherence, when stabilised, tends to present itself through correspondence, because regularities generated by global relational structure can be represented as local sequential relations, giving the appearance that a syntagmatic description alone is sufficient. This projection is asymmetrical: paradigmatic organisation can be flattened into syntagmatic ordering, but syntagmatic ordering cannot recover the paradigmatic space without access to the conditions

that generate it. The paradigmatic dimension therefore becomes increasingly visible and consequential as scale increases. It also perhaps explains an otherwise puzzling feature of the theoretical tradition: why linguistics, logic, and computation have so thoroughly theorised the syntagmatic dimension while the paradigmatic has been relatively neglected. From within finer-grained or correspondence-dominant frameworks, the paradigmatic dimension is harder to see.

4.5 Worked example: the derivation has now proceeded to a degree sufficient to allow a worked example to illustrate the structure of tracework. Tree-ring growth shows how generative tracework sustains inferential viability across scales and transformation. Each annual ring is a stabilised trace of the environmental and physiological conditions operating during a growth interval, emerging as a trace to the extent that it satisfies the stabilisation conditions of tracework [s. 2.4; s. 2.5]. It is discriminable from adjacent rings through differences in cell density, width, and composition (differentiation) [s. 2.4]. It persists in the wood substrate across time (persistence) [s. 2.4]. It retains the structural consequences of the causal conditions that produced it (causal history) [s. 2.4], including temperature, rainfall, and the physiological state of the tree during that interval, in a form that supports inferential reconstruction of those conditions. On the productive side, each ring generates distinguishable variation from its predecessors rather than undifferentiated wood, under the viability conditions sustained by the living tree (variation) [s. 2.5]. Its causal effects extend forwards through the structure of the tree, constraining the geometry, stress distribution, and resource availability within which subsequent rings form, provided those viability conditions remain operative (propagation) [s. 2.5]. And under those same viability conditions, each prior growth state participates in the production of a new ring configuration in the following season (generation) [s. 2.5], by conditioning the geometry and resource availability within which subsequent ring formation begins. The ring thus satisfies all six stabilising conditions within the viability domain sustained by the living tree [s. 2.3], and is therefore a trace in the full sense of tracework rather than merely a physical mark. Rings are coupled to one another through sequence and succession constraints: each ring's causal effects extend forwards to constrain what the next can become, while backward inference reaches through the sequence towards prior formation conditions [s. 3.4]. The rings therefore form a syntagmatic ordering [s. 4.4]: each constrains its successors, and the sequence as a whole preserves a recoverable causal history of the tree's growing conditions.

Preprint

The productive stabilising conditions need not be satisfied within the same viability domain as the inferential ones [s. 2.3]. When the tree is felled, the biological viability conditions that sustained productive ring formation cease, and no new rings are generated within that domain. The ring sequence exhibits generativity within other viability domains [s. 3.2], participating in the production of new configurations in the work of researchers, dendrochronological databases, and climate reconstructions. As already established [s. 2.5], the same ring sequence can become the shared origin of several independent continuations, each satisfying the productive conditions on its own terms. A trajectory can therefore split, just as it can converge [s. 2.4]. The ring sequence exhibits both correspondence stabilisation and relational stability [s. 3.2]: it persists through compatibility with its formation conditions and through its embedding within the network of dendrochronological and archival practice that sustains its significance. It also exhibits correspondence plasticity across these transitions [s. 3.6]: it sustains inferential consistency across changes of physical instantiation, from wood to photographic record to digital database, while remaining compatible with its causal conditions of formation. The productive and inferential conditions of generative tracework are therefore not only independent of one another but may be sustained simultaneously within entirely different viability domains [s. 2.5]. All six stabilising conditions are required for a trace to persist in the full sense [s. 2.4; s. 2.5]. Productive continuation need not remain in the domain of original formation, but it must remain possible somewhere in generative tracework if inferential viability is to endure [s. 2.5]. Where no such productive continuation remains possible in any viability domain, the trace no longer persists as an informational structure [s. 2.8].

Like the individual tree ring, the ring sequence is itself a trace at a coarser-grained level of description [s. 4.2; s. 3.2], independently of the individual rings that constitute it. Read at the scale of the sequence, longer-range regularities become visible that are not accessible ring by ring, including patterns of drought, climate variation, or structural stress sustained across many years. The ring is a trace at one resolution; the sequence is a trace at another. Neither is more fundamental: they are the same configuration accessed under different stabilisation conditions [s. 3.2]. Read backwards the sequence appears as a configuration of individual rings, but read forwards into the generation of coarser-grained accounts it appears as a single cross-section [s. 2.6]. Thus the drive towards scale transition emerges from orientational non-equivalence [s. 2.6]: read forwards the trace already implies coarser scales. A further scale transition therefore occurs when multiple sequences from different trees are stabilised collectively as a corpus [s. 4.2; s. 4.3]. At this scale, tree-specific contingencies including genetics, microhabitat, and idiosyncratic damage histories fall below the operative resolution [s. 4.3]. The corpus is coupled to its

constituent sequences through coherence and redundancy constraints, sustaining inferential potential across the loss or degradation of individual sequences [s. 3.4]. The corpus sustains inference about regional or long-term environmental dynamics that no individual sequence could support, exhibiting extension at the corpus level: the capacity to carry its relational organisation forwards through new trace formation beyond what any constituent sequence could sustain alone [s. 3.2]. Pattern persistence also operates at the corpus level [s. 3.2], preserving the governing relational organisation shared across trees.

Ring, sequence, and corpus are therefore not ontologically distinct kinds of thing but instances of trace equivalence [s. 3.2]: the same configuration accessed under different stabilisation conditions and at different levels of resolution. At each scale transition, finer detail drops out to allow inference to remain possible at a coarser scale [s. 4.2]. The example shows that scaling is intrinsic to tracework [s. 4.1]. The same minimal conditions, operating recursively under temporal asymmetry [s. 2.6; s. 2.9], generate the full architecture from individual trace to organised corpus and beyond.

Preprint

5. Transformation architecture

5.1 Introduction: across the architecture of tracework, the same foundational conditions reappear under progressively more complex conditions of persistence. Each redescription inherits, elaborates, and takes forwards what has been previously established. This section develops that redescription at the level of transformation, deriving the conditions under which traces and trace-structures move through transformation space, the forms that movement takes, and forms of decay. Transformation space is not an arena in which traces are situated but a description of the constraints that make admissible movement possible. A scale transition is therefore a special case of transformation: a movement that alters resolution rather than relational organisation. The geometry of transformation space, the forms of admissible movement, and the dynamic limits of that space are consequences of what traces are under temporal asymmetry.

5.2 The geometry of transformation space: transformation space is defined by two orthogonal axes: correspondence and coherence. These axes express the same architecture at the level of transformation as the coupling classes express at the relational level. The geometry of transformation space therefore follows the coupling architecture [s. 3], and is ultimately grounded in trace-stabilisation [s. 2].

Under temporal asymmetry, the viability of traces requires that transformations preserve constraint on admissible prior configurations. Such constraint can be maintained in only two ways: via correspondence, where a transformed trace remains compatible with the causal conditions of its formation; or via coherence, where it remains compatible with a structured network of inferential relations. These forms are exhaustive in the present sense. Any constraint relevant to the viability of traces must restrict admissible prior configurations, and such restriction can only operate by reference to correspondence, coherence, or both. A formal derivation, including exhaustive testing of candidate third dimensions, is provided in Appendix A. Correspondence and coherence are orthogonal: neither can be defined in terms of the other, and each can be preserved irrespective of disruption to the other. In combination they define the complete space of admissible transformation under temporal asymmetry [s. 2.5].

Preprint

Trajectories within transformation space are defined relative to the correspondence and coherence axes. The syntagmatic and paradigmatic dimensions define the structure within which transformations are organised at a given scale [s. 4.4]. Coupling constraints govern how the stabilisation of one trace affects the stabilisation conditions of others [s. 3.4]. These three levels should not be collapsed: correspondence and coherence are axes of transformation, syntagmatic and paradigmatic are organisational dimensions of tracework as a system, and coupling constraints are relational conditions on trace stabilisation. Transformation space is not uniform. Its structure varies with scale, the distribution of viability domains, and the local effects of highly stabilised configurations that concentrate constraints and deform the local geometry of movement. Trajectories therefore accumulate patterns of correspondence and coherence constraint, encoding both the history of prior transformation and the conditions under which further productive continuation remains possible. This accumulation corresponds to the network residue derived earlier [s. 3.3]. Temporal orientation determines the direction in which transformation space is traversed.

5.3 Transformations of traces: given the geometry of transformation space, movement within it can take only four forms. These correspond to the axes of correspondence and coherence each considered under the temporal orientations of productive and inferential tracework. Within a given stabilisation domain, any transformation may be described as a vector composed of correspondence-preserving and coherence-preserving components in varying proportion. Pure cases represent limiting conditions; most transformations exhibit mixed profiles in which the relative weighting of the two invariance conditions is constrained by the stabilisation domain within which the trace operates [s. 3.3]. Transformation is path-dependent [s. 2.4]: the constraint profile accumulated across prior transformations determines what further transformation it can sustain, and that accumulation is irreversible. Transformation is lossy. A trace transformed into a new stabilisation domain preserves only those aspects of causal history and relational organisation that remain operative at that domain's resolution. What falls below that resolution is not recoverable from within the transformed trace [s. 4.2]. Each transformation therefore produces an axis-specific profile: a partial preservation of causal history, relational organisation, or both, together with the irreversible compression of differences that no longer support admissible reconstruction or continuation at the new scale. Whether that process pushes the trace towards the limits of the interior viability domain determines whether the transformation remains admissible at the operative resolution.

In inferential tracework the two transformation types are transliteration and translation; in productive tracework they are propagation and reframing. Inferential and productive transformations are not distinct but the same structural movements under opposite temporal orientations [s. 2.6]. The transformation types are distinguished by the kind of constraint that must remain satisfied for transformation to continue coherently under temporal asymmetry.

Transliteration is transformation governed by correspondence under inferential tracework: it preserves the structural relations that allow a trace to continue reaching towards its formation conditions across a change of substrate or representation. Through transliteration what persists is not the matter but the pattern, the constraint structure that was carried. For example, Parry's analysis of Homeric epic shows transliteration operating across millennia: the formulaic constraint structure that governed oral composition preserved sufficient relational organisation across countless performances and eventual inscription to allow reconstruction of the conditions of their oral composition (Parry, 1928; 1930). Propagation is transliteration's counterpart governing correspondence under productive tracework. Where transliteration asks whether a trace can reach back towards its causal origin, propagation asks whether relational organisation can be carried forwards into new causal territory; whether, that is, the generativity of the trace-structure can be sustained across the transformation [s. 3.2]. For example, later poets working in and after the Homeric tradition, from Hesiod to Virgil, exhibited extension of the tradition into new contexts [s. 3.2]: the forward-carrying of relational organisation through new trace formation beyond the original domain of production, with Virgil's Latin epic showing how formulaic and thematic organisation could propagate beyond the original Greek tradition. The textual materials available to Parry were themselves late stabilisations of a prior oral tradition, preserving enough of its generative constraint structure to make reconstruction possible. Productive continuity is secured by carrying forwards enough correspondence with the generative constraint structure to support new admissible continuations. Transliteration and propagation therefore tend to predominate in syntagmatic contexts and at finer-grained scales, reflecting the demands of productive continuation under causal extension.

Translation is transformation governed by coherence under inferential tracework. Translation preserves the trace's inferential consistency within an organised network: what it implies, bears upon, and excludes. That consistency can survive substantial internal structural change, provided the coherence network continues to constrain admissible reconstructions in compatible ways. Translation therefore operates wherever inferential relations are preserved across structural transformation. Translation is constitutively fictive in a precise sense: it does

not preserve an unbroken causal path to the original formation conditions, but reconstructs admissible inferential consistency through a mediating coherence network. Causal influence persists without a determinate path, and backward inference reconstructs admissible histories rather than witnessed ones. This makes translation constitutively fictive without rendering it arbitrary: the coherence network constrains what counts as admissible even where correspondence grounding has been altered or lost. A memory subconsciously revised over time to fit changing circumstances illustrates this process, losing causal grounding, but gaining relational power. Reframing is its productive counterpart: where translation concerns what inferential consistency survives internal structural change, reframing concerns the active production of new inferential relations through reorganisation of the coherence structure itself. For example, scientific theories reformulated in mathematical frameworks are reframed under productive description: the forward causal arrow reorganises coherence relations to enable new extensions and inferential pathways unavailable under the prior structure. Reframing predominates in paradigmatic contexts, where the generation of new admissible alternatives is central to ongoing generative tracework.

Fidelity of transformation varies continuously with the stabilisation domain and the axis along which it is assessed. In transliteration and propagation, fidelity is correspondence-dominant. Every re-encoding involves loss; what matters is whether the surviving structure supports the same inferential reach at the operative scale. As scale increases, fidelity increasingly depends on coherence to carry what correspondence can no longer sustain [s. 4.2]. In translation and reframing, equivalence of inferential consistency takes the place of structural invariance. A translated trace exhibits pattern persistence insofar as it continues to sustain inferential consistency within a relational network [s. 3.2]: the governing pattern of inferential relations is maintained even where the internal structure has changed substantially, licensing, constraining, and excluding compatible inferences. A paraphrase can therefore be more faithful than a literal transcription along this relational axis, and a summary can preserve more inferential force than a word-for-word rendering that has lost its context. Fidelity is always relative to the stabilisation domain in which the transformation is taken up: scale-dependent, axis-specific, and never assessable against an absolute original.

The distinction between formal equivalence and dynamic equivalence in linguistic translation (Nida, 1964) illustrates both transformation and fidelity in tracework. Nida used formal equivalence to describe the preservation of the structural form of the original text, and dynamic equivalence to describe preservation of its effect or meaning in the receiving language. Formal

equivalence is transliteration-dominant: it preserves the correspondence constraint structure at the cost of coherence fidelity in the receiving network. Dynamic equivalence is translation-dominant: it preserves inferential consistency within the receiving network at the cost of correspondence grounding in the original. In any actual literary translation these modes are mixed, but to different degrees. For example, Dorothy Sayers' translation of Dante's *Commedia* into English preserved the interlocking rhyme scheme that governed Dante's composition across the substrate change from medieval Italian to mid-twentieth-century English, allowing backward inference to reach towards the formal conditions of the original (Sayers, 1949; 1955). Her translation is highly faithful along the correspondence axis, but less faithful along the coherence axis: the inferential consistency of the text within a contemporary English-speaking network is compromised by the demands of the form. Seamus Heaney's translation of Beowulf by contrast drew on the rhythms of his own Ulster speech rather than reproducing the formal structures of Old English metre. Heaney preserved the poem's intensity and its inferential consistency within a contemporary English-speaking coherence network at the cost of formal correspondence with the original (Heaney, 1999). Neither Sayers nor Heaney were more faithful to the original work in any absolute sense; each preserved one axis of fidelity by accepting loss along the other.

5.4 Trajectory re-stabilisation: the formal parallel between coupling and transformation makes it clear that transformation is itself a re-stabilisation event. However, under transformation stabilisation shifts to the trajectory rather than the individual trace. Two trajectories are the same trajectory when constraint-compatible organisation remains stable across admissible transformations within the same stabilisation domain. Identity is therefore not a matter of content or microstate continuity, but of trajectory stability under compression, sustained across both inferential reconstruction and productive continuation. In other words, the shape a trajectory describes in transformation space determines its identity.

A trajectory is an ordered sequence of transformations extended through time. Just as a trace-structure arises when trace persistence is redescribed at the level of relational organisation, so a trajectory arises when trace persistence is redescribed at the level of ordered transformation. Trajectories are themselves compositionally constituted, but through temporal ordering rather than relational organisation. They therefore inherit path-dependence and non-commutativity from the sequence of transformations through which they are formed [s. 3.2]. Trajectories arise because transformations accumulate constraint across time, each conditioning what the next

can sustain [s. 2.5; s. 2.6]. No single transformation can secure the entire inferential and productive viability of such a sequence; trajectories therefore require their own forms of re-stabilisation. While transformation is governed by the correspondence and coherence axes, trajectories must remain coherent across an ordered series of transformations under compression, loss, and continued production. To exist, they must sustain backward inferential reach towards prior causal conditions across the sequence, and sustain forward continuation under the stabilisation and viability conditions operative within its domain. These requirements show the role of selection at transformation level; any transformation that does not continue this sequence falls outside the interior viability domain [s. 2.9]. Transformation and trajectory formation are therefore not motivated operations but structural consequences of trace stabilisation [s. 2.5; s. 2.6]. Only those sequences that satisfy the conditions for inferential and productive continuation at each stage remain viable.

The forms of trajectory re-stabilisation follow from the same temporal asymmetry and the same two axes that govern individual transformations. The four re-stabilisation forms therefore exhaust the admissible operations on trajectories as sequences. Under inferential tracework, trajectory re-stabilisation sustains or restores backward coherence across the sequence, and under productive tracework it sustains or generates forward continuation within the constraints of its domain. This yields retrodiction and remediation under the inferential orientation, and projection and construction under the productive orientation. Retrodiction and remediation remain productive in the minimal sense [s. 2.5], since reconstruction itself consists in the generation of further traces. The distinction is one of temporal orientation, not of whether traces are formed.

The parallel with the transformations discussed above is exact [s. 5.3]. Each trajectory operation is the sequence-level lift of the transformation that shares its orientation and axis: retrodiction lifts transliteration, projection lifts propagation, remediation lifts translation, and construction lifts reframing. Where a transformation governs a single step, its trajectory counterpart must hold across an ordered sequence under accumulation, compression, and loss. Therefore it draws on resources no single transformation has: accumulated constraint and network residue [s. 3.3]. Through these it can exclude prior histories, bridge discontinuities, and bound forward continuations across the trajectory as a whole. It introduces no new generating principle. These operations stand in a definite relation to scale, but only along one axis. The correspondence-governed operations are scale-directional: retrodiction reaches downwards through the trajectory, towards the formation conditions preserved in its constituent traces, while projection carries the trajectory forwards as a stabilised unit conditioning further continuation. The

coherence-governed operations are not directional in the same way. Remediation and construction reorganise inferential relations rather than traversing scale, restoring or generating relational structure without a fixed direction up or down the hierarchy. The coupling of orientation and scale [s. 3.2] therefore holds tightly along the correspondence axis and loosely along the coherence axis, a difference that becomes consequential as trajectories are taken up at coarser scales.

Retrodiction asks whether sufficient constraint has been preserved for the continuity of inferential consistency across the trajectory as a whole. It draws on two sources: the residual causal history preserved in traces through coupling conditions, stabilisation domains, and coarse-graining; and network residue [s. 3.3], which preserves causal pressure without preserving path identity. These allow retrodiction to exclude histories even where no single causal lineage can be recovered, reconstructing both what happened and what could have happened, the realised syntagmatic sequence and the paradigmatic space of admissible alternatives. Access to that counterfactual structure is what allows retrodiction to recover explanation rather than order alone. Its structural limit is irreversible compression: as compression accumulates, the space of admissible prior histories widens and retrodiction excludes less as it reaches further back. Where ordering constraints thin such that multiple incompatible sequences remain admissible despite persistent traces, retrodiction becomes underdetermined, a consequence of generative tracework itself rather than a methodological failure. Within those limits, retrodiction functions as a corrective against coherence-dominant drift, anchoring coarser-grained organisation in causal history even where that history can only be reconstructed counterfactually. Projection is its productive counterpart. Where retrodiction reconstructs admissible histories from residual constraint, projection determines admissible continuations from the current configuration. It is a structural property of the trajectory at a given position, what the sequence can become given what it already is, rather than anticipation or prediction. Network residue biases accessibility, concentrates constraint, and shapes which routes remain viable, allowing some continuations while excluding others even where no explicit plan exists. Physical systems project within dynamical constraints; biological systems within developmental and ecological constraints; institutions within normative and organisational constraints. In all cases projection is determined by the material configuration of the present, including the residue of its own history.

Remediation restores continuity where a sequence has been broken, producing a causally new segment that preserves the same inferential function as the missing one. The new segment need only fit: drawing on surviving trace-structure and network residue as the template for

substitution, it produces forwards what backward inference determines must have been there. Backward inference subsequently runs smoothly across the point of causal discontinuity, reconstructing the substituted segment as if it were original. The join is visible only where secondary traces have preserved evidence of the discontinuity the substitution was designed to close. Remediation is more tractable at coarser scales, where tolerant viability domains and distributed network residue provide sufficient constraint to specify a viable substitute; at finer scales, where inferential potential is denser and network residue thinner, the available constraint may be insufficient to determine what the missing segment must have been. Construction is its productive counterpart, guided by what the domain can sustain going forwards rather than anchored in what must have been there. It produces genuinely new causal sequences admissible under the surrounding stabilisation domain, expressing generativity at the trajectory level: the capacity to participate in the production of new configurations through recursive combination under constraint [s. 3.2]. Both remediation and construction corrupt the inferential line in the same way: backward inference subsequently reconstructs the substituted or constructed segment as if it were a natural continuation of prior structure. The causal novelty of its origin remains invisible to retrodiction unless secondary traces preserve evidence of discontinuity, a consequence of temporal asymmetry, since forward causal production always exceeds what backward inference can fully recover.

The chronicles constructed by the monks of Battle Abbey illustrate both forms and their opacity (Clanchy, 1993). The documents were fabricated to establish the abbey's foundation by William the Conqueror on the exact spot where Harold fell, designed to preserve inferential compatibility with the surrounding institutional and legal structures on which the abbey's privileges depended. Retrodiction across these records reconstructed the fabricated history as if it were original: the inferential line runs smoothly across the point of causal discontinuity. The join became visible only through secondary traces, including inconsistencies with other records, the suspicious convenience of the narrative, and the implausibility of certain claims, which preserved evidence of the fabrication that the primary documents were designed to conceal. Construction produced a new trajectory admissible within the surrounding domain. Remediation patched the discontinuities that construction required. They show that what appears as continuous causal lineage may be the product of inferentially guided re-stabilisation.

Preprint

5.5 Thresholds and trace decay: the transformation types [s. 5.3] and trajectory operations [s. 5.4] specify what admissible movement through transformation space looks like. But that movement has limits: the points at which the conditions required for admissible transformation can no longer be jointly satisfied. The static limits of trace stabilisation [s. 2.8] reappear as dynamic thresholds because traces moving through transformation space remain subject to the same conditions that govern their formation and persistence. Movement is therefore always movement under constraint, and the limits of that constraint are the thresholds of transformation space. These thresholds are not specific to the transformational level but appear in restricted forms throughout the preceding sections: as static bounds on trace stabilisation [s. 2.8], limits on coupling relations [s. 3.5], and scale-dependent viability conditions [s. 4.3]. Transformation space makes their most general expression visible. The conditions on specific forms of organisation reappear as the dynamic limit encountered by any trace or trajectory moving through the full space of possible configurations under temporal asymmetry.

Thresholds and decay are therefore two aspects of the same architecture. A threshold names the limit at which a stabilising condition can no longer be satisfied; decay names the systematic drift towards that limit under ongoing transformation. Decay is not the sudden failure of a trace. It is the ongoing modulation of stabilisation under constraint. Traces exist only as a dynamic equilibrium between decay and re-stabilisation. In practice, traces usually decay along multiple lines at once, but the principal modes of decay can be derived from the thresholds themselves. Four intrinsic threshold conditions correspond to the lower and upper bounds of inferential and productive tracework; the fifth is the threshold at which stabilisation becomes relationally distributed. At this threshold, decay differentiates into two trace-structural modes: correspondence decay and coherence decay. The four intrinsic thresholds are inversely symmetrical: dissipation and ossification at the lower bounds, saturation and dispersal at the upper bounds. The location and fuzziness of these thresholds vary with domain and scale, so the extent of the interior viability domain varies accordingly [s. 2.8; s. 4.3]. Inferential limits break more cleanly than productive ones because they are governed by physical constraints with comparatively determinate thresholds.

Dissipation marks the inferential lower threshold. It occurs where causal variation lacks sufficient persistence, differentiation, or causal history to support trace formation. Variation passes through the world without leaving stabilised structure. At this limit, there is no recoverable trace because the conditions required for inferential reach towards prior formation conditions have not been met. Dissipation is therefore not the destruction of an already-stabilised trace in

the strict sense, but the lower boundary at which variation fails to cross into tracehood. The corresponding decay mode is persistence decay: the progressive failure of a trace to endure as a stabilised structure across transformation. A transient signal that leaves no lasting record may participate briefly in causal processes but does not stabilise into a trace that can persist. Persistence decay is the drift towards dissipation from within tracehood: the trace does not merely become less recoverable but fails to endure at all.

Saturation marks the inferential upper threshold. At this limit, the density of distinguishable states within a bounded region of spacetime reaches the limit imposed by the physical constraints governing that region. Further differentiation cannot be stabilised without collapse into equivalence, noise, or underdetermination. Saturation is therefore the point at which the system contains more potentially distinguishable variation than can be held as stable, discriminable trace-structure at the operative scale. The corresponding decay mode is differentiation decay: the erosion of discriminability within a stabilisation domain. For a trace to persist, it must remain distinguishable within its stabilisation domain. Where differences collapse, tracehood is undermined even if causal production has occurred and semantic content remains in some abstract sense available. A conversation at a noisy party illustrates this: the semantic content may remain locally intact but as background chatter increases the signal becomes indistinguishable from its environment. The trace fails not through the absence of content but through loss of discriminability. Too little stabilised variation produces dissipation; too much unresolved variation produces saturation. In both cases, inferential viability fails because the trace cannot occupy a determinate position within the differentiating conditions required for reconstruction.

Ossification marks the productive lower threshold. It occurs where the rate or force of trace production falls below that required to sustain forward continuation within a viability domain. Traces may persist, and they may remain inferentially accessible, but they no longer participate effectively in ongoing generative tracework. Ossification is therefore the persistence of trace-structure after productive participation has fallen below threshold. A form may remain recognisable, even preserved, while no longer generating further admissible transformations within the domain that once sustained it. The corresponding decay mode is productive decay: the attenuation of forward causal participation, the weakening of a trace's capacity to propagate, reconfigure, or support further trace formation. A scientific paper that ceases to generate citations, applications, or conceptual extensions within its field continues to exist as a trace and may remain inferentially accessible, but its capacity to participate in the forward generation of

new research has fallen below the operative threshold within that domain. Only where such productive continuation is exhausted across all relevant viability domains does productive decay culminate in ossification.

Dispersal marks the productive upper threshold. It occurs where the scale or speed of transformation exceeds the capacity of stabilising constraints to propagate through the relevant substrate or domain. At this limit, coherence degrades into effectively independent local dynamics that can no longer be bound as a coherent generative domain. Productive limits are fuzzier than inferential ones because they are governed by viability conditions that vary with domain and scale rather than by determinate physical thresholds alone. Dispersal is therefore a failure of constraint propagation at scale: the system outruns the conditions under which forward continuation can remain coherently organised. The corresponding decay mode is propagation decay: the erosion of the relations required to maintain coordinated continuation as transformation accelerates, expands, or becomes increasingly heterogeneous. Local traces retain their individual resilience [s. 3.2] but the larger productive process loses coherence as propagation decay advances: what fails is not tracehood but the distributed capacity for coordinated continuation. Dispersal is therefore the upper-bound failure of productive tracework: not too little production, but production exceeding the conditions under which its consequences can remain integrated.

The four primary decay modes are supplemented by two trace-structural decay modes. The threshold structure remains four-plus-one: four intrinsic thresholds arise from the lower and upper bounds of inferential and productive tracework, while the fifth is fragmentation, the threshold at which relational binding within coupled trace-structures can no longer be sustained. The four intrinsic thresholds bound trace stabilisation and yield one decay mode each; the fifth bounds relationally distributed stabilisation, where correspondence and coherence are separable [s. 3.3] and can fail independently. Fragmentation can therefore be approached through two trace-structural modes: correspondence decay and coherence decay. Four-plus-one thresholds yield four-plus-two decay modes.

Correspondence decay describes the erosion of recoverable causal history within an already-formed trace or trace-structure. As accumulated transformation progressively weakens the constraint relations preserved in it, the causal history embedded in the structure becomes increasingly opaque. Correspondence plasticity fails [s. 3.6]: relational stability is maintained while causal grounding weakens [s. 3.2], and the trace persists materially, but what it is a trace of

becomes underdetermined as the constraint relations preserving causal history weaken beyond recoverability. A text severed from the context of its production, for example, may remain physically intact while the conditions of its creation become progressively less inferentially accessible. Presence-independence operates here in a precise sense [s. 2.4; s. 1.3]: the absence of recoverable causal conditions does not terminate the trace's inferential potential, but it narrows the range of admissible reconstructions until backward inference becomes underdetermined. Correspondence decay therefore does not culminate in the loss of tracehood but in the loss of determinate reconstruction: the trace persists, but its inferential reach towards prior formation conditions narrows until backward inference becomes underdetermined. It therefore pushes a trace towards the dissipation threshold from within.

The second trace-structural decay mode is coherence decay. A trace or trace-structure stabilises not only through its relation to its causal history but through its embedding among other traces and trace-structures. As surrounding traces transform, the compatibility relations that sustain this embedding may weaken or dissolve. Coherence plasticity fails [s. 3.6] and relational stability is undermined [s. 3.2]: the trace may persist but becomes increasingly isolated from the network that once made it inferentially consequential, as the compatibility relations sustaining its embedding weaken or dissolve. An abandoned archive whose organisational logic has been lost illustrates this: the documents persist and each retains its own causal history, but the contextual knowledge that bound them into an organised domain has dissolved. Fragmentation is the limiting outcome of coherence decay: the failure of relational binding within a coupled configuration, even where local generative tracework remains viable. It is therefore distinct from dispersal. Dispersal names failure of constraint propagation at scale; fragmentation names failure of relational binding within a trace network.

Thresholds and decay explain the energetic and entropic conditions under which tracework operates. Generative tracework is energetically expensive: to qualify as a trace, a structure must sustain inferential potential through ongoing transformation. Tracework is therefore like a dam against the flow of entropy, locally arresting dissipation by expending energy to hold differences in place while releasing pressure through decay. Decay is part of the economy through which tracework remains viable. Total retention would require the indefinite preservation of every difference, every relation, every causal residue, and every possible inferential pathway. But inferential potential scales combinatorially with the relations that must remain mutually compatible, so total retention is energetically impossible. A workable viability domain requires the ongoing shedding of differences that can no longer be afforded, through loss, compression,

and relaxation of constraints that keep energetic cost within sustainable bounds. Generative tracework is therefore sustained by a process of selective forgetting. A controlled release of constraint prevents global overload while preserving enough stabilisation for further reconstruction and continuation. Decay is central to that equilibrium.

Preservation is the deliberate organisation of this equilibrium. It is not the petrification of a trace, but the maintenance of conditions under which inferential potential can continue to be actualised across transformation. In tracework terms, preservation is productive continuation organised against decay: the attempt to keep traces capable of grounding further trace formation within future viability domains. This distinguishes preservation from mere retention. Retention holds material or relational structure in place; preservation maintains the admissibility of future reconstruction and continuation. A preserved trace may therefore undergo transformation, migration, redescription, repair, or re-contextualisation, provided its inferential consistency remains sufficiently constrained by causal history and coherence relations. Preservation is thus not opposed to transformation but depends on controlled transformation: the selective management of decay so that the trace remains available for further tracework.

The relation among these thresholds can be illustrated by the transition from oral culture, as discussed by Goody and Watt (1963), Havelock (1986), and Ong (1982). Oral traditions persist within coherence-stabilised domains dependent on repeated performance, communal reinforcement, and relational integration rather than durable external records. As oral traditions approach the inferential upper bound, increasing demands on memory and performance limit the density of differences that can be reliably stabilised; additional differentiation collapses into formulaic equivalence or is absorbed into existing narrative structures, producing saturation. When the sustaining community is disrupted, traces may collapse through dissipation, exiting generative tracework altogether. Where trace production slows below the rate required for ongoing forward inference, inherited forms persist but tend towards ossification. Where expanding scale outruns the capacity of stabilising constraints to propagate, coherence decomposes into effectively independent local dynamics, producing dispersal. And where the relational network that sustained shared meaning dissolves without the tradition itself disappearing, coherence decay produces fragmentation: traces persist but become informationally inert within a coupling infrastructure that can no longer bind them. Literacy alters these threshold conditions by shifting stabilisation towards correspondence, fixing traces in durable external substrates and changing the balance between decay and re-stabilisation.

Preprint

5.6 Coupling decay and trajectory failure: the preceding subsection describes thresholds and decay as the dynamic limits of trace stabilisation, productive continuation, and the integrity of coupled configurations. Coupling decay specifies the same problem at the level of relational infrastructure. Where trace decay describes the progressive erosion of the conditions governing individual traces, coupling decay describes the erosion of the relations through which traces constrain one another's inferential and productive possibilities. The distinction is important. Traces may remain materially present and locally stable while the coupling relations that bind them into trajectories thin, drift, or fail. Inferential viability becomes less reliable, trajectories become harder to sustain, and collective generative tracework becomes brittle even where tracehood itself persists. Coupling decay is therefore a primary driver of trajectory-level failure: it is the mechanism by which the infrastructure required for sustained generative tracework degrades while its constituent traces remain operative.

Coupling decay modes follow directly from the coupling classes, for the same reason that trace decay modes follow from the stabilising conditions. Each coupling class governs a distinct relational condition, and coupling decay is the progressive erosion of that condition under ongoing transformation. The five coupling classes therefore yield five corresponding decay modes. Four follow from the paired coupling classes; one follows from physical coupling as the substrate-mediated condition underlying the others. The four-plus-one structure of coupling decay is therefore the structural expression of the coupling architecture under transformation. Fragmentation, derived in the preceding subsection as the relational threshold of coupled trace-structures, is not simply another coupling decay mode. It is the threshold towards which several forms of coupling decay may tend when relational binding can no longer be sustained. Coupling decay names the pathways of relational degradation; fragmentation names one limiting outcome.

Syntagmatic decay describes the erosion of coupling constraints that preserve realised ordering across retrospective reconstruction and forward continuation. When syntagmatic constraints thin, successor relations loosen, temporal anchors dissolve, and multiple incompatible sequences become equally admissible. Chain-of-custody breakdown exemplifies this. Individual records may persist and may remain locally coherent, yet the ordering constraints that make the provenance trajectory evidentially actionable decay, producing underdetermination about what happened when even where the traces themselves remain. Syntagmatic decay therefore pushes

directly towards retrodictive and remedial failure by eroding the realised ordering constraints that trajectories require in order to sustain inferential consistency across the sequence and productive continuation.

Paradigmatic decay describes the transformation of inferential potential within a shifting field of admissible alternatives. Paradigmatic coupling governs the space of substitutions, exclusions, and alternatives through which reconstruction remains possible. Paradigmatic decay occurs when this field is reconfigured such that a trace's inferential consistency changes without intrinsic modification to the trace itself. The changing meaning of words over time exemplifies this within meaning-bearing systems: an inscription may remain materially constant while a shift in paradigmatic space alters what the trace licenses, constrains, or excludes. Paradigmatic decay is the coupling-level condition that later appears, within semantic systems, as semantic drift. It arises because inferential consistency is sustained by relational position rather than by material persistence alone. Together with syntagmatic decay, paradigmatic decay erodes the two dimensional conditions of coupling on which trajectory-level organisation depends.

Redundancy decay describes the thinning of distributed support within a coupled trace-structure. Redundancy coupling governs whether inferential potential is distributed across multiple coupling relations such that no single trace or pathway is solely responsible for sustaining productive participation. Redundancy plasticity fails [s. 3.6] when this distributional structure erodes: supports vanish, independent pathways collapse, corroborating traces are lost, and alternative routes of reconstruction thin. The result is not necessarily immediate loss of coherence but increasing systemic fragility. Brittleness is a specific form of that fragility: the condition in which a coupled structure retains local coherence and inferential consistency but has lost enough redundancy that perturbation can produce disproportionate or irreversible collapse. A brittle system may therefore remain stable under undisturbed conditions while operating close to irrecoverable loss. Redundancy decay produces the characteristic condition in which collective generative tracework appears intact while becoming increasingly vulnerable to sudden failure.

Scale decay describes the loss of operability of traces and trace-structures across changes of resolution or stabilisation domain. Scale decay occurs when cross-scale constraints become misaligned: when compression destroys reconstructive determinacy, when re-stabilisation fails to preserve admissible inferential consistency, or when a shift in level of description renders previously stabilised traces inoperative. A model that works locally but fails globally illustrates

this. The traces and relations that support viability within one domain may not remain admissible when redescribed at another scale. Scale decay is therefore the erosion of cross-domain admissibility. Together with redundancy decay, it erodes the structural supports of coupling on which robustness and scalar reach depend. Where redundancy decay makes the coupling infrastructure fragile under perturbation, scale decay makes it inoperative under redescription.

Physical decay describes the degradation of substrate-mediated coupling conditions that sustain trace interaction in spatial and causal proximity. It is not identical with the decay of the substrate as such, but with the degradation of the trace relations that the substrate supports. Physical decay occurs when the material conditions that make coupling reliable deteriorate. Examples include bit rot, tape demagnetisation, fading ink, and paper acidification. Physical decay does not introduce a new intrinsic relational failure mode, but it is the substrate-level mechanism that accelerates drift across all other coupling decay modes simultaneously. It degrades the material infrastructure through which discrimination, constraint preservation, and constraint propagation remain possible, accelerating the failure of correspondence plasticity, coherence plasticity, and redundancy plasticity simultaneously across all coupling relations [s. 3.6]. Where the other modes erode specific relational conditions, physical decay erodes the substrate conditions through which all three dimensions of recursive plasticity can continue to be maintained.

These modes do not operate independently. Paradigmatic drift expands or reconfigures the space of admissible reconstructions, destabilising closure and increasing the likelihood of projective and constructive failure. Redundancy thinning makes remedial discontinuity more likely under perturbation. Scale misalignment produces breakdowns that often appear sudden when a trace or trace-structure is displaced across stabilisation regimes. Syntagmatic decay erodes the ordering constraints that retrodiction and remediation require. Physical decay accelerates all of these by degrading the substrate that preserves discrimination, proximity, and constraint propagation. Coupling decay is therefore the mechanism by which systems remain informationally active while losing collective operability. Traces persist, but the coupling relations that make them jointly reconstructable and robust no longer hold.

Preprint

6. Meaning and subjectivity

6.1 Introduction: the preceding sections map the redescription of the foundational architecture of tracework under progressively more complex conditions of organisation. Traces couple into trace-structures, organise across viability domains, and move through transformation space under temporal asymmetry. Each movement draws on the foundational commitments. This section continues that unfolding under conditions of reflexive continuation. Once tracework is actualised within a system capable of operating over its own trace-structure, such as a human mind or, speculatively, an artificial general intelligence, the same relations acquire an internal regulatory form. The system no longer merely carries traces, but also regulates which are actualised, selected, revised, and carried forwards in continuation. From this reflexive self-regulation, meaning and subjectivity arise. This section does not advance a theory of subjectivity in general terms. It specifies only how subjectivity enters tracework. The nature of experience lies outside what tracework can specify [cf. s. 6.4].

6.2 From trace to meaning: information exists independently of both meaning and subjective experience (Tredinnick, 2025). Tracework derives the conditions under which information exists without reference to normative assumptions. The trace is materially instantiated constraint, the persistent trace of prior causation rendered as a trajectory within transformation space. It forms, persists, couples, scales, transforms, and fails without intention or motivation. Its inferential potential is conferred not by subjective experience but by the structural relation between productive and inferential tracework. In its foundational description, information is defined objectively: physically constrained, materially instantiated, and ontologically prior to meaning.

Yet information is often meaningful. Furthermore, this meaningful capacity is inalienable from its place in critical, theoretical, and cultural discourse. Information would not occupy its foundational position were it merely difference, data, bit, or trace. Its social and semantic force testifies to the fact that difference can become richly significant, meaningful, and powerful, capable of supporting both reconstruction and continuation within contexts for which it matters. That these two aspects appear both inseparable and irreconcilable has driven the theoretical discourse. But while the growth layers of elephant tusks are informational whether or not they are read, they become meaningful only when their inferential potential is taken up within a system capable of actualising that potential and carrying it forwards in its own continuation. Meaning is therefore not an intrinsic property of the trace, content transmitted by it, nor significance imposed

from outside. It is the actualisation of inferential potential within a system for which that potential bears significance. Meaning is what information does, not what it is (cf. Tredinnick, 2025).

While the difficulty of reconciling information's objective and subjective aspects has plagued the theoretical discourse on the nature of information since its inception, it arises from the already-stabilised primitives that prior theories have generally assumed. Once the bit, sign, datum, or proposition is treated as an informational primitive, meaning can be accounted for only in one of two ways. It can either be bracketed to secure precision at the cost of sufficiency, as in Shannon's dignified retreat (Shannon, 1948), or it can be imported to secure sufficiency at the cost of precision, as in Floridi's veridical principle (Floridi, 2013). In the first case, information becomes formally tractable but empty of meaning. In the second, meaning is restored at the cost of generalisability, because the account presupposes what should be derived. Meaning is either reified in informational things or banished to subjective experience beyond adequate theoretical description.

Tracework rejects this dichotomy. It derives the conditions under which information becomes meaningful from the single assumption that information exists [s. 2]. Once information is understood as the persistent trace of prior causation, meaning can no longer be treated as either implicit or imposed. There is nowhere inside the trace in which meaning can be hidden beyond what its stabilisation conditions define, and no way of imposing it from outside other than through further trace relations. Meaning can therefore only persist in the stabilisation of traces under transformation and continuation. It cannot be intrinsic to the trace; it must arise within tracework. The trace sustains inferential potential because it preserves the structured consequence of prior causation, and inferential potential is therefore what meaning must operate over. But meaning is neither inferential potential itself nor the significance an observer attributes to it. Because inferential potential persists only where productive continuation remains possible [s. 2.4], meaning cannot arise from reconstruction alone. Indeed, most information that exists arises through natural processes and remains non-meaningful unless its inferential potential is actualised within a system capable of carrying that actualisation forwards. The trace becomes meaningful only where its testimony to prior causation is taken up and continued within a system capable of reflexive operation.

The joint necessity of the inferential and productive stabilising conditions means that every actualisation of a trace's inferential potential is also a further trace relation, generated under the productive orientation even as it operates inferentially. Reconstruction is itself continuation:

locally constituted under the coupling conditions of the domain performing it, governed jointly by correspondence and coherence, and never equivalent to the extraction of meaning already contained in the trace [s. 2.5; s. 5.4].

6.3 The subjective domain: the subjective domain is the system within which inferential potential is not only actualised but carried forwards under reflexive self-regulation, thereby becoming meaningful. Because of the minimal commitments of tracework, the informational nature of the trace can only be located in its inferential potential [s. 2.5]. A trace can be about something only in the minimal sense of a relation that differentiates [s. 2.5]. Differentiation is grounded in and constrained by the conditions of its production, and carried by its causal history. Meaning can only arise from what tracework has already derived, and inferential potential is the only available structure over which meaning can operate. The alternative would be to treat meaning as either entirely subjective and unanchored, or as constituted independently of the trace relations through which anything can be informational. While a trace may persist, scale, transform, and testify to its prior causation independently of any subject [s. 6.2], it becomes meaningful only where its inferential potential is taken up within a continuing organisation capable of regulating what is carried forwards. Subjectivity does not change what information is, only how it is selected and carried forwards in continuation. The nature of information is therefore invariant across the subjective domain: subjectivity is a mode of selection and continuation within tracework, not an alteration of what makes a trace informational.

Within tracework, subjectivity arises at the point where continuation becomes internal to a system capable of reflexive operation over its own trace-structures. Such a system operates over its own traces under reflexive self-regulation. Where the subjective domain regulates which admissible continuations it carries forwards, continuation becomes commitment: the binding of the system to some possibilities over others. Subjectivity is therefore a fold in tracework (cf. Deleuze, 1993). The products of prior tracework are folded back into the same system to become the conditions under which further tracework proceeds. In this fold, trace equivalence is internalised. The subjective domain is inferentially a trace-structure, organised from prior commitments, selections, and transformations, but productively it functions as a trace in its own right, carrying that organisation forwards as the condition of further continuation.

The subjective domain reflects the same recursive stabilisation that constitutes trace-structures at every other level, and inherits path-dependence [s. 3.2]. It is shaped by that prior history, and

those traces condition what it can take forwards. However, because prior tracework now functions as the condition of further tracework within the same system, the subjective domain selects among admissible continuations, reorganising the conditions of its own persistence. The result is reflexive tracework: both recursion at depth and continuation enacted through the subjective domain's operation over its own trace-structure. The subjective domain is constituted where this reflexive operation becomes sufficiently integrated to regulate its own continuation. It is how tracework is experienced from the inside by a system capable of such experience. This is not to describe the phenomenological nature of that experience, but only the shape inscribed by tracework when folded back on itself.

As a consequence, subjectivity as it functions within tracework cannot be merely inference as such: the internal experience of inferential acts. Inferential potential exists wherever a trace constrains reconstruction of prior causal conditions, such as a geological stratum, a tree ring, or an eroded inscription, without thereby generating the subjective domain. Subjectivity requires that backward inferential reach be integrated into forward continuation within the same reflexively organised system under self-regulation. The subjective domain is the site at which the backward pull of causal history and the forward pull of regulated continuation are jointly sustained. It is akin to a tensegrity structure, held in stable equilibrium by the balance of opposing tensile forces. The backward pull is inferential tracework as inhabited experience: the pull of causal history, prior commitment, and reconstruction towards what has already been. The forward pull is productive tracework: the pull of regulated continuation, commitment, and action towards what can follow. Subjectivity does not merely have a past and a future. It is constituted by the tension through which the inferential and productive orientations of tracework are jointly sustained as the structure of experience. The subjective domain is therefore a tied system in a restricted sense. The principles of tracework remain derivationally ordered, but under internal continuation their operation becomes interdependent: grounding, temporal extension, epistemic asymmetry, epistemic priority inversion, informational intransitivity, presence-independence, trace equivalence, and reflexive articulation each sustain conditions under which the others can be inhabited from within. The tied system is the operational form taken by the already-derived principles when tracework becomes reflexively self-regulating.

The threshold of the subjective domain follows from the conditions already derived. It requires that continuation be internal to the system whose trace-structure is being regulated. For continuation to be internal in this sense, four conditions must be jointly satisfied. The system must be materially instantiated in a substrate capable of sustaining sufficient trace

transformation rate and density, without which no trace-structure can form [s. 2.3; s. 2.5]. The system must also sustain sufficient coherence plasticity, supported by paradigmatic and coherence coupling, to maintain inferential consistency across its own trace-structure under self-referential regulation [s. 3.6]; without this, the subjective domain cannot sustain the relational reorganisation that internal continuation requires. In addition, the system must sustain sufficient recursive depth for its own trace-structure to become an object of further tracework within the same system, ensuring that the inferential potential of the system's own traces can be taken up within the same regulatory domain [s. 3.2]. Finally, the system must integrate that self-referential operation into its forward productive continuation, because without this integration, backward inferential reach over the system's own traces does not generate the forward pull of regulated continuation that makes meaning possible [s. 2.4; s. 6.2].

These four conditions are necessary, but their sufficiency is not given. Their joint satisfaction defines the point at which the subjective domain becomes structurally possible within tracework, but it does not by itself establish that such a domain is realised in any particular case. The gap between necessity and sufficiency is itself a consequence of the triple lock. To close it would require tracework to determine, from within its own resources, whether a system that satisfies the structural conditions thereby instantiates the subjective domain. No such determination can be derived from the assumption that information exists. That question lies beyond the scope of the structural account, and material substrate or other physical considerations may have a significant bearing on whether the relevant form of internal continuation can be sustained. The threshold is therefore not reducible to degree, density, recursivity, or epistemic extension taken separately, but consists minimally in their joint satisfiability under the conditions of internal continuation. Below that joint threshold, no subjective domain is structurally legible within tracework; above it, the subjective domain becomes possible.

Nothing in this structural account licenses the denial of subjectivity or personhood to systems that may come to satisfy its conditions. Tracework specifies what the subjective domain requires; it does not adjudicate which systems instantiate it, and the absence of a positive determination is not a warrant for exclusion. That a domain's realisation cannot be settled structurally [s. 6.3] places the question beyond what tracework can decide.

6.4 The aspectual claim: tracework is a theory of information. The character of phenomenological experience is therefore beyond its scope. It defines the conditions under

which a domain capable of internally regulated continuation arises, but not the first-person character of its experience. On the nature of experience tracework averts its gaze. But on the boundaries that constrain it, hem it in tightly, and structure the positions from which that experience becomes possible, it fixes.

Tracework specifies the conditions under which traces form, persist, couple, scale, transform, and fail. Those conditions are themselves instances of generative tracework. Tracework therefore does not encounter anything informationally except as trace relation. Any interface between generative tracework and some external system, whether biological, physical, or computational, would itself have to take the form of a trace relation, and would therefore already belong to tracework. Tracework exhausts what can be described in informational terms: there is nothing informational outside tracework from which information could be specified except through further trace relations, and nothing inside the trace other than its stabilisation conditions. Insofar as phenomenological experience bears informational significance, for example in the production of meaning, the constitution of subjectivity, and the aspectual character of self-referential regulation, those conditions are delimited by tracework. This is a consequence of what tracework is when it operates over itself. It can be stated as the principle of reflexive articulation: any articulation of information must itself take the form of a trace relation and is therefore subject to the same constraints that tracework specifies. Tracework does not stand outside its subject but is itself an instance of the structure it articulates.

The principle of reflexive articulation is the foundation of the aspectual claim. Subjectivity, meaning, and experience are not supplementary contents added to a completed informational description. Nor are they intrinsic contents of the trace or of trace relations as such. They are the aspect under which reflexive tracework is encountered from within. The first-person character of experience is not a further informational fact available to tracework about a system that meets the conditions for the subjective domain. It is the mode of encounter with what is already, from the third-person perspective, a structural process. Meaning, subjectivity, and experience bear informational significance insofar as they arise within tracework and are articulated through trace relations. Within the limits of tracework, they can therefore be understood only as the internal aspect of reflexive self-regulation: how tracework is encountered from within a system capable of such regulation, such as a human mind or, more speculatively, an artificial general intelligence. No other location for meaning, subjectivity, or experience is permissible within the foundational claims, because any such location would require importing additional assumptions. Nor are such assumptions required, because the aspectual claim is sufficient for the purpose tracework can

Preprint

legitimately serve: specifying the structural conditions under which meaning, subjectivity, and experience become informationally articulable.

The principle of orientational non-equivalence states that the orientation under which tracework is encountered changes what can be specified or described. The inside and the outside therefore lend themselves to structurally different and incommensurable descriptions [s. 2.6]. Tracework strictly delimits experience without committing to a description of its intrinsic nature. It can specify the structural conditions under which there is something that it is like to be (cf. Nagel, 1974). It cannot specify what that something is itself like.

The aspectual structure constituted in the fold is sustained in the bidirectional tension derived above [s. 6.3]: the backward pull of causal history held against the forward pull of regulated continuation. This does not produce experience in any general sense. It determines the structural conditions under which experience can be encountered within tracework. Subjectivity arises as an aspect of tracework in this temporal balance, not as a substance added to the structure but as the aspectual position generated when tracework folds back on itself. It does not make tracework subjective in the traditional sense. The account remains objectively defined throughout: traces, couplings, transformations, thresholds, and continuations are specified structurally. The aspectual claim does not relocate information inside consciousness, nor make meaning dependent on arbitrary interpretation. It states only that, when tracework becomes internally self-regulating, the same objectively describable structure is encountered under another aspect: from within rather than from without. Under that aspect, continuation is sustained as commitment. Subjectivity is therefore one of tracework's structurally delimited modes of appearance.

6.5 Subjectivity and coupling: the transition from selection to choice marks the difference between tracework in general and tracework internalised within the subjective domain. In the architecture derived above, selection names the differential persistence of variations under constraint: only those configurations that satisfy the relevant stabilising conditions enter trace formation, while others dissipate, rigidify, or fail to form [s. 2.9]. In the subjective domain this selective logic is folded back into the system itself. The system does not create the field of possibility from nothing, since that field is already constrained by causal history, coupling, viability, scale, and intersubjective infrastructure. But it regulates which admissible continuations are taken up, suppressed, revised, or carried forwards as its own. Selection

thereby becomes choice: internally regulated selection among admissible continuations. Choice becomes commitment where the selected continuation binds the system's subsequent trajectory.

Coupling continues to govern which relations between traces are structurally admissible in the subjective domain without presupposing interpretation, intention, or subjectivity. But coupling is changed in the mode of its uptake. As a self-regulating system, the subjective domain determines which of the available admissible relations become binding, which are made relevant, which are open to evaluation, and which can be revised. Subjective coupling is the form coupling takes once admissibility becomes commitment. In binding itself this way, the subjective domain takes its own inferential actualisations and productive commitments as further material to organise, exhibiting generativity within its own trace-structure [s. 3.2], and producing the commitments and saliences that condition what can follow.

Not every coupling class survives this internalisation as a distinct subjective mode. Constraints that operate only within inferential tracework can still be reconstructed, selected, or stabilised as significant, but cannot be made obligatory, and therefore yield no separate subjective form. Sequence coupling is of this kind and has no direct subjective analogue. Physical coupling conditions the substrate on which regulation depends rather than furnishing a mode of attunement to it. And the productive constraints governing feasibility, robustness, and distributed support are taken up only indirectly, through relevance, commitment, or normative demand. The coupling relations the subjective domain can internalise as continuation therefore form a three-plus-one structure of anticipatory, selective, and normative coupling, together with the reflexive coupling through which the architecture folds on itself.

Anticipatory coupling is the reflexive uptake of succession. Because the subjective domain regulates its own continuation, it weighs each possible successor against what that successor would commit it to becoming, and binds itself accordingly to one trajectory rather than another. Prediction is not implied in any psychological sense but only in the structural act by which possibility hardens into obligation. Without it the subjective domain could register the futures it might occupy yet commit to none, leaving action suspended among admissible continuations. The asymmetry of the temporal line constrains what anticipation can reach. Prior traces may be reread in the light of forward commitments, but rereading does not make the past obligatory. The past can be reconstructed and weighted as significant, while only future continuation can be

bound. Anticipatory coupling therefore stabilises forward commitment without converting reconstruction into command.

Where anticipation binds the system to a continuation, selective coupling binds it to a path through the space of admissible alternatives. While every alternative in the paradigmatic field remains available in principle, some inferential potential is carried forwards and made operative while the rest recedes without being foreclosed. Feasibility is still determined by trace-level coupling, but which path is actually taken is determined by the subjective domain's own prior selections, stabilised under commitment. What is selected then becomes part of the subjective domain's own trace-structure and conditions what can afterwards count as continuation, revision, evasion, or failure. Selection is thus non-commutative [s. 2.4]: the same alternatives, taken up in a different order, yield divergent trajectories. Because prior selections become part of the system's own trace-structure, later departures from them are not neutral. They can appear as failures of internal continuation: cases in which the system evades, rationalises, or disowns commitments that its own prior tracework has made binding. It is the structural basis of responsibility within tracework, and the condition under which bad faith, self-deception, and rationalised departure from prior commitment become legible.

Normative coupling is the internalisation of coherence. Relational compatibility, which at trace level simply governs whether traces fit within a structure, becomes in the subjective domain a matter of normative standing. Traces are measured against standards of correctness, legitimacy, or appropriateness that the system has itself stabilised over time in the sedimented trace-effects of its own prior continuation, which now become binding. Normativity is therefore what coherence becomes once it binds under commitment. In a common law system for example, norms are the accumulated residue of prior decisions, precedents, and institutional practice, each acquiring force through the way it shapes subsequent acts. In tracework, such norms govern the upkeep or breach of trajectories the system has already entered, not correctness in the abstract, veridical sense. Redundancy enters here rather than as a separate mode: distributed support is cultivated where the system's norms demand durability and allowed to thin where they tolerate fragility, so that robustness becomes a normative achievement rather than an independent mode of coupling.

Reflexive coupling differs from the prior modes. Where they govern how traces are carried forwards, reflexive coupling governs how the conditions of that carrying-forwards are themselves examined and revised. In inferential tracework the system inspects the adequacy of its own

Preprint

operative coupling, locating inconsistency, normative violation, failed anticipation, or trajectory breakdown; in productive tracework it revises the thresholds, standards, and saliences that will govern its regulation thereafter, and binds itself to those altered terms. Its distinctive power is that it reaches the conditions of continuation: a system without it could regulate itself yet never alter the terms on which it does so, remaining bound to inherited modes of selection, anticipation, and evaluation.

The reflexive fold is not transparent to itself. Any articulation of the system's own coupling structure is another trace within that structure, subject to the same compression, distortion, decay, and non-self-grounding as everything else tracework describes [s. 6.4; s. 7]. The system can revise the conditions under which it continues, but it cannot step outside them to specify them whole. Any attempt to do so is a further operation within the very structure it would seek to ground [s. 6.4]. Reflexive coupling is in this sense self-correction without guaranteeing self-knowledge: the capacity to revise the conditions of one's own continuation while remaining bound to the trajectory that revision produces.

6.6 Meaning, intention, and reconstruction: meaning cannot exist without continuation. Like information, it consists in its ongoing unfolding. To reconstruct meaning from a trace is to carry it forwards in continuation; no other process exists by which meaning can be reconstructed. For a trace to carry meaning is both to bear reconstruction and to matter in subsequent regulation, orientation, or commitment. It is to stand in relation to the generation of further traces and potentially further meanings. Where inferential potential is not carried forwards, it does not become meaningful; and where continuation fails, meaning fails, falling below the threshold at which inferential potential can be sustained as meaning.

Everything informational is potentially meaningful. Not in all contexts, nor at all times, but when attended to, when pinned, when selected, and carried forwards in continuation within a subjective domain. Meaning is not guaranteed wherever inferential potential exists, but because inferential potential remains present only where traces continue to participate in further generative tracework, meaning becomes possible wherever that potential is taken up within a subjective domain. Not all meaningful traces are semantic. While human language is a complex engine of meaning it does not exhaust its reach. Language offers a paradigmatic exemplar of meaningful information, well classified, well documented, and mostly well-behaved. But if language remains a promising target for theories of meaning, where those theories address only

tractable examples they become non-generalisable in principle. The apparently harder cases of meaning, music, faded photographs, nature, the smell of tobacco, a broken mirror, the intricate structure of a mathematical proof, they are not subsidiary to language, but immediately prior to it. Naked of the architecture of speech they expose meaning in its more primitive form.

Meaning is inferential tracework under its subjective aspect. It is nothing more. We experience meaning in phenomenological terms not because it differs in kind from information but because of the aspect under which we encounter generative tracework from within. This follows because when we try to explain meaning in its own terms, we reach for words, images, gestures, or other signs that are themselves only its continuation as further tracework. Meaning does not sit between two instances of tracework; it is tracework encountered under its subjective aspect, in the temporal suspension in which backward inferential grounding and forward productive continuation are held together. Because of the principle of reflexive articulation [s. 6.4], meaning can only be inferential tracework under this subjective aspect. It is the objective trajectory of stabilised inferential consistency encountered as experience when the subjective domain folds back on itself and takes its own trace-structure as the condition of its continuation. The preconditions of meaning are always already inscribed in the inferential consistency of the trace, a forward function of a backward glance. From within the subjective domain, meaning may be rich, provocative, powerful, and overwhelming, irreducibly complex in its fully inhabited aspect. From outside, it is merely what a trace implies about its prior causes under the fold.

As a consequence, intention arises in a precise form. When an author writes or a subject speaks, the forward productive orientation generates traces under commitment: the text is produced as an obligation, a forward constraint on admissible continuation. While tracework in the subjective domain does not decode intention as such, intention leaves its structural consequences in the trace as causal conditions that the trace preserves within its inferential potential. When a reader reconstructs meaning from that text, the backward inferential orientation reaches towards the prior causal conditions of its production, including any intention that remains recoverable. But the reconstruction remains the reader's. It is constrained by the trace and by the network residue of intermediate transformations, but not determined by them; and it is also shaped by the coherence network of the subjective domain within which it is carried forwards. The author's intention is one of many prior causal conditions to which the trace testifies but not the origin of meaning itself. Intention and meaning are therefore the productive and inferential orientations of the same trace relation, encountered from opposite temporal directions but never equivalent.

Preprint

Meaning is not determined but remains tightly constrained. The trace preserves the structural consequences of its production: correspondence grounding reaches towards prior causal conditions, coherence embedding positions the trace within a relational network of implications and exclusions. These constraints narrow the space of admissible reconstructions. What counts as a faithful reconstruction is not determined by comparison with the author's intention but by whether pattern persistence is maintained across the transformation [s. 3.2]: the inferential consistency operative at the relevant scale must be preserved even where the internal structure has changed. A reading that loses fine-grained biographical detail while preserving the inferential force of the argument is faithful at the scale of meaning and unfaithful at the scale of material inscription. However, this does not license all or any interpretation. While we cannot anchor meaning in absolute terms, there remain good reasons for preferring one interpretation over another.

Meaning is therefore always locally constituted in reconstruction and continuation rather than extracted from the trace. The coupling conditions of the subjective domain performing the reconstruction, its paradigmatic space, coherence network, and network residue, all shape what admissible reconstructions are available. Different subjective domains will constitute different meanings from the same trace, not because meaning is arbitrary, but because reconstruction is locally conditioned by the structures that make it possible. There is a correspondence with experience: the same text, gesture, image, sound, or object can sustain different meanings for differently situated subjects without thereby becoming unconstrained. There is therefore no original meaning preserved in the trace awaiting recovery. There are original conditions of production, and these include intention, but meaning is constituted only in reconstruction. This expresses the principle of presence-independence [s. 2.4; s. 1.3]: the absence of original conditions does not terminate inferential potential, since absences that are the stabilised consequence of prior causal conditions remain informationally operative within the coupling network that sustains them. The untethering of meaning from intention is a structural consequence of what meaning is and how communication works. In this respect tracework occupies the same structural terrain as Hall’s encoding/decoding model (Hall, 1980), but without retaining a preferred meaning as the privileged reference point for negotiated or oppositional readings.

Because meaning is inferential tracework under subjective regulation, it exhibits the general failure modes of tracework. Retrodictive failure appears where correspondence decay renders causal history unrecoverable, producing underdetermined interpretation. The trace persists but

what it is about becomes increasingly opaque. Remedial failure arises where continuity of inferential consistency cannot be sustained, interrupting the capacity to carry meaning forwards under commitment. Projective failure occurs where forward continuation cannot be stabilised as bounded interpretation, giving rise to proliferating ambiguity or indeterminate extension. Constructive failure appears where transformation of the inferential space produces meanings that were not merely unforeseen but unanticipatable within prior trajectories. Constructive transformation is not a failure in any absolute sense. Construction may become generative, creative, or newly viable once re-stabilised within a later domain. It is a failure only relative to the prior trajectory because the previous structure of admissible continuation could neither anticipate nor integrate the new meaning-space it produced. These are not exceptional cases but structural limits of meaning itself. Meaning remains rigorously constrained by tracework even where it cannot be determined by any single trace, emerging across the whole system in its continuing operation rather than isolated in the sign, token, or symbol.

The collapse of meaning is held at bay by reflexive self-regulation, which selects among admissible reconstructions, commits to particular trajectories, and maintains those commitments under ongoing decay. The subjective domain is what allows meaning to sustain productive continuation beyond local reconstruction [s. 2.5]: the forward-carrying of inferential potential into sustained trajectories that reach back towards causal conditions while continuing forwards under commitment. At larger scales this regulatory function develops external supports in institutional form, such as administrative rationality (Weber, 1922), archives and libraries (Tredinnick, 2026), invented traditions (Hobsbawm and Ranger, 1983). Meaning is maintained against the constant pressure of decay, drift, and fragmentation. Where that maintenance fails, meaning fails with it, not suddenly but progressively, as the conditions that sustain its continuation thin and dissolve.

6.7 Creativity and productive continuation: creativity is encountered in the subjective domain as a particular mode of productive continuation. Creative outputs are generally novel, but not all novel traces are creative, since novelty may arise through other means including noise, error, recombination, or stochastic variation. Nor is creativity reducible to unexpectedness or surprise. Error and construction can be extensionally similar while remaining structurally opposed: both may produce outcomes that were not straightforwardly anticipated, but only construction preserves sufficient coherence to extend the conditions of continuation [s. 5.4]. Within

tracework, creativity therefore occurs where existing trace-structures are transformed in such a way that their inferential potential is extended, displaced, or reorganised while remaining sufficiently coherent to be taken up within ongoing tracework. It depends on the same conditions that make forward continuation possible: transformation [s. 5.3], coherence stabilisation [s. 3.3], and the capacity of a produced trace to alter the space of subsequent tracework [s. 5.4].

Creativity is a crucible for the structure of tracework, drawing together transformation and fidelity [s. 5.3], construction [s. 5.4], coherence stabilisation [s. 3.3], viability domains [s. 2.3], paradigmatic organisation [s. 4.4], recursive depth [s. 6.3], and forward continuation [s. 2.5]. It tests whether the distinction on which tracework depends at each of these levels, between deviation and construction, can hold without collapse when it recurs at the level of output and again at the level of system. Creativity arises in cases where fidelity cannot mean simple preservation, because a transformation may remain faithful to a trace-structure by altering the conditions under which that trace-structure can continue. The transition from selection to choice is therefore central to creativity in the subjective domain. In the subjective domain, selection is internalised as choice: the system regulates which admissible continuations are taken up, suppressed, revised, or carried forwards as its own [s. 6.5]. Creativity depends on this transformation. Novelty becomes creative only where a constructed trace is integrated into the system's continuation in a way that alters what further continuation can become. Creativity is therefore not unconstrained invention, but internally regulated selection among admissible transformations whose products modify the field of future admissibility.

This approach reframes a persistent difficulty in the theoretical literature: differentiating creativity from processes that can look superficially similar but do not alter the conditions of further continuation. Boden's influential distinction between exploratory, combinatory, and transformational creativity (Boden, 1998; 2004) describes what kinds of output creativity produces but does not specify the structural conditions under which a system can be a creative subject. Exploratory creativity operates within a given conceptual space; combinatory creativity recombines existing elements; transformational creativity alters the space of possibility itself. Tracework preserves these distinctions, but grounds them structurally rather than descriptively. Exploratory creativity is productive continuation within an existing stabilisation domain. Combinatory creativity is construction operating through redundancy coupling: recombination that alters the conditions of further production rather than merely rearranging existing elements within them. Transformational creativity is construction at the boundary of the viability domain: the production of causal sequences that are new relative to that domain and that alter what is

subsequently admissible within it. The distinction between these forms of creativity in tracework is not found in creative outputs but in whether the produced trace alters the conditions of further production within the operative stabilisation domain.

Creativity framed in this way connects to what Floridi has called the scandal of deduction (Floridi, 2009; 2010): the apparent paradox that deductive reasoning, whose conclusions are already implicit in its premises, nevertheless appears to produce knowledge. Tracework distinguishes logical containment from trace-productive transformation. Deduction is non-ampliative with respect to entailment, since a valid conclusion adds nothing not already contained in the premises. Yet deduction may still be productive with respect to trace-structure, because it makes what was implicit available as a trace that can be stabilised and used as a condition of further continuation. Deduction transforms the organisation and availability of traces within a reasoning system. What construction produces is no less fully caused than a deductive conclusion, but it is not contained in its antecedent conditions in the way a conclusion is contained in its premises. Its product is a new trace whose admissibility is determined by what the viability domain can sustain, and whose significance lies in whether it changes the field of further continuation.

Turing's response to Lady Lovelace's objection is useful here because it separates determination from anticipation. Turing argued that the assumption that machines cannot surprise us rests on "the assumption that as soon as a fact is presented to a mind all consequences of that fact spring into the mind simultaneously with it" (Turing, 1950). In tracework terms, the fact that an output is determined by prior conditions does not mean that it was projectively available in advance. What surprises is not anything internal to the machine's own production, but the gap between the conditions from which an output is constructed and the limited inferential reach through which those conditions can be anticipated. The constructed trace may be fully caused, but not therefore already legible as a possible continuation.

In the subjective domain, the continuation produced by construction is internal to the system whose trace-structure is transformed. The creative act does not simply result in a creative product, but modifies the system's own field of possible continuation. It feeds back into the self-referential structure of the subjective domain, altering not only what is produced but the conditions under which future production becomes possible. This addresses what Kronfeldner identifies as the exceptionalist and obscurantist attitudes towards creativity: the assumption that creativity depends on exceptional persons or exceptional cognitive processes, and the related assumption that genuine creativity remains inexplicable from a naturalistic point of view

(Kronfeldner, 2018). Tracework shows that creativity need not be treated as inexplicable. It is structurally derivable as construction operating within a reflexively organised system whose outputs alter the conditions of its own further continuation. The sense of inexplicability belongs not to creativity itself but to the aspectual position from which it is encountered from within [s. 6.4].

6.8 Subjective non-self-grounding: the subjective domain has been derived as the form tracework takes when continuation becomes internal to a reflexively self-regulating system. However, a system that regulates its own continuation by operating over its own trace-structure does so under conditions it did not itself produce. Tracework ontologically precedes the subjective domain. The subjective domain is constituted through prior tracework it did not produce and cannot fully reconstruct from within its own regulatory operations. The fold in tracework that constitutes the subjective domain is itself the product of prior causal processes operating outside that fold. The system that emerges from those processes cannot reach back through them from within its own regulatory operations to specify the conditions of its own emergence. Any attempt to do so would itself be a regulatory operation over traces, subject to the same constraints as all other tracework, including the constraint that traces cannot fully recover the causal conditions of their own formation. The subjective domain therefore cannot ground itself. Its regulatory operations presuppose conditions they cannot themselves derive. The subjective non-self-grounding argument anticipates the more general principle given later [cf. s. 8.2], where non-self-grounding at the level of the framework as a whole is derived. At the foundational level, its consequence is the limit of the derivation itself.

Subjective non-self-grounding has a precise consequence. If the subjective domain cannot ground itself from within its own regulatory operations, then the conditions required for its continued viability cannot be supplied entirely from within. The redundancy plasticity of the coupling network through which it sustains inferential potential [s. 3.6], the coherence plasticity through which it maintains inferential consistency across its own trace-structure [s. 3.6], the normative standards through which it evaluates productive continuation, and the paradigmatic space within which it selects admissible continuations all depend on a shared trace infrastructure that extends beyond any individual subjective domain. The intersubjective does not provide an absolute outside to tracework, but differential positions within the same trace infrastructure: positions from which the same causal conditions can leave different traces, from

which correspondence grounding can be restored where it has begun to drift, from which coherence drift can be challenged where it has outrun causal recoverability, and from which the space of admissible continuations can be extended beyond what any single domain can sustain alone.

Intersubjectivity does not guarantee correction. The same shared infrastructure that makes correction possible also makes distortion, contestation, reinforcement, and fragmentation possible. Other subjective domains may restore correspondence grounding, but they may also intensify coherence drift; they may extend the paradigmatic space of continuation, but they may also stabilise incompatible trajectories. What matters is that no subjective domain can supply these possibilities wholly from within itself. The subjective domain requires differential externality as a structural condition of its continued viability. Intersubjectivity is therefore the condition under which subjective non-self-grounding becomes corrigible, extensible, contestable, and capable of sustained continuation. Intersubjectivity is required by the non-self-grounding of the subjective domain.

Preprint

7. The intersubjective domain

7.1 Introduction: the intersubjective domain is the final redescription of tracework available at the human scale. Intersubjectivity in tracework is defined by the coupling of subjective domains, rather than by mutual understanding or consensus. It describes what happens when multiple subjective domains coordinate, contest, or fragment one another's continuations through shared trace infrastructure. The intersubjective domain extends beyond the timescale of any individual subjective domain, because the shared trace infrastructure through which it operates is inherited, maintained, and projected across generations. It reaches back into the historical past and forwards towards anticipated futures. This section outlines the conditions under which intersubjectivity becomes possible.

7.2 Intersubjectivity: the intersubjective domain arises from the non-self-grounding of subjective domains [s. 6.8], but its structure and regulatory role derive from the subjective coupling conditions [s. 6.5]. The minimal threshold for its formation is co-commitment: the coupling of two or more subjective domains through a shared field of continuation in which what one domain carries forwards becomes materially consequential for what another can sustain, contest, revise, or carry forwards. It does not require consensus or even mutual understanding. It requires only that subjective domains are mutually implicated within a shared trace infrastructure, so that what one subjective domain does with a trace can shape what another is subsequently able to do with it. Conflict and disagreement therefore become intersubjective only where co-commitment has already established a shared field of continuation within which disagreement can be recognised as such. Two civilisations that never meet do not share an intersubjective domain merely by virtue of possessing incompatible belief systems. The intersubjective domain remains viable only while divergence remains below the threshold at which coordinated continuation becomes impossible.

Subjectivity requires shared trace infrastructure because it cannot ground itself, but that infrastructure is not a uniform external resource. It is constituted through the prior tracework of other subjective domains accumulated across generations, conditioning what any individual domain can sustain, reconstruct, and carry forwards. No subjective domain operates within a coupling network that is exclusively its own; this parallels Wittgenstein's private language argument (Wittgenstein, 1953), although it is derived from the architecture of tracework rather than from the grammar of rule-following. At the human scale, this shared infrastructure takes the

form of languages, archives, legal systems, institutions, and other symbolic and normative arrangements. These are not merely resources available to subjects; they are inherited trace-structures through which subjective domains become capable of continuation. Shared infrastructures exhibit trace equivalence at the intersubjective scale [s. 3.2]: inferentially, they are trace-structures sedimented from prior collective activity; productively, they function as traces that condition future coordinated continuation. Each such infrastructure is the product of intersubjective tracework, precedes, and cannot be unilaterally revised by any individual domain. This shared infrastructure conditions meaning, reception, and regulation.

The shared trace infrastructure is the material and institutional substance of information work. The information professions, including librarianship, archival practice, documentation, knowledge organisation, and information science, operate on the traces through which subjective domains become intersubjectively extended. Their practices preserve, organise, and reconfigure the trace infrastructure through which meanings remain reconstructable beyond the lifespan and memory of any individual subject. They therefore form part of the regulatory infrastructure of the intersubjective domain. To describe information at the social level is therefore to describe this intersubjective tracework: the collectively sustained organisation of traces through which meaning remains available for further reconstruction and continuation.

Meaning requires continuation within a subjective domain [s. 6.2], but the conditions under which that continuation is possible depend on the shared trace infrastructure through which the domain's coupling network has been formed. A subjective domain confined to its own traces would be informationally impoverished to the point of non-viability. Its correspondence grounding would be narrow, its coherence network thin, and its redundancy insufficient to preserve inferential potential under perturbation. The meanings such a domain could sustain would therefore be correspondingly narrow, held only as far as a single trajectory can carry them without correction, extension, or intersubjective grounding. Intersubjectivity is therefore the scalar architecture of meaning: it supplies the shared trace infrastructure through which individual domains sustain and extend their continuations beyond what any one domain could carry alone. At human scale, every subjective domain inherits a trace infrastructure it did not create and sustains one it will not survive.

Intersubjective tracework is irreducibly divergent. Multiple subjective domains operating over shared trace-structures reconstruct those structures from different positions and under different histories of commitment. Their paradigmatic spaces and normative standards do not coincide,

and there is no position from which the shared trace infrastructure is identically accessible to all domains at once. The intersubjective domain is characterised by permanent underdetermination and ongoing negotiation over what counts as common ground. Alignment is sustained as a dynamic equilibrium rather than achieved once and for all. Shared infrastructures such as archives, libraries, legal systems, and languages keep divergence within tolerable bounds. Intersubjective tracework succeeds not by removing disagreement but by holding it below the threshold at which coordinated continuation becomes impossible.

The intersubjective domain highlights the precise moment at which theories grounded in already-stabilised primitives lose their grip on meaning. Shannon's framework requires a code to specify what counts as a message, but the code itself is not derived within the framework; it is presupposed as already given. Shannon's machinery therefore defines a formal mapping, but not its status as a code. A code is a stabilised regime within which marks, signals, or states count as messages for a domain. It cannot be grounded by the same machinery that presupposes it. Any such grounding would require a further coding regime, and that regime would require another. The code is therefore the site at which semantic content has been quietly deposited before the formal account begins. Qvortrup has observed that both the selection of messages and the code imply subjective selection in Shannon's model (Qvortrup, 1993). This concealed subjectivity is, from the tracework perspective, the intersubjective trace infrastructure that Shannon's framework presupposes without acknowledging. The code is not informationally neutral. It is the stabilised residue of prior intersubjective tracework, comprising inherited standards, paradigmatic relations, and coupling networks formed through accumulated continuation. It carries the meaning and significance that the bit cannot carry because it is embedded within shared trace infrastructure that no individual subjective domain produces alone. Any code is itself embedded in other codes ad infinitum, and there is no foundational code that grounds them all. Theories grounded in already-stabilised primitives therefore cannot reach meaning because the stabilised unit presupposes the intersubjective infrastructure that makes stabilisation possible, and that infrastructure is where meaning has been bracketed.

7.3 Regulation and alignment: the intersubjective domain has a regulatory function that operates through alignment: the coordination of subjective coupling constraints at the scale of multiple domains, together with the reflexive revision of the conditions under which that coordination is maintained. Alignment does not require consensus. It requires only sufficient

coordination across shared trace infrastructure for continuations to remain mutually legible, corrigible, and contestable.

Alignment follows from trace equivalence operating at the intersubjective scale [s. 3.2]. Once subjective domains are coupled through shared trace infrastructure, each domain functions productively as a trace capable of conditioning further continuation and inferentially as a trace-structure whose commitments, selections, and norms are reconstructed from prior tracework. The intersubjective domain is therefore a coarser-grained trace-structure composed of subjective domains, and it inherits the coupling architecture already derived for the subjective domain [s. 6.5]. What changes is not the underlying structure but the scale and conditions of its realisation: subjective coupling becomes intersubjective alignment. Anticipatory coupling becomes the coordination of projected continuations across domains; selective coupling becomes alignment over what is made relevant or salient; normative coupling becomes alignment over shared or intersecting standards; and reflexive coupling becomes the capacity of shared infrastructures to revise the procedures by which these alignments are maintained. The three-plus-one structure is therefore preserved because no additional coupling class is introduced at the intersubjective scale; the existing subjective coupling conditions are re-stabilised across multiple domains. Whether coordination remains possible under conditions of permanent partial misalignment depends on the degree to which subjective domains maintain compatible trajectories in line with intersubjective alignment.

Anticipatory alignment is projection operating at the intersubjective level [s. 5.4]: the structural property of a shared trajectory at a given position, determining what the sequence can become given what it already is, now distributed across multiple subjective domains. It holds where domains can project one another's futures within a shared horizon of admissible continuation, making joint action, coordinated planning, and shared expectations possible. Its failure is a collapse into surprise, mistrust, and strategic opacity, as domains lose the capacity to anticipate one another's trajectories.

Relevance alignment is the intersubjective expression of selective coupling. At the subjective level, relevance follows from selective coupling, where some admissible continuations are chosen and carried forwards rather than others [s. 6.5]. At the intersubjective level, the issue is not which continuation a single domain selects, but which traces, saliences, and alternatives are sufficiently shared to make selection mutually intelligible. Relevance alignment concerns what is salient: it occurs where domains converge sufficiently on what is stabilised as significant within

their respective trajectories. It does not require agreement on truth or value, only enough convergence on what counts as relevant to sustain engagement, and it is what allows domains to communicate across difference. Its failure is subtler than anticipatory collapse, because domains may continue to share traces, environments, and even institutions while organising their inferential spaces around incompatible commitments: co-presence without mutual intelligibility. Indexing provides a practical expression of relevance alignment. It does not simply extract what a document is about, but stabilises particular inferential pathways through which a trace can be taken up within a domain. Index terms make some continuations available, others less visible, and thereby participate in the intersubjective regulation of relevance.

Normative alignment binds where anticipatory and relevance alignment coordinate. It occurs where domains recognise and partially bind themselves to shared standards as extensions or intersections of their existing commitments, sustaining institutions, practices, and collective norms as structures capable of holding divergent trajectories within tolerable limits. Its failure is the fragmentation of norms into mutually incompatible trajectories, each sustaining its own standards of evidence, legitimate authority, and admissible argument, so that conflict and institutional crisis emerge from an excess of incompatible normativity rather than its absence.

Reflexive alignment is the intersubjective expression of reflexive coupling, and operates upon the other three rather than alongside them. It appears where shared trace infrastructures revise the conditions under which anticipatory, relevance, and normative alignment are maintained. Legal appeal, scientific peer review, and constitutional amendment are examples of this reflexive function. They do not merely coordinate existing continuations; they revise the procedures, categories, evidential standards, or institutional forms through which coordination itself remains possible.

The first three forms of alignment are mutually reinforcing and mutually destabilising. Anticipatory alignment depends on sufficient relevance alignment to make others' projections legible; relevance alignment depends on sufficient normative alignment to sustain shared standards of significance; normative alignment depends on sufficient anticipatory alignment to render shared commitments actionable. Because the three are mutually dependent, failure in one intensifies strain in the others, and coordination can remain remarkably robust under ordinary conditions and then fail not by gradual degradation but by sudden threshold collapse, as systems that appear functional cease to be so all at once. Contemporary digital information environments exhibit this kind of cross-dimensional threshold pressure, as algorithmic curation

and the erosion of shared evidential infrastructure push alignment towards the limits of what intersubjective tracework can ordinarily manage (Tredinnick, 2023a; 2025).

7.4 Robustness, correction, and conflict: alignment alone does not explain the relative robustness of the intersubjective domain. That robustness depends on redundancy, network residue, and the distribution of correction across multiple positions. Multiple subjective domains operating over shared trace infrastructure generate overlapping, competing, and mutually corrective reconstructions of the same causal conditions from different positions. Redundancy is constitutive rather than incidental: it is what intersubjectivity structurally produces wherever multiple domains operate over shared traces. The network's redundancy plasticity sustains what no individual domain could sustain alone [s. 3.6], preserving inferential potential against decay, drift, and fragmentation across scales and generations. Network residue accumulates within this shared infrastructure [s. 3.3]: the concentrated deformations of prior intersubjective tracework bias which reconstructions are accessible, which pathways remain viable, and which continuations are most likely to be stabilised across the network as a whole.

Intersubjective alignment also provides corrective functions that individual reflexive coupling cannot supply on its own. It sustains correspondence grounding at the intersubjective level through a process structurally equivalent to retrodiction [s. 5.4], recovering causal reach where it has begun to drift beyond the scope of self-revision alone, even where the relevant history can only be reconstructed counterfactually. It also performs remediation at the intersubjective level [s. 5.4], producing new causal segments that restore the continuity of inferential consistency where individual trajectories have been broken or distorted, drawing on network residue as the template for substitution. Neither function eliminates the underlying epistemic gap: a subjective domain can only know through traces that never fully ground themselves. In combination they distribute epistemic correction across multiple positions, extending abductive reach beyond what any individual trajectory could sustain.

Calder's corrective reappraisal of the Blitz spirit illustrates both functions. The popular narrative of universal stoic defiance during the London Blitz was sustained through coherence within the commemorative traditions of national identity, selectively compressing a complex and contradictory historical experience into a dominant configuration. As Calder notes, the Blitz was named before it occurred, and "was instantaneously and spontaneously mythologised" (1991: 10). Documented evidence of panic, looting, class resentment, and psychological breakdown

remained available in the wider trace infrastructure but was compressed below the operative resolution of the coarser-grained narrative. Calder's social-historical analysis (Calder, 1991), operating from a different intersubjective position and drawing on traces the commemorative tradition had not integrated, performed retrodiction by recovering formation conditions the dominant narrative had rendered difficult to reconstruct, and remediation by restoring the continuity of inferential consistency with the documented experience. The correction depended on traces distributed across positions that no individual trajectory could hold together: it came from the intersubjective infrastructure rather than from self-revision within the commemorative tradition. Yet the epistemic gap persists. Calder's account is itself a reconstruction, and the Blitz spirit narrative continues to circulate within the coherence network it has always sustained.

The corrective power of intersubjectivity therefore depends on disagreement as well as alignment. Multiple domains reconstructing shared traces from different positions do not merely introduce noise into a system that would otherwise be stable. They provide the redundancy through which drift can be detected, correspondence can be restored, and broken continuities can be remediated. The intersubjective domain is robust because no single subjective trajectory exhausts the shared infrastructure. Different domains preserve different pathways, retain different residues, and sustain different possibilities of reconstruction. What appears as disagreement at one level may therefore function as redundancy at another.

Intellectual and ideological conflict is the normal expression of intersubjective tracework under conditions of partial misalignment, and reflects competition between stabilisation domains. When shared trace infrastructure is reconstructed under different coupling conditions, prior commitments, and normative standards, divergent trajectories compete to establish their own correspondence grounding and coherence organisation as the dominant inferential profile within the shared domain. Conflict is this competition surfacing openly. It is neither irrational nor accidental, but the structural consequence of multiple subjective domains operating over shared traces under conditions that cannot produce identical reconstructions. Resolution, where it occurs, is the provisional stabilisation of one trajectory as sufficiently correspondence-grounded, coherence-sustaining, and dominant within the shared coupling network to carry coordinated continuation forwards. Its dominance must still preserve inferential reach, sustain relational compatibility, and hold divergent domains within a common field of contestability. For that reason it is never final, remaining exposed to paradigmatic drift, further perturbation, and renewed competition from alternative stabilisation domains.

Preprint

All of this depends on the maintenance of the shared trace infrastructure within which conflict can remain productive rather than fragmentary. Where that infrastructure degrades, conflict ceases to re-stabilise, as can be seen across contemporary social media platforms. In these circumstances, what remains is the continuation of tracework within increasingly isolated domains, each internally coherent and mutually unintelligible, each sustaining meanings that can no longer be corrected by the perspectives of others.

7.5 Speech, writing, and stabilisation: regulation and alignment are at their most visible in speech, writing, and language. These are not merely forms of communication, but also stabilisation technologies through which shared trace infrastructure is produced, maintained, and carried forwards across subjective domains and generations. Tracework shows why speech and writing are structurally distinct rather than merely materially different. They stabilise traces under different viability conditions, and therefore support different forms of intersubjective regulation. As Poster observes, "each method of preserving and transmitting information profoundly intervenes in the network of relationships that constitute a society" (Poster, 1990: 7). In tracework terms, this is because every stabilisation infrastructure conditions what trajectories remain viable.

Speech in its immediate situational context is syntagmatically dominant. It exists only in enactment and is sustained through ephemeral physical coupling, immediate succession, and the shared conditions under which it takes place. Its correspondence grounding is narrow and its decay rapid, both because the spoken trace remains bound to the contexts within which it occurs and because its material substrate dissipates almost as soon as it is produced. Its inferential force lies instead in immediate coherence stabilisation. What becomes meaningful in speech is not determined by semantic content alone, but by the whole field of enactment in which an utterance is produced and taken up. Tone, gesture, expression, timing, social markers, and phatic connection all help determine what can be reconstructed and carried forwards. Speech is therefore powerful in the moment of production, but once the situation has passed, the spoken trace no longer retains that stabilising field in itself. What remains must be reconstructed through memory, report, recording, institutional consequence, or further tracework.

Writing stabilises differently. It decouples the trace from the viability domain of its production. Correspondence grounding can persist in a durable substrate independent of the producing system; the syntagmatic sequence can be fixed and returned to; the paradigmatic dimension

becomes more accessible; and revision becomes possible outside the constraints of synchronous production. Writing therefore extends inferential reach across absence. It allows traces to remain available for reconstruction after the disappearance of speaker and context. The written trace persists not because meaning has been transmitted into a durable object, but because the conditions for reconstruction have been re-stabilised in a form that can survive the absence of the producing system. Speech mobilises the intersubjective domain, but writing gives that domain sedimentary persistence. However, writing sacrifices the immediacy of speech, and the situational context within which speech is stabilised.

This historically decisive role of writing has not been matched by an adequate theoretical account of writing itself. The integrational linguist Roy Harris has written:

> It is a striking paradox that the Western academic tradition, which has relied so extensively on writing, as opposed to oral transmission, for its very existence, has so far produced no comprehensive theory of writing itself. That is not to say that important questions about writing have not been raised and debated. But they have usually been raised as side-issues relevant to other concerns rather than as a series of coherent questions in their own right (Harris, 1995: 1).

Harris later observed that "we are still lacking a viable semiology of writing [...] The rethinking of writing that was already needed in Saussure's day still remains to be done" (Harris, 2000: 16). A quarter of a century later the point stands. This matters because writing is a distinct stabilisation regime, altering the viability conditions under which traces persist, are compared, and remain available for reconstruction. That includes spoken language itself. Once writing becomes embedded within shared trace infrastructure, it can stabilise language, regulate correctness, fix forms of authority, and shape what becomes permissible within speech. In this sense, within literate cultures, writing can come to precede speech as a regulatory condition of its reconstruction and evaluation (cf. Derrida, 1967). Harris's observation therefore marks the failure to theorise one of the central infrastructures through which intersubjective tracework has historically extended itself.

The valorisation of speech is visible across the philosophy of language. Saussure privileges the spoken sign while assigning writing a secondary status (Saussure, 1916). Austin's speech acts are named for their medium, yet their theoretical force depends heavily on the repeatability of examples stabilised in writing (Austin, 1962). Chomsky's generative grammar abstracts language from particular realisations, treating competence as prior to performance and thereby bracketing

many of the stabilisation conditions through which actual traces become meaningful (Chomsky, 1965). Davidson's radical interpretation presupposes shared linguistic regularities whose persistence depends on repeatable and socially stabilised forms (Davidson, 2005). Wittgenstein's language games are directed against abstraction from use, yet they too are constructed for philosophy as written examples of situated speech (Wittgenstein, 1953). In each case, speech is treated as theoretically primary while the theory itself depends on the stabilising affordances of writing: repeatability, comparison, citation, abstraction, and return. Tracework makes this dependence explicit by treating speech and writing as distinct viability domains rather than as transparent media for a single underlying linguistic object. It also exposes the danger that speech may be reconstructed under the conditions of writing, so that what appears to be a theory of spoken language is already shaped by the stabilisation regime of inscription.

Speech has its own mechanisms for achieving temporal extension independently of writing. Austin's performative utterances, for example, do not transmit content but stabilise traces whose inferential consequences extend forwards beyond the moment of production (Austin, 1962). The spoken utterance is materially ephemeral, but what it produces may persist. Where felicity conditions are satisfied, the utterance becomes a forward constraint on admissible continuation: a promise, oath, or declaration that binds subsequent trajectories even after the sound-event has decayed. Medieval practices of memorialising commitment show how speech could be extended through symbolic and institutional supports before written forms became dominant in many regimes of validation. Clanchy's account of medieval England describes this transitional world, in which memory, oral testimony, ritual, seals, objects, and written instruments formed overlapping stabilisation regimes (Clanchy, 1993). The classical and medieval arts of memory likewise show that memory was a trained infrastructure for preserving, ordering, and retrieving traces in the absence of durable inscription (Yates, 1966; Carruthers, 2008). The study of oral tradition shows how formulaic composition, mnemonic patterning, rhythm, repetition, and conventional narrative structures allowed speech to acquire durability through highly stabilised forms of recurrence before and alongside writing (Parry, 1928; 1930; Ong, 1982; Havelock, 1986). These mechanisms extend speech through other forms of tracework. They provide supplementary coupling structures through which ephemeral utterance can acquire durable inferential consequences. Where they hold, the spoken trace becomes durable within the network of social, symbolic, mnemonic, and institutional tracework, binding later continuation under commitment.

Preprint

Writing extends temporal reach by different means. It strengthens the capacity of traces to survive absence, displacement, and return. It allows prior commitments to be revisited, compared, cited, contested, preserved, copied, and recontextualised across domains and generations. Its regulatory power lies in the stabilisation of retrospective reach. Administrative rationality depends on writing because writing can preserve traces beyond the immediate continuity of speech. This makes writing a primary technology of normative alignment. It anchors standards against drift, holds commitments beyond the moment of utterance, and makes prior conditions available for later reconstruction. Speech regulates through situated anticipation and immediate commitment; writing regulates through durable correspondence grounding and extended normative reach. Relevance alignment requires both: the situational salience and responsiveness that speech supplies, and the persistence, comparison, and return that writing makes possible.

The history of writing is therefore not merely a history of communication technology. It is a history of changing stabilisation regimes. Ong has argued that literacy transforms the structures of thought (1982); however, tracework suggests that a different mechanism is at play. Manuscript, print, and digital writing each define distinct viability domains (cf. Tredinnick, 2008). Manuscript stabilisation depends heavily on custodial continuity, material singularity, and path-dependent reproduction. Inferential consistency is difficult to sustain across changes of physical instantiation without institutional support. Print intensifies repeatability and standardisation. It increases correspondence plasticity across dispersed domains because multiple physical instances can sustain compatible inferential potential at the same time. Digital writing expands searchability, recombination, circulation, and scale, but often loosens the relation between trace, substrate, provenance, and shared context. It can increase coherence reach while weakening correspondence grounding to the conditions of original production. Each transition alters the balance of correspondence and coherence, and therefore changes the forms of anticipatory, relevance, and normative alignment that writing can sustain. What appear as historically situated modes of thought are historically situated modes of information.

Speech and writing differ not in the sense that one is primary and the other subordinate, nor in the sense that one is natural and the other technological. Their difference lies in the viability domains through which each sustains inferential potential and productive continuation. Speech and writing are therefore not rival expressions of language but distinct stabilisation regimes within intersubjective tracework. Together, they constitute language as a system. In combination they become a powerful and pervasive expression of the intersubjective domain.

7.6 Meaning and ambiguity: the shared trace infrastructure of intersubjective meaning is never perfectly aligned. Communication therefore takes place within a space of admissible reconstruction, where correspondence grounding and coherence organisation jointly constrain what can be meant without determining it in advance. The potentiality of intersubjective ambiguity is realised whenever traces are reconstructed under the coupling conditions of a particular communicative event. Ordinary communication often conceals this structure because its conventions narrow the admissible field quickly enough for reconstruction to appear immediate. Figures of speech make the operation visible. They delay closure, multiply admissible pathways, or shift reconstruction from correspondence-dominant to coherence-dominant stabilisation. In ambiguity, meanings multiply from a single form; in malapropism, meaning is constituted despite the wrong form; in irony meaning works against the surface form; in puns meaning emerges twice from the same form. If meaning were intrinsic to linguistic form, what is stated should be the end of the matter.

Under tracework these cases dissolve cleanly. Ambiguity is not a failure of linguistic encoding but a structural feature of the correspondence-coherence plane: it is the condition under which any meaning can emerge at all, resolved when additional coupling constraints narrow the space of admissible reconstructions through prosodic, syntactic, institutional, or situational context. Malapropism demonstrates redundancy coupling in action: the coherence network distributes inferential potential across multiple relations, none of which depends on the lexical form alone, allowing meaning to be reconstructed despite correspondence failure at the level of the word. Redundancy plasticity absorbs the correspondence failure [s. 3.6]: no new causal segment is required because sufficient inferential potential is distributed across the coupling network to sustain reconstruction despite the local disruption. Puns demonstrate the paradigmatic dimension: the same position in the syntagmatic chain sustains multiple distinct inferential relations concurrently, each drawing on different regions of the coherence network under shared formal conditions. Irony is paradoxical only under transitive assumptions; in tracework it is entirely expected because the coherence network is an independent axis of stabilisation that can override surface correspondence where the surrounding relational structure is sufficiently constraining. It is normal linguistic function made transparent.

Metaphor extends the same architecture beyond local figures of speech into transformation itself. Lakoff and Johnson argue that:

> Our concepts structure what we perceive, how we get around the world, and how we relate to other people. Our conceptual system thus plays a central role in defining our everyday realities. If we are right in suggesting that our conceptual system is largely metaphorical, then the way we think, what we experience, and what we do every day is very much a matter of metaphor (Lakoff and Johnson, 1980: 3).

Erickson has noted that spatial metaphors 'can have remarkable effects on our perceptions' (2012: 912), and Boroditsky that such metaphors 'provide relational structure to an abstract domain' (2000: 1). Tracework suggests that metaphor has transformative power over the shape of discourse. It is not merely decorative but constitutive and restructuring, and therefore part of what successful communication requires. Metaphor is a reframing operation: a transformation along the coherence axis that generates new inferential pathways not previously accessible within the trace-structure. It does not merely select among admissible reconstructions; it alters the field within which future reconstructions become admissible. It is construction operating at the semantic level [s. 5.4]: producing new causal sequences admissible under the surrounding coherence domain without being guided by an existing inferential template, occupying the boundary between viable innovation and collapse. The new coherence relations it constructs condition subsequent trajectories in ways that cannot simply be undone. Meaning is never intrinsic but always reconstructed within the coupling conditions of the system that stabilises it as meaningful. Metaphor works not despite the mismatch between form and sense but because the reframing it performs generates new coherence relations that can be taken forwards as meaning.

Figures of speech therefore do not sit at the margins of meaning. They disclose the normal architecture of communicative reconstruction. What is usually called literal meaning is not a default state from which ambiguity, irony, malapropism, pun, and metaphor depart, but a stabilisation effect: coupling constraints close the space of admissible reconstruction quickly enough for meaning to appear contained in linguistic form. Where closure is disrupted the hidden structure becomes visible. Communication always proceeds through a field of possible reconstructions shaped by correspondence grounding and coherence organisation; figures of speech reveal how that field is transformed. They are not deviations from meaning but demonstrations of the conditions under which meaning becomes possible at all. This explains why the same architecture that makes communicative meaning possible makes it vulnerable. If coherence can repair malapropism, sustain irony, and generate metaphor, it can also stabilise reconstructions whose correspondence grounding has weakened or failed. The productive

Preprint

ambiguity of communication and the pathology of misinformation are therefore not unrelated phenomena. They arise from the same correspondence-coherence architecture under different conditions of regulation, commitment, and decay.

7.7 Misinformation, disinformation, and phatic coupling: misinformation and disinformation show the same architecture operating at the level of informational pathology. Both are possible because meaning is reconstructed under the conditions of the subjective domain rather than transmitted as content. Were meaning contained in the trace, falsehood would be a simple defect of correspondence: the trace would either match its object or fail to do so. But under tracework, meaning arises from the actualisation of inferential potential within a subjective domain, and that actualisation is governed by the relation between correspondence and coherence. A trace may therefore remain meaningful even where its correspondence grounding has weakened, drifted, or been deliberately severed. Coherence can sustain meaning beyond the point at which causal recoverability remains secure.

Misinformation is a trace or trace-structure that sustains coherence-dominant organisation after its correspondence grounding has decayed below the level required for reliable reconstruction of prior causal conditions. The inferential potential remains active and the meaning is genuine, but the trace no longer testifies adequately to the causes it purports to disclose. The regulatory actualisation holds because the receiving coherence network can still stabilise the trace as meaningful: it fits, explains, confirms, affiliates, or extends existing commitments. Misinformation is therefore not simply failed information, nor merely false content. It is information whose coherence has become detached from sufficient correspondence. It remains meaningful because it can still be carried forwards within a subjective or intersubjective domain; it is misinformational because the continuation it supports no longer remains adequately constrained by causal history.

It is for this reason that misinformation is epistemically sticky. Remediation fails [s. 5.4] not because the receiving subject is simply irrational, but because the misinformation has already entered a trajectory of internal continuation: it has become bound to prior commitments, reinforced by network residue, stabilised through normative coupling, and made salient within a community in ways that the corrective trace cannot displace. A corrective trace that restores correspondence may nevertheless fail to become meaningful if it cannot be carried forwards within the subjective domain's existing commitments. The correction is informational; it may

even be correspondence-rich. But if it cannot enter the coherence structure through which the subject regulates continuation, it remains inert. Misinformation persists because it is not merely believed; it is taken forwards.

Disinformation is structurally distinct. It is the deliberate construction of traces designed to induce coherence-dominant drift in reconstruction while the producing system maintains a different regulatory actualisation. Disinformation exploits the gap between correspondence and coherence. It is construction deployed strategically [s. 5.4]: producing traces whose inferential potential is structured to be actualised in one way by the receiver while being grounded in another trajectory of production by the sender, occupying the admissible space of the receiving domain while corrupting its inferential line. The disinforming system therefore operates productively over the reconstructive system's coherence conditions, constructing traces that are admissible within that domain while weakening, obscuring, or redirecting its correspondence grounding. Disinformation is not only a false claim but a strategic intervention in another system's continuation.

Disinformation is not reducible to lying. A lie still assumes a relatively direct relation between assertion, belief, and truth: the speaker asserts what they know to be false, intending the hearer to believe it. Disinformation may include lies, but its structure is broader. It need not aim only at false belief. It may aim at confusion, exhaustion, distrust, or the saturation of a domain with incompatible reconstructions. Its target is not always a proposition but the coupling architecture of the subjective domain: what can be trusted, what can be taken forwards, which sources remain admissible, which institutions retain coherence, and which continuations remain binding. In tracework terms, disinformation attacks the conditions under which correspondence can be restored.

Misinformation also has a phatic function that content-based approaches generally miss (Tredinnick, 2023a). Its circulation frequently functions not primarily as a truth claim but as a coherence coupling event within a social network: an act of affiliation, recognition, provocation, or boundary maintenance. The informational significance of the shared trace lies less in what it corresponds to than in what it does to the coupling structure. Sharing a misleading trace may stabilise group membership, signal loyalty, intensify antagonism, or reproduce a coherence network in which certain continuations become obligatory and others inadmissible. The content is the vehicle; the coupling is the dominant meaning. Correction at the level of content often fails

Preprint

because it addresses the correspondence claim while leaving untouched the coherence function that made the trace meaningful in the first place.

The phatic dimension also explains why misinformation can persist even when its propositional content is weak, unstable, or repeatedly contradicted. If the trace functions primarily as a coupling event, its correspondence failure may be secondary to its role in sustaining social coherence. What matters is not only whether the trace accurately testifies to prior causal conditions, but whether it can be taken forwards as a sign of belonging. Under these conditions, debunking may even strengthen the trace's coherence function by confirming the subjective domain's prior commitments about hostility, exclusion, or institutional unreliability. The corrective trace is then reconstructed not as restoration of correspondence but as further evidence within the existing coherence network.

Misinformation and disinformation therefore reveal the full force of the correspondence-coherence architecture under subjective and intersubjective regulation. Rather than being non-informational, they expose the precise boundaries of information's role within the intersubjective domain. Meaning can be genuine while correspondence is weak. Coherence can stabilise traces that no longer remain adequately grounded. Network residue can bias which reconstructions become available. Prior commitments can determine which traces are taken forwards. Meaning does not always function primarily in relation to truth; it also functions in relation to argument, standing, affiliation, and obligation. For this reason, traces that fail as correspondence may still succeed as instruments of justification, persuasion, or group maintenance. The same structural point is reached from the cognitive side by Mercier and Sperber, on whose account reasoning is organised for justification and persuasion rather than for tracking truth (Mercier and Sperber, 2017); a faculty of this kind would resist a correspondence-restoring correction not through malfunction but because correspondence was never its primary function. Communication can therefore fail even where accurate information is available, because meaning is not transmitted but locally reconstructed within domains of internal continuation. These phenomena are not anomalies at the edge of information theory. They are what meaning looks like when inferential potential remains active under conditions of correspondence decay, coherence dominance, and contested continuation.

7.8 Pathologies of regulated continuation: because the subjective domain has been situated within the intersubjective domain, the systemic constraints on subjective regulation can now be

stated. Like every level of organisation in tracework, subjective and intersubjective regulation have limits. These limits mark the conditions under which the subjective domain sustained through shared trace infrastructure can continue to regulate its own commitments under internal continuation, intersubjective correction, and ongoing transformation. Their failure produces identifiable pathologies of regulation: collapse, distortion, or drift in the capacity of a subjective domain to carry forwards its own commitments.

Because the principles of tracework form a tied system that bifurcates along the temporal line, subjective regulation remains viable only while the inferential and productive orientations are held together within a single trajectory. The constraints mark characteristic failures in holding those orientations together. One orientation dominates, loses operative force, collapses into the other, or generates demands the other cannot reconcile. Each constraint is anchored in one of the foundational principles [s. 1.3], but none is reducible to the failure of that principle in isolation.

The principles do not change at the subjective level. They are encountered from within, under the orientation of internal continuation. Not every principle generates a distinct pathology. Some remain background conditions of the subjective domain. Orientational non-equivalence supplies the bifurcation of productive and inferential orientations. Informational intransitivity prevents meaning from being transmitted as content. Presence-independence allows absence to remain operative within subjective reconstruction. Trace equivalence permits the same configuration to function across scales as trace, trace-structure, or trajectory. The five constraints considered below arise where particular principles become unstable under reflexive continuation. Temporal extension becomes the problem of persistence under changed conditions. Reflexive articulation becomes the problem of self-revision without regress. Epistemic priority inversion becomes the problem of reconstruction acquiring the force of obligation. Informational grounding becomes the problem of incompatible commitments. Finally, epistemic asymmetry becomes the problem of forward continuation ceasing to answer to causal history.

The first constraint concerns the viability of coherence. Rigidity is a failure in the tension between the two orientations of temporal extension. The backward inferential orientation remains operative: the system can preserve and reconstruct its causal history and maintain inferential consistency with prior tracework. But the forward productive orientation loses recursive plasticity: the system can no longer revise, reorganise, or carry itself forwards under changed conditions while preserving the continuity of its own commitments. The result is a subjective

domain that remains grounded but cannot continue adaptively, because temporal extension has become asymmetrically weighted towards backward grounding at the cost of forward continuation. Aesop's bow gives the image. An unstrung bow cannot act; a bow held always under full tension weakens and breaks. The point is not to maximise coherence but to maintain the tension between holding and releasing that makes continuation possible.

The second constraint concerns the limits of reflexivity. Reflexive regress is the backward inferential orientation of reflexive articulation operating without sufficient productive closure. The system can examine its own operations, reconstruct its own commitments, and reopen its own conditions, but it cannot bind itself to a forward trajectory. Every attempt at self-articulation becomes another object of examination rather than a basis for action. The result is a capacity for self-examination that has lost the capacity for enacted continuation. The centipede gives the image. It walks effortlessly until asked how it coordinates its hundred legs, at which point the attempt to explain arrests the possibility of walking. Action requires not ignorance of regulation but the suspension of complete self-explication at the point of continuation.

The third constraint concerns temporal governance. Retrospective capture is a collapse in the asymmetry between the two orientations of epistemic priority inversion. The backward inferential orientation acquires the force of forward obligation: what should function as reconstructed prior condition becomes binding obligation. The past ceases to operate as inferential grounding and begins to govern continuation directly. The result is not memory, tradition, or accountability, but a distortion in which reconstruction usurps production. Orpheus gives the image. He is permitted to lead Eurydice out of the underworld on one condition: he must not look back. He walks forwards under commitment, taking the past as inferential resource rather than binding obligation. The turning back destroys the path itself. Forward continuation depends on keeping the backward orientation in its proper relation to production.

The fourth constraint concerns normative load. Normative overload is the failure of the tied system to reconcile multiple forward obligations within a backward inferential structure capable of ordering them. The forward productive orientation of multiple commitments generates incompatible demands on continuation, while the backward inferential orientation cannot reconstruct a coherent trajectory through which those demands can be jointly stabilised. The system does not lack normativity; it is overloaded by normativity that can no longer be integrated into viable continuation. Dorigen in *The Franklin's Tale* gives the image. She is bound both by her commitment to her husband and by her word to Aurelius once the impossible condition she had

set him is apparently fulfilled. Neither obligation is false or externally imposed; both are sustained within her own trajectory. Yet they pull continuation in incompatible directions. She does not lack a normative framework. She is overloaded by one. It is Arveragus who releases one burden and restores compatibility from outside the overloaded domain. A viable subjective domain must sustain not only commitments but their compatibility as a condition of action.

A fifth constraint is the mirror of the first. Confabulatory drift is a failure in the tension between the two orientations of epistemic asymmetry. Here the forward productive orientation remains operative, even vigorous: the system continues to commit, act, and carry itself forwards. But the backward inferential orientation loses its grounding, degrading to coherence alone, so that reconstruction returns whatever fits the forward trajectory rather than what causal history will bear. The result is a domain that continues fluently while ceasing to be answerable to what produced it: productive continuation outrunning correspondence, the inner correlate of the coherence-dominant drift that produces misinformation [s. 7.7]. It is not that the system has lost its past; it is that the forward orientation has begun to write it, fabricating through internal remediation and construction [s. 5.4] a history that backward inference then reconstructs as though it were original. This is distinct from the mere absence of grounding, which does not disorder the subjective domain but dissolves it, dropping the system below the threshold of subjectivity altogether [s. 6.3]. Drift is the disorder of a domain that stays constituted while its grounding thins. Don Quixote gives the image: a fully constituted subject whose reconstruction has gone coherence-dominant. Windmills are smoothly resolved into giants by a chivalric frame robust enough to sustain action long after correspondence has drifted away. The point is that coherence can carry a subject forwards on its own, and a viable domain must keep forward continuation answerable to backward grounding rather than letting it supply its own.

None of these five constraints introduces a new kind of failure beyond the four trajectory-level modes derived later [cf. s. 8.3]. Each is the same threshold-crossing redescribed under reflexive self-regulation. Rigidity and reflexive regress provide sufficient conditions for remedial and projective failure respectively: the first a continuity that cannot be reorganised under changed conditions, the second a continuation that cannot close into action. Confabulatory drift is retrodictive failure occurring from within: correspondence decay that the system's own forward continuation conceals by manufacturing the history reconstruction then treats as given. Retrospective capture and normative overload are configurations rather than single modes: the first a retrodictive operation usurping a projective one, the second multiple forward trajectories failing jointly to close into one viable continuation. The systemic constraints are therefore not a

fifth category of failure standing alongside stabilisation, coupling, and transformation, but the same four modes encountered from within a system answerable to its own continuation.

These systemic constraints delimit the conditions under which the subjective domain can sustain reflexive self-regulation without distortion. To sustain commitment across time requires the maintenance of backward inferential grounding and forward productive continuation: enough coherence to remain oriented without hardening into rigidity; enough reflexivity to permit self-correction without collapsing into regress; enough temporal asymmetry to let the past ground continuation without capturing it; enough normative compatibility to bind action without overloading the trajectory; and enough correspondence grounding to prevent forward continuation from rewriting the history to which it should remain answerable. Where these conditions fail, subjectivity does not disappear, but its regulation becomes structurally disordered.

These pathologies emerge in ordinary distortions of regulated continuation. Rigidity appears in institutions that preserve inherited procedures after the conditions that made them viable have changed. The organisation remains coherent with its past but loses the capacity to adapt. Reflexive regress can appear in data-driven organisations where dashboards, KPIs, audits, and analytics proliferate until decision-making is displaced by the continual examination of the conditions for decision. The organisation can inspect itself ever more finely, but cannot convert that self-inspection into adaptive continuation, and data becomes a basis for deferral, self-justification, and defensive self-description (Tredinnick, 2023b). Retrospective capture appears where tradition, precedent, trauma, or origin narrative ceases to ground continuation and begins to command it, binding the present to a reconstructed past as though no alternative future were admissible. Normative overload appears where legal, professional, familial, or political commitments remain individually binding but become collectively incompatible, leaving action suspended between obligations that cannot be jointly stabilised. Confabulatory drift appears where a person, institution, or community continues fluently while revising its own history to fit its current trajectory, preserving coherence at the cost of correspondence. These cases differ in content, scale, and social form, but their structure is the same: subjective or intersubjective regulation remains active while the relation between backward grounding and forward continuation becomes distorted. Stat.

Preprint

8. Grounding and failure modes

8.1 Introduction: this section closes the foundational account. It derives the principle of non-self-grounding [s. 8.2], specifies the four irreducible failure modes that constitute the falsifiability condition of tracework [s. 8.3], and tests tracework against the boundaries of a foundational theory [s. 8.4].

8.2 The principle of non-self-grounding: tracework begins with the single assumption that information exists, and derives everything that follows as a consequence. But that assumption itself cannot be grounded from within the framework. Any attempt to specify the conditions under which information exists would itself presuppose information's existence and take the form of information, reproducing the assumption rather than escaping it. This is a structural consequence of proceeding from any single foundational assumption. A framework cannot simultaneously occupy the position of its own foundation and the conditions that precede it, without collapsing the very asymmetry that makes derivation possible in the first place. Tracework therefore closes as a structural account. Within the space it defines, nothing further is required. But it does not, and cannot, close as a self-grounding system. Why there is information rather than nothing lies outside what tracework can specify. This is the principle of non-self-grounding: not an addition to the framework, but its final derivation, expressing at the level of the whole the same asymmetry that governs every step within it.

8.3 Failure modes: tracework yields a precise and restrictive account of where and how trajectory viability can fail. Each failure mode is the breakdown of one trajectory operation, occurring where decay along one or more structural axes crosses the threshold at which stabilisation can no longer be maintained. Failure is therefore not an independent phenomenon but the limiting case of stabilisation under sustained transformation: not the absence of structure, but the inability to sustain a trajectory within the structure that exists. The space of possible failures is correspondingly restricted. Because decay is structured and constrained by the core conditions of trace persistence, correspondence, coherence, differentiation, persistence, and productive capacity, transformation space does not admit arbitrary breakdown. It admits only the failure of trajectories to remain viable under these constraints.

The number of irreducible failure modes is therefore strictly bounded. Each corresponds not to a primitive breakdown but to a typical configuration of decay across core structural axes.

1. Retrodictive failure arises where reconstruction cannot adequately recover the admissible causal conditions of a trace. This may occur because correspondence decay erodes recoverability, leaving multiple causal histories equally admissible, or because reconstruction prematurely commits to a particular causal history despite residual underdetermination. The defining condition is therefore a failure of retrodictive adequacy: either the trace no longer sufficiently constrains its origin, or the admissible space of origins is incompletely represented.

2. Remedial failure arises where persistence decay crosses the threshold of continuity, typically combined with coherence disruption and insufficient correspondence to support reconstruction. Forward trajectories require that successive states remain linked under shared constraints. Where this continuity is broken and cannot be re-established through backward reconstruction, the trajectory cannot be sustained. The defining condition is the absence of any viable continuation consistent with both correspondence and coherence constraints. It marks the limit of forward extension where continuity itself cannot be recovered.

3. Projective failure arises where productive decay or constraint insufficiency prevents the closure of forward trajectories. Forward propagation requires that the space of possible continuations be sufficiently constrained to form a stabilisable path. Where the generative space expands beyond the capacity of available constraints to stabilise it, trajectories fail to close. The defining condition is non-closure: continuation remains possible but cannot be bounded into a coherent trajectory.

4. Constructive failure arises where forward production alters the space of admissible continuations itself. Trajectories presuppose a bounded structure of possibility defined by prior constraints. Where transformation produces structures that redefine those constraints, previously valid anticipations no longer apply. The defining condition is endogenous expansion of the space of possibility. Novelty here is not simply unexpected but cannot be anticipated relative to prior trajectory structure.

These four modes describe failures of trajectory viability within a stabilisation domain. They do not describe the case in which trajectories succeed under conditions of systematic bias internal to the domain itself. An archive that systematically excludes a category of voices may remain correspondence-grounded, coherence-stable, and projectionally closed within its own terms

while sustaining systematically distorted reconstruction. Such distortion is not a trajectory failure but a property of the domain's constitution, recognisable only from a differential position within shared trace infrastructure. It is therefore an intersubjective phenomenon, addressed through the corrective functions of alignment [s. 7.3] and the presence-independent informativeness of stabilised absence [s. 2.4], rather than a fifth failure mode.

These four modes are jointly exhaustive. No additional independent failure mode is possible without introducing new structural conditions of tracehood, new temporal orientations, or relaxing the requirement that trajectories preserve inferential and productive viability. All apparent complexity, hybrid failures, domain-specific breakdowns, gradual or cascading collapse, can be analysed as multi-axis trajectories across these dimensions. The space of failure is closed, but internally articulated.

Tracework is falsified if there exists a well-formed case in which trajectory viability fails that cannot be expressed as threshold-crossing along one or more of these structural axes, or as a configuration reducible to the canonical forms derived from them. Such a case would imply either an additional irreducible condition of trace persistence or a failure in the constraint structure defining transformation space itself. In the absence of such a case, all limits on inference, reconstruction, and continuation must fall within the bounded space defined by stabilisation, decay, and their interaction under temporal asymmetry.

8.4 The boundaries of a foundational theory: section 1.3 set out the boundaries within which a foundational theory of information would need to operate, drawn from a prior thematic review (Tredinnick, 2025) and those from Floridi's open problems in the philosophy of information not already implied by that review (Floridi, 2004a; cf. Crnkovic and Hofkirchner, 2011). That mapping distinguished what was known and generally agreed, what remained unknown and contested, and what appeared currently intractable, setting a deliberately high bar for what a foundational theory should be expected to resolve or show to be irresolvable. These questions were not chosen to suit tracework. They catalogue the field's own points of agreement, dispute, and intractability as independently identified in the prior literature, established before this paper's derivation and without reference to it. Nor was tracework shaped to answer them: it proceeds by derivation from the single assumption that information exists, under a triple lock that permits no step introduced merely to secure a particular result [s. 1.3]. What follows is therefore a test of scope against an independent standard. Tracework meets that standard in the following ways:

- **The precise nature of information (15).** This is the paper's central claim, and everything above follows from it. Information is the persistent metaphysical trace of prior causation, the structured residue of causal processes that sustains the possibility of reconstructing admissible prior conditions [s. 1.1], where reconstruction is understood as continuation. It is neither primitive nor transmitted, neither a static unit nor an abstract entity, but a trajectory constituted under temporal asymmetry, testifying to its own causation under continuation [s. 2.5]. Every other item in this list is a consequence of this single specification.

- **Whether a grand unified theory of information is possible (19).** Tracework shows that a unified theory of information is possible if unification is understood structurally rather than reductively. It does not reduce information to data, signal, entropy, meaning, computation, or communication. Instead, it derives the minimal conditions under which any of these can function informationally. The result is a foundational architecture within which diverse informational phenomena can be located, related, and delimited.

- **Information is a real phenomenon with material instantiation rather than wholly abstract or ideal (1).** This is derived, not merely assumed: a difference, whether presence or absence, can only be realised within a material substrate [s. 2.4], and the four intrinsic limits of trace stabilisation are themselves physical limits, corroborated by Bekenstein's bound and Landauer's principle [s. 2.6; s. 2.8]. Information is not an abstract entity standing free of physical constraint; it is inevitably materially inscribed.

- **Whether information's material instantiation matters on an epistemic level (14).** It matters, but not uniformly: its epistemic significance is scale-dependent. At finer-grained scales, where stabilisation is correspondence-dominant with respect to detailed causal history, material instantiation constrains backward inference tightly. At coarser scales, fine-grained correspondence is compressed, and inferential consistency depends increasingly on coherence among the relations that survive. This does not mean that coarser structures cannot be correspondence-dominant; rather, correspondence is re-anchored at the scale of the compressed causal profile. Material instantiation therefore matters epistemically in a way tracework specifies precisely rather than leaving it a matter of degree [s. 4.2; s. 4.3].

Preprint

- **What difference is, and the decisive role it plays in the nature of information (4, 11).** Difference is non-identity under transformation, distinguishable within a stabilisation domain [s. 2.3]; it is the minimal necessary condition of information, but not sufficient on its own. Bateson's aphorism captures the necessity of differentiation without specifying the conditions under which it is non-arbitrary, persists, or is sustained; tracework derives these further conditions as the inferential and productive stabilising requirements [s. 2.2; s. 2.4; s. 2.5].

- **The role of time in information (13).** Time is the dimension within which information exists and is constituted. Temporal asymmetry is the condition under which traces exist [s. 2.6].

- **Whether information is a thing, a structure, or a process (8).** Information is a trajectory: an ordered sequence of transformations constituted under temporal asymmetry [s. 5.4]. This places information closest to process, but not understood as mere activity or flow. A trajectory is process constrained by persistence, transformation, and inferential recoverability.

- **Whether information is intrinsically or situationally meaningful (9).** Meaning is not intrinsic content contained in the trace, nor significance imposed from outside. It is the actualisation of inferential potential within a system capable of carrying that actualisation forwards, locally constituted in reconstruction rather than extracted from the trace or transmitted by it [s. 6.2; s. 6.6].

- **How meaning arises in informational things (16).** Meaning is what reflexive self-regulation makes of inferential potential. It arises where inferential potential is actualised within a subjective domain and carried forwards under internal continuation [s. 6].

- **The apparent incommensurability of information's objective and subjective characteristics (5, 12, 17).** Orientational non-equivalence specifies that the inside and outside of generative tracework are structurally different, incommensurable descriptions of a single process. The apparent incommensurability is not a failure of explanation but its content [s. 6.4].

Preprint

- **The relationship between computation and information (18), developed further in Part Two.** Computation is tracework that is syntagmatically bound: causal constraint is so tight, and ordering so fully determinate, that paradigmatic structure, the space of admissible alternatives a system could itself draw on and revise, is foreclosed rather than merely narrow [s. 4.4]. A computational system is therefore neither excluded from subjectivity by definition nor granted subjectivity by complexity of output alone. The relevant structural question is whether an internal paradigmatic dimension is present and available for reflexive continuation.

- **Why Shannon's account of information has decisive reach (7).** This is explained by the principle of Shannon availability equivalence: any theory that grounds itself in already-stabilised elementary units, of which Shannon's is the purest case, inherits the bit at the level of formal informational availability regardless of the theoretical vocabulary subsequently added to it. The reach is not contingent influence but a structural lower bound every already-stabilised-unit theory shares [s. 2.2].

- **Whether data, information, and knowledge are distinct, related, or identical phenomena (10), developed fully only in Part Two.** What ordinary usage calls data is the assumption of an already-stabilised unit, the very thing tracework's derivation explains the stabilisation of rather than presupposes [s. 2.2]; the trace is the properly derived replacement for that assumption. From there, information, meaning, and knowledge follow as successive reorganisations of the same architecture under increasingly complex conditions of persistence, uptake, regulation, and commitment, not as separate kinds of thing [s. 2.9; s. 6.5].

- **The relationship between information and knowledge or knowing (3), developed fully only in Part Two.** Information is the trace prior to meaning; meaning is inferential potential actualised within a subjective domain; knowledge is meaning that has passed into the system's normative and selective coupling, where selection becomes choice and is stabilised as binding on further continuation rather than merely entertained [s. 6.5].

- **Information's objective characteristics are independent of observation (6), including the status of environmental information.** These are both what the trace is prior to

uptake within a subjective domain. The trace's inferential potential, and the causal history it preserves, are fixed by the conditions of its formation, not by whatever later comes to observe or measure it. Growth rings, geological strata, and eroded inscriptions are informational, and have the determinate structure they have, whether or not anyone reads them. Tracework derives this rather than treating environmental information as a special case [s. 2.4; s. 6.2].

- **The claim that information carries, transmits, or conveys meaning (2).** Informational intransitivity shows that this claim is false in the form given. What appears as transmission is the preservation of inferential potential across transformation. However, the implicit assumption that information is bound up with meaning is preserved: information becomes meaningful where inferential potential is actualised and carried forwards within a subjective domain [s. 1.3; s. 2.6; s. 6].

- **The problem of false information (20).** Floridi's veridical principle treats truth as a condition of informational status, making "false information" a contradiction in terms that ordinary usage nonetheless tolerates without difficulty. Tracework imposes no such condition: meaning is sustained by the relation between correspondence and coherence, and a trace may remain meaningful even where correspondence has weakened, drifted, or been severed. In misinformation, for example, coherence sustains meaning beyond the point at which correspondence secures it [s. 7.7]. The category is not paradoxical; it is what the architecture predicts.

- **Levels of abstraction (21).** Floridi treats the level of abstraction as a methodological choice made by an observer constructing a model (Floridi, 2013). Tracework derives why levels of description exist and how they relate. Coarse-graining is a stabilising operation required by the conditions of generative tracework itself, and the correspondence-coherence balance shifts systematically, not arbitrarily, as scale coarsens [s. 4.2; s. 4.3]. Levels of abstraction are not stipulated; they are derived consequences of what persistence requires at a given resolution.

- **Whether natural intelligence can be implemented non-biologically (23).** Tracework does not decide which systems instantiate the subjective domain. It specifies the structural threshold such systems would need to satisfy and refuses both automatic

inclusion and automatic exclusion [s. 6.3]. Therefore this question remains open. However, tracework does make a strong prediction: systems whose operation remains syntagmatically bound cannot instantiate natural intelligence in the relevant sense, because their admissible continuations are determined by ordered causal constraint without an internally available paradigmatic dimension over which they can exercise reflexive self-regulation. Current computational approaches are therefore excluded only insofar as they remain syntagmatically bound. A non-biological system would become a candidate for natural intelligence only where it sustained internal continuation, recursive depth, coherence plasticity, and a paradigmatic field of admissible alternatives available for choice, commitment, and self-revision.

- **The value of information and the proper scope of information ethics (22, 24).** These remain open. They are not addressed by tracework and could not be derived from the triple lock as it stands, since that admits no commitments beyond those required to describe tracework itself. Questions of value and ethical standing require normative commitments that cannot be derived from the single assumption that information exists. They are therefore shown to be irresolvable within tracework, though not necessarily irresolvable in other philosophical, ethical, or political registers.

The pattern of resolution presented above is what should be expected of a sufficiently foundational theory of information. A theory that derives everything from a single minimal assumption, under a lock that permits no additional principle to be introduced to secure a particular result, cannot resolve them piecemeal. They either all fall together, or none fall at all. Tracework's reach lies not in extending information by analogy, but in showing how the same minimal structure reappears under progressively more complex conditions of stabilisation, transformation, and continuation.

That these problems remained open for as long as they did is largely explained by the dominant influence of Shannon's model on the field, which made their resolution more difficult to see. The already-stabilised primitives of information, bit, ban, sign, datum, and natural sign, inherited the same prior assumption and therefore implied the same cluster of consequences: time entering only as an external parameter rather than a constitutive condition, absence treated as a limiting case rather than a mode of information in its own right, transitivity assumed rather than derived, and meaning either bracketed or imported because the unit's own stabilisation was never

explained [s. 2.2]. The field's unresolved questions accumulated because each was a symptom of the same unexamined starting point, encountered from a different angle. They were not problems to be resolved, but boundaries that delimited how implicit assumptions were doing unacknowledged work.

Preprint

9. Conclusion: the limits of information science

Information science has long been uncertain of its own status (cf. Goffman, 1970; Belkin, 1975; Farradane, 1976; Roberts, 1976). Described as “a tangled skein of many intertwining threads” (Shera and Cleveland, 1977), it has from its inception asked what kind of discipline it is. Its early decades were marked by a pronounced anxiety about its standing as a scientific discourse (Belkin, 1978), an anxiety partially alleviated by the later cognitive turn (Belkin, 1990; Ellis, 1992; Buckland and Liu, 1995). Yet that shift did not resolve a deeper difficulty: whether the objects of information science possess sufficient coherence to sustain a unified field. In recent decades information science has struggled to demonstrate continuing relevance at a time when information itself has come to dominate social, cultural, and political organisation. A quarter of a century later, Warner’s question remains pointed and newly urgent: w(h)ither information science? (Warner, 2001).

Brookes called the nature of information the fundamental problem of information science (1975: 44). That problem has remained open for half a century because unlike librarianship, which is anchored in an institutional history and tradition (cf. Tredinnick, 2026), information science has had no fixed class of objects around which its identity can stabilise. Documents, records, data, systems, practices, classifications, collections, infrastructures, and users do not form a settled domain in the usual sense. In tracework’s terms, they are configurations within transformation space: stabilised, coupled, transformed, preserved, accessed, and reconstructed under different conditions. The instability this produces is structural. Information science has seemed dispersed because its objects are themselves distributed across transformation space, and information science occupies a position in relation to them, rather than outside them.

Tracework redescribes information science from the inside. Developed within its disciplinary context, and committed to its concerns and history, tracework describes the transformation space within which the objects and practices of information science become operationally intelligible. The selective examples developed throughout this paper indicate only part of this potential. Provenance stabilises causal history as an inferential condition, preserving the relation between records and the structured circumstances of their production, custody, and transformation [s. 3.6]. Domain analysis describes the domain-relativity of reconstruction, which tracework derives from informational intransitivity: inferential potential is not transmitted intact across contexts but reconstituted under the coupling conditions of the domain in which a trace is taken up [s. 2.6]. Bibliographic control operationalises scale-relative trace equivalence, treating copies, editions, works, and items as equivalent or distinct according to the resolution at

which inferential potential is stabilised [s. 3.2; s. 4.2; s. 4.3]. Classification organises paradigmatic and combinatorial coupling, imposing coherence structures that determine which relations among traces become admissible within a domain [s. 3.4; s. 4.4]. Indexing re-stabilises inferential potential as points of future reconstruction, making some relations, aspects, and pathways available while leaving others less visible [s. 7.3]. Serendipity marks the appearance within a trajectory of an unanticipated but admissible continuation, made possible by paradigmatic structure and network residue [s. 4.4]. Preservation sustains inferential consistency across time under conditions of decay and transformation, not by freezing traces in place but by maintaining the possibility of further trace formation within future viability domains [s. 5.5]. Information overload marks the point at which the rate at which traces require uptake, reconstruction, and integration exceeds a domain's capacity to incorporate them sufficiently into ongoing continuation [s. 4.3; s. 5.5]. These are not disparate activities gathered under a loose professional heading, but different interventions within the same architecture of trace stabilisation, coupling, scaling, transformation, and continuation. What defines information science as a discipline is not a fixed class of objects but the scope of transformation space over which these interventions operate.

The structure of transformation space in which information is realised has appeared throughout the field's history. Otlet's Mundaneum (Rayward, 1994; 2003), the Memex (Bush, 1945), the World Brain (Wells, 1936; 1937), hypertext (Nelson, 1974; 1982), and the World Wide Web (Berners-Lee, 1989; 1999) all point towards information as a distributed relational process extended across time. These projects differ historically, technically, and institutionally, but each addresses the same implicit problem: how information can be organised, connected, retrieved, transformed, and carried forwards across domains of use. The discipline has therefore not lacked a foundation so much as it has operated within a structure difficult to identify from inside its own trajectory. Tracework is a theory of information in this tradition that emerges from information science itself. It therefore addresses a question answered in practice long before it could be answered in theory. If information science is to persist as more than a historical aggregation of methods, institutions, and inherited problems, it must recognise the structure that has always organised its practice and claim it as a foundational theory of information.

Preprint

References:

- Adams, F. and Aizawa, K. (2021) 'Causal theories of mental content', in Zalta, E.N. (ed.) The Stanford Encyclopedia of Philosophy. Available at: https://plato.stanford.edu/archives/fall2021/entries/content-causal/ (Accessed: 25 May 2026).
- Anderson, P.W. (1972) 'More is different', Science, 177(4047), pp. 393–396.
- Austin, J.L. (1962) How to Do Things with Words. 2nd edn. Edited by J.O. Urmson and M. Sbisà. Oxford: Oxford University Press, 1975.
- Bates, M.J. (2005) 'Information and knowledge: an evolutionary framework for information science', Information Research, 10(4).
- Bates, M.J. (2006) 'Fundamental forms of information', Journal of the American Society for Information Science and Technology, 57(8), pp. 1033–1045.
- Bateson, G. (1972) Steps to an Ecology of Mind: Collected Essays in Anthropology, Psychiatry, Evolution and Epistemology. Northvale, NJ: Jason Aronson.
- Beavers, A.F. (2016) 'A brief introduction to the philosophy of information', Logeion: Filosofia da Informação, 3(1), pp. 16–28.
- Bekenstein, J.D. (1981) 'Universal upper bound on the entropy-to-energy ratio for bounded systems', Physical Review D, 23(2), pp. 287–298.
- Belkin, N.J. (1975) 'Towards a definition of information for informatics', Proceedings of Informatics, 2, pp. 50–56.
- Belkin, N.J. (1978) 'Information concepts for information science', Journal of Documentation, 34(1), pp. 55–85.
- Belkin, N.J. and Robertson, S.E. (1976) 'Information science and the phenomenon of information', Journal of the American Society for Information Science, 27(4), pp. 197–204.
- Belkin, N.J. (1990) 'The cognitive viewpoint in information science', Journal of Information Science, 16(1), pp. 11–15.
- Berners-Lee, T. (1989) Information Management: A Proposal. Available at: http://www.w3.org/History/1989/proposal.html (Accessed: 16 October 2021).
- Berners-Lee, T. (1999) Weaving the Web: The Past, Present and Future of the World Wide Web by its Creator. London: Orion Business Press.

Preprint

- Bird, A. (2007) Nature's Metaphysics: Laws and Properties. Oxford: Oxford University Press.
- Boden, M.A. (1998) 'Creativity and artificial intelligence', Artificial Intelligence, 103(1–2), pp. 347–356.
- Boden, M.A. (2004) The Creative Mind: Myths and Mechanisms. London: Routledge.
- Boden, M.A. (2009) 'Computer models of creativity', AI Magazine, 30(3), pp. 23–34.
- Boroditsky, L. (2000) 'Metaphoric structuring: understanding time through spatial metaphors', Cognition, 75(1), pp. 1–28.
- Brier, S. (2008) Cybersemiotics: Why Information Is Not Enough!. Toronto: University of Toronto Press.
- Brillouin, L. (1962) Science and Information Theory. 2nd edn. New York: Academic Press.
- Brookes, B.C. (1975) 'The fundamental problem of information science', Informatics, 2, pp. 42–49.
- Brookes, B.C. (1980) 'The foundations of information science. Part I. Philosophical aspects', Journal of Information Science, 2(3–4), pp. 125–133.
- Buckland, M.K. (1991) 'Information-as-thing', Journal of the American Society for Information Science, 42(5), pp. 351–360.
- Buckland, M.K. (1997) 'What is a "document"?', Journal of the American Society for Information Science, 48(9), pp. 804–809.
- Buckland, M.K. (1998) 'What is a "digital document"?', Document Numérique, 2(2), pp. 221–230.
- Buckland, M.K. and Liu, Z. (1995) 'History of information science', Annual Review of Information Science and Technology, 30, pp. 385–416.
- Burgin, M. (2003) 'Information theory: a multifaceted model of information', Entropy, 5(2), pp. 146–160.
- Burgin, M. (2004) 'Data, information and knowledge', Information, 7, pp. 47–58.
- Burgin, M. (2005) 'Is information some kind of data?', in Petitjean, M. (ed.) Proceedings of the Third Conference on the Foundations of Information Science (FIS 2005). Basel: MDPI.
- Burgin, M. (2009) Theory of Information: Fundamentality, Diversity and Unification. Singapore: World Scientific.
- Bush, V. (1945) 'As we may think', The Atlantic Monthly, 176(1), pp. 101–108.
- Calder, A. (1991) The Myth of the Blitz. London: Jonathan Cape.
- Capurro, R., Fleissner, P. and Hofkirchner, W. (1997) 'Is a unified theory of information feasible? A trialogue', World Futures, 49(3–4), pp. 213–234.

- Capurro, R. and Hjørland, B. (2003) 'The concept of information', Annual Review of Information Science and Technology, 37(1), pp. 343–411.
- Carruthers, M. (2008) The Book of Memory. 2nd edn. Cambridge: Cambridge University Press.
- Chomsky, N. (1965) Aspects of the Theory of Syntax. Cambridge, MA: MIT Press.
- Cilliers, P. (1998) Complexity and Postmodernism. London: Routledge.
- Cilliers, P. (2005) 'Complexity, deconstruction and relativism', Theory, Culture and Society, 22(5), pp. 255–267.
- Clanchy, M.T. (1993) From Memory to Written Record. 2nd edn. Oxford: Blackwell.
- Cornelius, I. (2002) 'Theorizing information for information science', Annual Review of Information Science and Technology, 36(1), pp. 392–425.
- Crnkovic, G.D. and Hofkirchner, W. (2011) 'Floridi's open problems in philosophy of information, ten years later', Information, 2(2), pp. 327–359.
- Davidson, D. (2005) Truth, Language and History. Oxford: Clarendon Press.
- Deleuze, G. (1993) The Fold. Minneapolis: University of Minnesota Press.
- Derrida, J. (1967) Of Grammatology. Baltimore: Johns Hopkins University Press, 1976.
- Dretske, F. (1981) Knowledge and the Flow of Information. Oxford: Blackwell.
- Dretske, F. (1983) 'Précis of Knowledge and the Flow of Information', Behavioral and Brain Sciences, 6(1), pp. 55–90.
- Duguid, P. (2021) 'Introduction', in Blair, A. et al. (eds.) Information: A Historical Companion. Princeton: Princeton University Press.
- Eames, C. and Eames, R. (1977) Powers of Ten. [Film]. Santa Monica: Eames Office.
- Edmunds, A. and Morris, A. (2000) 'The problem of information overload in business organisations', International Journal of Information Management, 20, pp. 17–28.
- Ellis, D. (1992) 'The physical and cognitive paradigms in information retrieval research', Journal of Documentation, 48(1), pp. 45–64.
- Erickson, M. (2012) 'Network as metaphor', International Journal of Criminology and Sociological Theory, 5(2), pp. 912–921.
- Fairthorne, R.A. (1967) 'Morphology of information flow', Journal of the ACM, 14(4), pp. 710–719.
- Farradane, J. (1976) 'Towards a true information science', The Information Scientist, 10(3), pp. 91–101.
- Floridi, L. (2003) 'From data to semantic information', Entropy, 5(2), pp. 125–145.

Preprint

- Floridi, L. (2004a) 'Open problems in the philosophy of information', Metaphilosophy, 35(4), pp. 554–582.
- Floridi, L. (2004b) 'Information', in Floridi, L. (ed.) The Blackwell Guide to the Philosophy of Computing and Information. Oxford: Blackwell.
- Floridi, L. (2009) 'Philosophical conceptions of information', in Sommaruga, G. (ed.) Formal Theories of Information. Berlin: Springer.
- Floridi, L. (2010) Information: A Very Short Introduction. Oxford: Oxford University Press.
- Floridi, L. (2013) The Philosophy of Information. Oxford: Oxford University Press.
- Foster, A. and Ford, N. (2003) 'Serendipity and information seeking', Journal of Documentation, 59(3), pp. 321–340.
- Goffman, W. (1970) 'Information science: discipline or disappearance', Aslib Proceedings, 22(12), pp. 589–596.
- Goody, J. and Watt, I. (1963) 'The consequences of literacy', Comparative Studies in Society and History, 5(3), pp. 304–345.
- Gould, S.J. (1996) Life's Grandeur: The Spread of Excellence from Plato to Darwin. London: Jonathan Cape.
- Hall, S. (1980) 'Encoding/decoding', in Hall, S., Hobson, D., Lowe, A. and Willis, P. (eds.) *Culture, Media, Language: Working Papers in Cultural Studies, 1972–79*. London: Hutchinson, pp. 128–138.
- Harris, R. (1981) The Language Myth. London: Duckworth.
- Harris, R. (1995) Signs of Writing. London: Routledge.
- Harris, R. (2000) Rethinking Writing. London: Continuum.
- Hartley, R.V.L. (1928) 'Transmission of information', Bell System Technical Journal, 7(3), pp. 535–563.
- Havelock, E.A. (1986) The Muse Learns to Write: Reflections on Orality and Literacy from Antiquity to the Present. New Haven: Yale University Press.
- Hawking, S.W. (1992) 'The chronology protection conjecture', Physical Review D, 46(2), pp. 603–611.
- Heaney, S. (1999) *Beowulf: A New Translation*. London: Faber and Faber.
- Hjørland, B. (2023) 'Information', *ISKO Encyclopedia of Knowledge Organization*. Available at: https://www.isko.org/cyclo/information (Accessed: 28 July 2025).
- Hjørland, B. and Albrechtsen, H. (1995), Toward a new horizon in information science: domain-analysis, *Journal of the American Society for Information Science*, 46(6), pp. 400–425.

- Hobsbawm, E. and Ranger, T. (eds.) (1983) The Invention of Tradition. Cambridge: Cambridge University Press.
- Hofkirchner, W. (2009) 'How to achieve a unified theory of information', tripleC, 7(2), pp. 357–368.
- Hofkirchner, W. (2012) Emergent Information: A Unified Theory of Information Framework. Singapore: World Scientific.
- Hofkirchner, W. (2013) 'Emergent information: when a difference makes a difference', tripleC, 11(1), pp. 6–12.
- Kronfeldner, M. (2018) 'Explaining Creativity', in Gaut, B. and Kieran, M. (eds.) *Routledge Handbook on Creativity and Philosophy*. London: Routledge, pp. 213–229.
- Krzanowski, R. (2023) Information: What We Do and Do Not Know — A Review. Available at: https://www.researchgate.net/publication/370105722 (Accessed: 30 July 2025).
- Lakoff, G. and Johnson, M. (1980) Metaphors We Live By. Chicago: University of Chicago Press.
- Landauer, R. (1991) 'Information is physical', Physics Today, 44(5), pp. 23–29.
- Lewis, D. (1973) 'Causation', Journal of Philosophy, 70(17), pp. 556–567.
- Logan, R.K. (2012) 'What is information? Why is it relativistic and what is its relationship to materiality, meaning and organization', Information, 3(1), pp. 68–91.
- Manson, S.M. (2000) 'Simplifying complexity: a review of complexity theory', Geoforum, 32(3), pp. 405–414.
- Mercier, H. and Sperber, D. (2017) *The Enigma of Reason*. Cambridge, MA: Harvard University Press.
- Mingers, J. (1999) 'Information, meaning and communication: an autopoietic approach to linking the social and individual', Cybernetics and Human Knowing, 6(1), pp. 25–42.
- Mingers, J. (2004) Information, Knowledge and Truth: A Polyvalent View. Report No. 77. Canterbury: Kent Business School. Available at: https://kar.kent.ac.uk/25462/1/77_Web_version.pdf (Accessed: 29 July 2025).
- Nagel, T. (1974) 'What is it like to be a bat?', The Philosophical Review, 83(4), pp. 435–450.
- Nelson, T.H. (1974) Computer Lib / Dream Machines. Self-published.
- Nelson, T.H. (1982) Literary Machines. Sausalito, CA: Mindful Press.
- Nida, E.A. (1964) *Toward a Science of Translating: With Special Reference to Principles and Procedures Involved in Bible Translating*. Leiden: Brill.
- Ong, W.J. (1982) Orality and Literacy: The Technologizing of the Word. London: Methuen.

Preprint

- Parry, M. (1928) 'Studies in the Epic Technique of Oral Verse-Making. I: Homer and Homeric Style', *Harvard Studies in Classical Philology*, 41, pp. 73–147.
- Parry, M. (1930) 'Studies in the Epic Technique of Oral Verse-Making. II: The Homeric Language as the Language of an Oral Poetry', *Harvard Studies in Classical Philology*, 43, pp. 1–50.
- Poster, M. (1990) The Mode of Information: Poststructuralism and Social Context. Cambridge: Polity Press.
- Qvortrup, L. (1993) 'The controversy over the concept of information: an overview and selected annotated bibliography', *Cybernetics and Human Knowing*, 1(1), pp. 3–24.
- Raber, D. (2003) The Problem of Information: An Introduction to Information Science. Lanham, MD: Scarecrow Press.
- Ranganathan, S.R. (1933) Colon Classification. Madras: Madras Library Association.
- Rayward, W.B. (1994) 'Visions of Xanadu: Paul Otlet (1868–1944) and hypertext', Journal of the American Society for Information Science, 45(4), pp. 235–250.
- Rayward, W.B. (2003) 'Knowledge organisation and a new polity: the rise and fall of the ideas of Paul Otlet', Transnational Associations, 1–2, pp. 4–15.
- Roberts, N. (1976) 'Social considerations towards a definition of information science', Journal of Documentation, 32(4), pp. 249–257.
- Ruyer, R. (2024) Cybernetics and the Origin of Information. Translated by A. Berger Soraruff, A. Iliadis, D.W. Smith and A. Woodward. New York: Rowman & Littlefield.
- Salmon, W. (1994) 'Causality without counterfactuals', Philosophy of Science, 61(2), pp. 297–312.
- Saussure, F. de (1916) Course in General Linguistics. Translated by W. Baskin. New York: Philosophical Library, 1959.
- Sayers, D.L. (1949) *The Comedy of Dante Alighieri the Florentine: Cantica I — Hell*. Harmondsworth: Penguin.
- Sayers, D.L. (1955) *The Comedy of Dante Alighieri the Florentine: Cantica II — Purgatory*. Harmondsworth: Penguin.
- Shannon, C.E. (1948) 'A mathematical theory of communication', Bell System Technical Journal, 27(3), pp. 379–423.
- Shannon, C.E. (1953) 'The lattice theory of information', Transactions of the IRE Professional Group on Information Theory, 1(1), pp. 105–107.
- Shera, J.H. and Cleveland, D.B. (1977) 'History and foundations of information science', *Annual Review of Information Science and Technology*, 12, pp. 249–275.

- Stonier, T. (1990) Information and the Internal Structure of the Universe: An Exploration into Information Physics. New York: Springer.
- Stonier, T. (1991) 'Towards a new theory of information', Journal of Information Science, 17(5), pp. 257–263.
- Stonier, T. (1992) Beyond Information: The Natural History of Intelligence. Berlin: Springer.
- Tredinnick, L. (2006) Digital Information Contexts: Theoretical Approaches to Understanding Digital Information. Oxford: Chandos Publishing.
- Tredinnick, L. (2008) *Digital Information Culture: The Individual and Society in the Digital Age*. Oxford: Chandos Publishing.
- Tredinnick, L. (2015) The Nature of Information: An Analysis of the Historically Situated Socio-Cultural Assumptions Concerning the Nature of Information under Changing Technological Conditions of its Production, Reproduction, Dissemination and Use. Doctoral thesis. London Metropolitan University. Available at: https://repository.londonmet.ac.uk/633/ (Accessed: 20 May 2026).
- Tredinnick, L. (2023a) 'The intricate web: network and rhizome metaphors in hypertext and the web and the epistemic challenge of fake news', Journal of Documentation, 79(6), pp. 1485–1501.
- Tredinnick, L. (2023b) 'Dangerous data: analytics and information behaviour in the commercial world', *Business Information Review*, 40(1), pp. 10–20. doi: 10.1177/02663821231160964.
- Tredinnick, L. (2025) 'Theory and metatheory in the nature of information: review and thematic analysis', Information, 16(9), p. 791.
- Tredinnick, L. (2026) 'Early library history and the invention of tradition', Library & Information History, 42(1), pp. 3–25.
- Tredinnick, L. and Laybats, C. (2018) 'Information, communication and overload', *Business Information Review*, 35(3), pp. 96–98. doi: 10.1177/0266382118792762.
- Tredinnick, L. and Laybats, C. (2022) 'Editorial: Serendipity and information discovery', *Business Information Review*, 39(1), pp. 6–8. doi: 10.1177/02663821221088892.
- Turing, A.M. (1941/2012) The Applications of Probability to Cryptography. arXiv:1505.04714. Available at: https://arxiv.org/pdf/1505.04714.
- Turing, A.M. (1950) 'Computing machinery and intelligence', in Copeland, B.J. (ed.) The Essential Turing. Oxford: Oxford University Press, 2004, pp. 441–464.
- Urry, J. (2005) 'The complexity turn', Theory, Culture and Society, 22(5), pp. 1–14.

Preprint

- Van Rijsbergen, C.J. and Lalmas, M. (1996) 'Information calculus for information retrieval', Journal of the American Society for Information Science, 47(5), pp. 385–398.
- Warner, J. (2001) 'W(h)ither information science?/!', *The Library Quarterly*, 71(2), pp. 243–255.
- Weber, M. (1922) Economy and Society. Translated by G. Roth and C. Wittich. Berkeley: University of California Press, 1978.
- Wells, H.G. (1936) The Idea of a World Encyclopaedia: A Lecture Delivered at the Royal Institution November 20th 1936. London: Hogarth Press.
- Wells, H.G. (1937) 'The world brain: the idea of a permanent world encyclopaedia'. Available at: https://sherlock.ischool.berkeley.edu/wells/world_brain.html (Accessed: 16 May 2021).
- Wersig, G. and Neveling, U. (1975) 'The phenomena of interest to information science', Information Scientist, 9(4), pp. 127–140.
- Wiener, N. (1961) Cybernetics, or Control and Communication in the Animal and the Machine. Cambridge, MA: MIT Press.
- Williams, R. (1977) *Marxism and Literature*. Oxford: Oxford University Press.
- Wittgenstein, L. (1953) Philosophical Investigations. Translated by G.E.M. Anscombe, P.M.S. Hacker and J. Schulte. Revised 4th edn. Oxford: Wiley-Blackwell, 2009.
- Woodward, J. (2003) Making Things Happen: A Theory of Causal Explanation. Oxford: Oxford University Press.
- Yates, F.A. (1966) *The Art of Memory*. London: Routledge and Kegan Paul.
- Yovits, M.C. (1969) 'Information science: towards the development of a true scientific discipline', American Documentation, 20(4), pp. 369–376.
- Zins, C. (2007) 'Conceptual approaches for defining data, information, and knowledge', Journal of the American Society for Information Science and Technology, 58(4), pp. 479–493.

Preprint

Appendix A: The Two-Dimensionality of Transformation Space

A1 Introduction: the claim that transformation space is exactly two-dimensional, defined by correspondence and coherence as orthogonal axes, is a necessary consequence of the conditions under which traces remain inferentially legible under temporal asymmetry. What follows derives this result.

A2 The core argument: a trace is inferentially legible if and only if backward inference can reach the causal conditions that produced it, that is, if reconstruction of prior causal conditions remains possible from the current configuration. For this to hold, transformations of the trace must preserve constraint on admissible prior configurations. A transformation that preserves no such constraint renders all prior configurations equally admissible and destroys inferential legibility entirely. Any constraint relevant to inferential legibility must therefore function by restricting which prior configurations count as admissible predecessors. If it does not do this, it does not contribute to inferential legibility.

A current configuration can be connected to its admissible predecessors in two ways. Either admissibility is determined by the generative process that produced the configuration, by what caused it, or it is determined by compatibility with the structured network of relations in which the configuration participates, by where it sits. These two forms are exhaustive. Any constraint on admissible prior configurations must operate either by restricting which configurations could have produced the current one, or by restricting which configurations are compatible with it within a relational structure. There is no third form in the present architecture. Any proposed alternative must either specify conditions of production or conditions of compatibility. If it does neither, it imposes no restriction on admissible predecessors and contributes nothing to inferential legibility.

Preservation of constraint by reference to causal production is correspondence. A transformation preserves correspondence when the resulting configuration remains compatible with the causal constraints that produced the original, when admissible prior configurations remain restricted by generative continuity. Preservation of constraint by reference to relational position is coherence. A transformation preserves coherence when the resulting configuration remains compatible with the relational structure in which the original participated. The relational structure that grounds coherence is constituted by the coupling constraints established in section 3. Coherence is therefore not reducible to correspondence.

Preprint

Any form of constraint preservation sufficient to maintain inferential legibility must either reduce to correspondence, reduce to coherence, or fail to preserve constraint on admissible prior configurations, in which case it is irrelevant to inferential legibility. No independent third axis exists.

A3 Testing the candidates: the argument can be tested against the obvious candidates for a third axis.

Modal constraints, physical or logical possibility, appear to constrain configurations without reference to production or structure. But physical possibility is determined by causal laws and therefore by generative compatibility: preserving it is correspondential. Logical possibility is determined by inferential relations: preserving it is coherential. Modal constraints reduce to correspondence and coherence and define no independent axis.

Probability preserves constraint through statistical regularities of generative processes. The restriction on admissible predecessors is therefore correspondential. Probability is correspondence operating at the level of distributions.

Scale constrains admissible configurations at a given level of description through both causal production and relational compatibility jointly. Scale parameterises the geometry of the correspondence-coherence plane; it does not add a dimension to it.

Temporal ordering maintains relations of before and after. These arise from causal production. Temporal ordering is a form of correspondence.

Energy and complexity both restrict which prior configurations are physically or generatively admissible. Both are correspondential.

Similarity, identity, and functional equivalence appear to constrain predecessors without reference to production or relational position. But similarity between configurations is either grounded in shared causal production or in shared relational position. Pure similarity without either is a property of configurations in themselves rather than a restriction on admissible predecessors, and contributes nothing to inferential legibility. The same applies to identity and functional equivalence.

Forward predictive capacity, the ability to project admissible successors, appears to define a forward constraint distinct from backward inference. But forward projection is the continuation

of the same constraints under the opposite temporal orientation. Correspondence under productive tracework yields propagation; coherence under productive tracework yields reframing. Direction is a parameter of traversal along the existing two axes. The non-commutativity of tracework arises from the sequential accumulation of constraints along these axes, not from an independent third dimension.

A4 Orthogonality: correspondence and coherence are orthogonal in the sense that neither is definable in terms of the other. Correspondence constrains admissible prior configurations by restricting which configurations could have produced the current one under causal transformation. Coherence constrains admissible prior configurations by restricting which configurations are compatible with the current one within a structured relational domain.

The irreducibility of each to the other follows from the asymmetry of causal constraint. Causal production is directional: the trace was produced by prior conditions but does not produce them. Relational compatibility is non-generative: it constrains co-membership within a structure without specifying a directional relation of production. Reducing correspondence to coherence eliminates the asymmetry that makes backward inference operative. Structural causal models such as Pearl's do-calculus confirm this non-reducibility: the explicit intervention operator required to distinguish causal from correlational constraints corresponds to the correspondence-coherence distinction in tracework.

The independence of the two dimensions is demonstrated by transformations that preserve one while disrupting the other. Physical copying preserves correspondence while potentially disrupting coherence: the copy retains causal compatibility with the same prior conditions while losing relational position within the network in which the original participated. Translation preserves coherence while disrupting correspondence: the translated text retains its inferential consistency within a relational network while losing causal grounding in the original production conditions. These are not edge cases but paradigm instances of how the two dimensions operate independently.

A5 Conclusion: transformation space is exactly two-dimensional. Correspondence and coherence are the only independent axes of constraint preservation sufficient to maintain inferential legibility under temporal asymmetry. Any proposed third axis either reduces to one of these, represents a special case, or parameterises their interaction without introducing a new dimension. The two-dimensionality is not a classificatory choice but a structural consequence of the conditions under which constraints on admissible predecessors can be maintained.

Appendix B: Glossary of original theoretical terms

Anticipatory alignment: alignment that occurs when domains coordinate their projections of future states under the constraints of their prior commitments, such that joint action, coordinated planning, and shared expectation become possible.

Anticipatory coupling: the subjective domain's attunement to succession under productive tracework; the mode in which future continuation becomes obligatory, through which a system binds itself to one trajectory rather than another.

Aspectual claim: the claim that subjectivity, meaning, and experience are the aspect under which self-referential tracework is encountered from within; they describe how tracework is experienced in the delimitation of that experience, and the claim accordingly forecloses any inference from structural description to the nature or moral, legal, or political status of the system described.

Brittleness: the structural consequence of redundancy decay: the condition of a trace-structure that retains inferential consistency and local coherence but has lost sufficient redundancy to absorb perturbation without collapse.

Causal history: the stabilising condition under which a trace retains the structural effects of prior causal interaction in a form that can support inferential potential. It is a structured sequence of constraint-forming transformations through which a trace's inferential profile is established.

Co-commitment: the minimal threshold condition for the intersubjective domain; the coupling of two or more subjective domains through a shared field of continuation in which the commitments, selections, revisions, or refusals of one domain become materially relevant to what another can sustain, contest, revise, or carry forwards.

Coarse-graining: the process by which detail is compressed into coarser structures as the scale of description increases; a stabilisation process arising naturally from the dynamics of tracework itself, by which traces and trace-structures are necessarily re-stabilised at coarser scales under the appropriate stabilisation domain.

Coherence plasticity: the form of recursive plasticity corresponding to coherence stabilisation; the capacity of a trace to reorganise its relational structure without losing inferential consistency within the network in which it is embedded.

Coherence coupling: the inferential constraint arising when traces are embedded within a shared relational network whose internal consistency conditions constrain the admissible stabilisations of each constituent trace; increases backward inferential determinacy by enforcing internal consistency while simultaneously reducing forward generativity.

Coherence decay: the loss of relational compatibility within a trace network; a trace persists but becomes increasingly isolated as the compatibility relations that sustain its embedding weaken or dissolve.

Coherence stabilisation: the condition under which a trace's inferential potential is sustained by its embedding within a network of other traces; not internal consistency in the abstract but continued relational compatibility within a trace network.

Combinatorial coupling: the productive counterpart of paradigmatic coupling; governs whether traces may be jointly stabilised so as to participate in tracework, determining which traces are admissible as co-present, co-selectable, or co-activatable within a stabilisation domain.

Compositionality: the constitution of trace-structures through recursive stabilisation rather than their presupposition as given configurations; names the fact that trace-structures are themselves outcomes of tracework and therefore inherit path-dependence and non-commutativity at the level of organisation.

Construction: the productive trajectory operation that generates new trajectories under constraint; forward causal production guided by the viability and stabilisation conditions of its domain rather than by an existing inferential template or by deductive entailment from prior configuration.

Correspondence plasticity: the form of recursive plasticity corresponding to correspondence stabilisation; the capacity of a trace to sustain inferential consistency across changes of physical instantiation while remaining compatible with its causal conditions of formation.

Correspondence decay: drift away from constraint compatibility with generative conditions; accumulated transformation progressively erodes constraint relations such that the causal history embedded in the trace becomes increasingly opaque.

Preprint

Correspondence stabilisation: the condition under which the structural effects of prior causal interaction remain recoverable; continued constraint compatibility across transformation, neither representation nor a relation between traces and any external reality.

Coupling constraints: the relational conditions that determine how the stabilisation of one trace affects the stabilisation conditions of others; a constraint on the inferential potential that each trace can sustain given the fact of the other.

Coupling decay: the ways in which trace coupling constraints drift, thin, or fail under ongoing transformation; the progressive degradation of the conditions of mutual operability rather than the disappearance of traces.

Differentiation: the inferential stabilising condition under which a trace is distinguishable as a variation in state relative to a given stabilisation domain, subject to scale-dependent limits of resolution and discrimination.

Differentiation decay: the erosion of distinctness between traces; traces lose their separability, becoming indistinguishable from one another or from background variation.

Dispersal: the upper-bound productive threshold failure, arising when the finite speed at which physical constraints can propagate through a substrate is exceeded, such that coherence degrades into effectively independent local dynamics.

Dissipation: the lower-bound inferential threshold failure, occurring when causal variations lack sufficient persistence, differentiation, and causal history to support the formation of a trace, leaving variation transient and unable to function informationally.

Domain coupling: the productive counterpart of resolution coupling; governs whether tracework can continue under forward transformation when scale, context, or descriptive framework shifts.

Extension: the capacity of a trace-structure to extend its causal history forwards through new trace formation, rather than merely preserving existing organisation; continuity of organisation through further production in productive tracework.

Fragmentation: the relational threshold condition marking the limit of relational co-existence, the point at which coherence coupling can no longer bind a trace-structure as a single coupled domain and the structure decomposes into disconnected components that cannot be reintegrated under the operative coupling constraints.

Generation: the productive stabilising condition determining whether prior causal conditions, under the relevant viability constraints, yield a configuration capable of grounding further trace formation; the productive counterpart of causal history.

Generativity: the capacity of a trace-structure to participate in the production of new configurations through recursive combination under coupling, functioning as a stabilised organisation that conditions, channels, and constrains how new traces are generated, combined, or transformed.

Inferential consistency: the structured and necessarily partial preservation of causal constraints that operate at the relevant stabilisation domain; the condition under which a trace's ontological persistence is realised, rather than through its physical substrate.

Inferential potential: the capacity of a trace to generate further traces under appropriate conditions; domain-relative rather than agent-relative, defined with respect to the stabilisation constraints and resolution limits under which a structure is discriminable and persistent.

Interior viability domain: the zone between the four threshold-adjacent regions in which trace formation, persistence, and transformation are jointly sustained under compatible stabilisation constraints, and within which inferential and productive tracework remain simultaneously viable.

Network residue: the distributed deformation of connectivity produced by prior tracework through the system; consists of concentrations of constraint, asymmetries of accessibility, and scar-like alterations of the network that preserve causal pressure without preserving path identity.

Non-commutativity: the property of tracework by which different sequences of constraint formation constitute different inferential structures, even where they converge on equivalent outcomes; tracework is not commutative.

Normative alignment: alignment that occurs when domains recognise and partially bind themselves to shared standards as extensions or intersections of their existing commitments, such that institutions, practices, and collective norms persist as shared structures capable of binding divergent trajectories within tolerable limits.

Preprint

Normative coupling: the subjective domain's attunement to coherence under standards of admissibility internal to its own trajectory; traces are evaluated against standards of correctness, legitimacy, or appropriateness that the system has itself stabilised over time.

Ossification: the lower-bound productive threshold failure, arising when trace production falls below the rate required to sustain forward projection of backward inference, such that traces persist but cease to participate productively in tracework and become effectively irrecoverable as generative structures.

Paradigmatic coupling: the inferential constraint governing which traces could have occupied a given position while preserving the inferential consistency of the surrounding structure; operates through structural compatibility across a space of possible substitutions.

Paradigmatic dimension: the space of inferentially admissible alternatives at each position within organised tracework, governing what could occupy a given position while preserving inferential consistency; the dimension of admissible alternatives.

Path-dependence: the property by which inferential consistency is determined by both the set of causal interactions from which a trace arises and the order in which those interactions are stabilised; identical sets of traces stabilised in different sequences produce structurally divergent trace-structures.

Pattern persistence: the capacity of a trace-structure to maintain its governing relational organisation across transformation and variation in its constituent traces, so long as the inferential constraints defining the structure remain operative.

Physical coupling: the constraint governing the mutual constraint of traces that co-stabilise in spatial and causal proximity within the same material substrate or shared causal medium; operates through substrate-mediated, proximity-dependent effects.

Physical decay: the degradation of substrate-mediated coupling conditions that sustain trace interaction in spatial and causal proximity; a coupling-level mechanism that accelerates drift towards multiple thresholds by degrading the material infrastructure through which coupling relations are preserved.

Presence-independence: the principle that information does not require presence to be informational; wherever an absence is the stabilised consequence of prior causal conditions, it

functions informationally in the same structural sense as a presence, and may sustain inferential potential within the relational network in which it remains a nexus.

Productive decay: the attenuation of forward causal participation; a trace's ability to sustain further propagation falls below the threshold required for continuation within a given viability domain.

Productive tracework: tracework considered under the forward temporal orientation; the processes of causal generation, propagation, and continuation through which traces participate in future tracework.

Projection: the productive trajectory operation that determines admissible continuations from the current configuration of constraints; not anticipation, prediction, or representation, but a structural property of the trajectory at a given position.

Propagation: the productive stabilising condition under which causal effects, under the relevant viability conditions, extend forwards far enough to sustain further trace formation; the productive counterpart of persistence.

Re-stabilisation: the process by which competing inferential potential profiles are compressed into a single dominant one while consolidating productive pathways into a coherent continuation; occurs at coarser scales as a response to the fundamental instability generated by increasing density and heterogeneity of constituent traces.

Reconstruction: the inferential activity through which prior causal conditions are recovered from a trace under the backward orientation. Reconstruction is itself a form of continuation: because any articulation of information must itself take the form of a trace relation, every actualisation of a trace's inferential potential is itself a further trace relation, generated under the productive orientation even as it operates inferentially.

Recursive plasticity: the umbrella property comprising the three forms, correspondence plasticity, coherence plasticity, and redundancy plasticity, through which a trace or trace-structure sustains inferential consistency across change of instantiation, relational reorganisation, and distributed perturbation respectively.

Redundancy plasticity: the form of recursive plasticity corresponding to redundancy coupling; the capacity of a trace-structure to distribute inferential potential across multiple coupling

relations such that no single pathway failure is sufficient to undermine inferential consistency at the operative scale.

Redundancy coupling: the productive coupling constraint governing whether inferential potential and further trace formation are distributed across multiple coupling relations within a trace-structure, such that no single trace or pathway is solely responsible for sustaining productive participation.

Redundancy decay: the thinning of distributed support within a coupled trace-structure; the mechanism by which systems shift from stable to brittle without any dramatic local collapse.

Reflexive coupling: the subjective domain's attunement to the organisation of its own coupling architecture; the basis of self-correction and the coupling-level correlate of wisdom, making the coupling architecture itself an object of ongoing regulation.

Reframing: the productive counterpart of translation; movement along the coherence axis under productive orientation in which the forward causal arrow generates new relational organisation within a coherence network, producing new inferential relations through reorganisation of the coherence structure itself.

Relevance alignment: alignment that occurs when domains converge sufficiently on what is stabilised as salient, meaningful, or significant within their respective trajectories, such that domains can communicate across difference because what matters in one trajectory remains recognisable as mattering in another.

Remediation: the inferential operation that restores continuity in a disrupted trajectory by substituting a new causal segment under constraint; forward causal production guided by backward inference, preserving inferential compatibility with the surrounding tracework rather than causal history as such.

Resolution coupling: the inferential constraint governing how stabilised differences at one level constrain, and are constrained by, their coarser- or finer-grained decompositions; a relation across scales and viability domains governing the preservation of inferential consistency under redescription.

Preprint

Retrodiction: the inferential reconstruction of causal history or restoration of the continuity of inferential consistency under irreversible loss; applies to entire trajectories, reconstructing both what happened and what could have happened.

Saturation: the upper-bound inferential threshold failure, arising when the maximal density of distinguishable states that can be supported within a bounded region of spacetime is reached, constraining further differentiation.

Scale decay: the loss of operability of traces and trace-structures across changes of resolution or levels of description; the erosion of cross-domain admissibility when cross-scale constraints become misaligned.

Selective coupling: the subjective domain's attunement to paradigmatic admissibility under commitment; renders continuation path-dependent by binding the system to a path within the admissible space and generating forward obligation from prior selective stabilisation.

Sequence coupling: the inferential coupling constraint governing what may have preceded a trace given the partial and compressive preservation of causal history; determines a reconstructive ordering over traces rather than a complete recoverable chronology.

Stabilisation domain: the narrower operative regime, relative to a given scale or resolution, within which traces and trace-structures remain discriminable and persistent.

Subjective domain: the domain that arises where tracework operates at sufficient recursive depth for a system to sustain the backward inferential and forward productive orientations together, operating over its own trace-structure under the same constraints it applies to all others and integrating that self-referential operation into its continuation.

Subjective regulation: the organisation required to maintain trajectories within the interior domain despite ongoing drift towards threshold breakdown; regulates which portions of inferential potential become operative, which are taken forwards under commitment, and which are allowed to decay below threshold.

Succession coupling: the productive coupling constraint governing what can follow a trace under productive continuation.

Preprint

Syntagmatic decay: the erosion of coupling constraints that preserve temporal ordering across retrospective reconstruction and forward continuation; ordering constraints thin or become unreliable such that multiple incompatible sequences become equally admissible.

Syntagmatic dimension: the realised organisation of productive constraint, ordering what follows from what under forward constraint; the dimension of realised continuation within organised tracework.

Trace: the minimal stabilised variation capable of preserving both backward inferential potential and productive continuation; its ontological persistence is realised through its inferential consistency rather than through its physical substrate.

Trace equivalence: a stabilised configuration may function as either trace or trace-structure depending on orientation and scale of access. Productively, it functions as a trace, a unit conditioning further continuation; inferentially, it functions as a trace-structure, an organisation whose consistency is reconstructed from constituent traces. Trace equivalence follows from the coupling of orientation and scale in the trace/trace-structure relation, while orientational non-equivalence prevents the two descriptions from becoming interchangeable

Trace-structure: a stabilised organisation of traces in which the relations among constituent traces persist under transformation; an independently persistent configuration whose inferential consistency no longer depends on the continued presence of the particular traces from which it was originally formed.

Tracework: the recursive processes through which stabilised traces are organised, transformed, and sustained under temporal asymmetry; the theoretical framework specifying the conditions under which information exists, persists, and transforms.

Trajectory: an ordered, constraint-governed sequence of transformations; not spatial motion but the structured movement of traces through transformation space under temporal asymmetry.

Transformation space: the full space of possible configurations and their transformations; characterised by two orthogonal axes, correspondence and coherence, which exhaust the independent directions along which admissible transformation may occur.

Preprint

Translation: transformation governed by coherence; preserves a trace's inferential consistency, what it implies, what it bears on, and what it excludes, across substantial internal structural change.

Transliteration: transformation governed by correspondence; preserves relational organisation under re-encoding while remaining compatible with the causal conditions of the trace's formation.

Variation: the productive stabilising condition under which causal production, under the relevant viability conditions, yields distinguishable configurations rather than undifferentiated flux; the productive counterpart of differentiation.

Viability domain: a bounded region of transformation space within which the minimal conditions of tracehood can in principle be realised.

Preprint

Appendix C: Named principles of tracework

The following principles derive from the single assumption that information exists. Temporal asymmetry follows as a structural consequence of the minimal conditions for information's possibility, and the named principles below are in turn derived from that constraint in the sequence given. Temporal asymmetry is not itself a principle but the constraint from which the principles of tracework derive. Each entry states the principle and the step by which it is derived; the full derivations are given in the body of the paper as specified.

P1. The principle of informational grounding. *Statement:* information is not given but produced; the foundational question shifts from what information is to how information comes to be, as a condition arising when causal processes leave sufficiently stabilised structure for generative tracework to begin. *Derivation:* if information exists, something must account for its existence; to treat it as a given elementary unit presupposes the persistence, differentiation, and structured consequence that make it identifiable as information, and so assumes an unexamined prior process. Grounding the assumption therefore requires specifying the conditions under which anything comes to exist informationally [s. 2.2].

P2. The principle of temporal extension. *Statement:* a static description of information is always a cross-section of something fundamentally extended in time; information must exist for an interval at least sufficient for it to function informationally, or it cannot exist as information; time is not the medium through which information moves but the dimension within which information exists and within which its full extension must be defined. *Derivation:* production is a process and unfolds across an interval; persistence requires that a difference endure long enough to be stabilised. A configuration existing for no interval could neither be produced nor persist, and so could not satisfy the conditions for tracehood [s. 2.6].

P3. The principle of orientational non-equivalence. *Statement:* within tracework, orientation is not a neutral perspective on an invariant object; because trace relations are constituted under temporal asymmetry, the same architecture functions differently depending on which orientation is operative; the difference is operational rather than merely perspectival, determining what can be produced, reconstructed, bound, stabilised, or carried forwards; this applies to the productive and inferential orientations, to internal and external continuation, to scalar and relational description, and to aspectual and structural specification. *Derivation:* because production and persistence are not equivalent operations, transformation is directionally ordered; what is fixed under the productive orientation at the point of formation cannot be recovered under the

inferential orientation by running that operation in reverse, so the two orientations are not two descriptions of one operation [s. 2.6].

P4. The principle of epistemic asymmetry. *Statement:* because causal processes unfold forwards but inference must proceed backwards along a chain of prior conditions, any account of information that is not purely stipulative must derive its structure from a basic asymmetry in the relationship between information and time; this asymmetry is a necessary but not sufficient condition, constraining any admissible reconstruction without generating a specific solution. *Derivation:* the causal processes that produce a trace are complete at the moment it exists, so there is no longer a forward process for inference to follow; inference about those processes must instead proceed backwards, from the trace towards its conditions. This is the epistemic expression of temporal asymmetry [s. 2.6].

P5. The principle of epistemic priority inversion. *Statement:* while on an ontic level cause precedes trace, on an epistemic level trace precedes cause; the arrow of epistemic priority runs from consequence to cause rather than from cause to consequence. *Derivation:* given epistemic asymmetry and orientational non-equivalence jointly, the epistemic starting point is the trace rather than the source; the source, including communicative intention, is not the origin from which information flows but a causal condition reconstructable in the trace [s. 2.6].

P6. The principle of informational intransitivity. *Statement:* information does not move; what appears as transmission is the preservation of sufficient causal structure across a sequence of transformations to maintain inferential potential about prior conditions at each stage; informationally, the future cannot bind the past. *Derivation:* if epistemic priority runs from trace to cause, what is fixed in production cannot determine what is reconstructed in inference; information therefore cannot be transmitted as self-identical content but propagates only through transformation, inferential potential being constituted locally wherever sufficient causal history is preserved [s. 2.6].

P7. The principle of presence-independence. *Statement:* wherever absences are the stabilised consequence of prior causal conditions, they become informational in the same structural sense as a presence; traces are not inherently present but may also be structurally absent. *Derivation:* the three inferential stabilising conditions, differentiation, persistence, and causal history, make no reference to presence; an absence that satisfies them therefore functions informationally in the same sense as a presence, requiring only that it remain materially determined within a relational network [s. 2.4].

P8. The principle of reflexive articulation. *Statement:* any articulation of information must itself take the form of a trace relation and is therefore subject to the same conditions that tracework specifies; tracework does not stand outside its subject but is itself an instance of the structure it articulates; generative tracework can therefore articulate only with tracework, and whatever bears informational significance does so only under the conditions tracework derives. *Derivation:* the conditions tracework specifies are themselves instances of generative tracework; any interface with an external system would itself have to take the form of a trace relation and so would already belong to tracework, which therefore exhausts what can be described in informational terms [s. 6.4].

P9. The principle of trace equivalence. *Statement:* any stabilised configuration may be regarded as either a trace or as a trace-structure without remainder, depending solely on whether it is accessed productively or inferentially and the level of inferential resolution at which it is taken up. *Derivation:* once a configuration's continuation is jointly constrained across the relations from which it is organised, it conditions further continuation (and so functions as a trace) and is derivative of prior traces (and so functions as a trace-structure); which description applies depends only on temporal orientation and resolution, not on any ontological difference [s. 3.2].

P10. The principle of non-self-grounding. *Statement:* no framework grounded in the assumption that information exists can specify from within its own resources the conditions under which that assumption holds; the structure is complete, its grounding is not. *Derivation:* any attempt to specify the conditions under which information exists would itself presuppose information's existence and take the form of information, reproducing the assumption rather than grounding it; a framework cannot occupy both the position of its own foundation and the conditions that precede it [s. 8.2].

Derived principle. The following principle is derived from the architecture of tracework rather than from the foundational assumption directly. It names a structural consequence that becomes visible only once the conditions of stabilisation, coupling, and transformation have been established.

DP1. The principle of Shannon availability equivalence. *Statement:* any theory of information that grounds itself in already-stabilised elementary units is structurally equivalent to Shannon's framework regardless of the theoretical vocabulary in which it is expressed or the semantic conditions it subsequently adds, and the fundamental units in use can always be described as Shannon bits. *Derivation:* once treated as already-stabilised, any informational unit assumes the

formal structure of a bit, a discriminable difference between alternatives; whatever is subsequently added (sign, datum, semantic content) leaves that underlying binary availability untouched, so the apparent incommensurability of such theories arises from post hoc extension rather than from any difference at the level of the unit [s. 2.2].

Preprint