

THE ARCHIVAL AGE · Volume I

One Crack in Every Cathedral

Moreno Nourizadeh

ORCID: 0009-0006-0174-2585

Independent Researcher, Stockholm

July 2026 | Version 1_3 | Preprint draft

One Crack in Every Cathedral

Moreno Nourizadeh

ORCID: 0009-0006-0174-2585

Independent Researcher, Stockholm

Preprint draft — V1_3 — July 2026

© 2026 Moreno Nourizadeh. All rights reserved.

A Note on This Edition

This is a volume of THE ARCHIVAL AGE, issued in preprint reader form. The text is set as a book for continuous reading. A full manual citation check is ongoing (later revisions will be issued under successive version numbers). Cross references between the volumes are given by volume and section rather than by page.

Contents

General Introduction to the Series	10
1. The theorem	10
2. The four incapacities	11
3. The expenditure booked as performance	12
4. The defence loop	12
5. The four movements of the series	13
6. The reader's contract	13
The Post-Cybernetic Organisation: Paper 3 - The Dissolution of Orders	15
<i>Abstract</i>	145
1. Introduction: The Question of Orders	16
1.1 The Proliferation of Reflexivity	16
1.2 Methodological Ground	17
1.3 Plan of the Paper	18
2. First-Order Cybernetics: The Observer Outside	18
2.1 The Founding Concepts	18
2.2 The Observer's Position	19
2.3 The Achievements of First-Order Cybernetics	20
3. Second-Order Cybernetics: The Observer Included	21
3.1 The Reflexive Turn	21
3.2 Autopoiesis and Operational Closure	22
3.3 The Conversation and the Observer of Observers	22
3.4 The Retained Commitments	23
4. The Exhaustion of Reflexivity	24
4.1 The Regress Problem	24
4.2 The Persistence of Exteriority	25
4.3 The Formal Limitation	26
5. Toward the Chiasmic Turn	26
5.1 The Insufficiency of Reflexivity	27
5.2 Merleau-Ponty's Chiasm	27
5.3 From Reflexivity to Reversibility	28

6. The Seams, Deferred	28
12. The Insufficiency of Reflexivity	29
12.1 Why More Orders Cannot Help	29
12.2 The Structure of Reversal	29
12.3 The Topology of the Chiasm	30
13. Merleau-Ponty's Ontology of Flesh	30
13.1 The Concept of Flesh	30
13.2 Reversibility as Ontological Structure	31
13.3 The Visible and the Invisible	31
13.4 Flesh and Language	32
14. From Flesh to Organisation	32
14.1 The Organisational Chiasm	32
14.2 The Organisational Flesh	33
14.3 Patterns as Organisational Flesh	33
15. The Dissolution of Orders	33
15.1 Why Orders Dissolve	33
15.2 What Replaces Orders	34
15.3 Post-Cybernetic Instead of Third-Order	34
16. Implications for Organisational Ontology	35
16.1 Organisations as Flesh	35
16.2 Incalculability as Ontological Feature	35
16.3 Navigation Instead of Control	35
17. The Requirements, Stated Once	36
25. The Mathematical Shield	36
25.1 Mathematics Against Mathematisation	36
25.2 The Cascade Structure	36
25.3 The Impossibility Theorems	36
26. The Viability Constraint	37
26.1 Finitude as Boundary Condition	37
26.2 Path Integral Formulation	37
26.3 Ecstatic Temporality Formalised	37
27. What Replaces Metrics Is the Work of the Fourth Book	37
**Paradox as Diagnosis: Foundational Contradictions from Physics to Artificial	
Intelligence**	45
<i>Abstract</i>	145
1.1 Paradox as Philosophical Method	47
1.2 What This Paper Is and Is Not	47
1.3 The Structure of This Paper	48
2.1 The Concept of a Seam	48
2.2 Historical Precedents	48
3.1 Seam 1: The Isolated System and the Black Hole Information Paradox	49
3.2 Seam 2: The Measurement Problem	49

3.3 Seam 3: The Arrow of Time	50
3.4 Seam 4: The Vacuum and Potentiality	51
3.5 Seam 5: Non-Locality	51
4.1 Seam 6: The Continuum	52
4.2 Seam 7: Gödel's Incompleteness	52
4.3 Seam 8: Set-Theoretic Foundations	53
4.4 Seam 9: The Axiom of Choice	54
5.1 Seam 10: The Hard Problem of Consciousness	54
5.2 Seam 11: The Binding Problem	54
5.3 Seam 12: The Frame Problem	55
6.1 Seam 13: The Halting Problem	56
6.2 Seam 14: The Symbol Grounding Problem	56
7.1 The Common Pattern	56
7.2 The Single Rupture	57
8.1 The Final Seam	58
8.2 Why This Seam Is the Keystone	58
8.3 Historical Pattern	59
9.1 Ecstatic Temporality	59
9.2 Metabolic Transformation	60
9.3 Autopoietic Closure	60
9.4 Embodied Situatedness	60
9.5 Field Integration	60
9.6 Constitutive Finitude	60
 The Archival Regulator	 121
<i>Abstract</i>	145
1.1 A Quarterly Review	123
1.2 The First Result: The Ledger Is Blind to the Joining	123
1.3 The Second Result: What the Instrument Cannot Capture, Someone Absorbs	125
1.4 The Third Result: The Instrument Arrives After the Horizon	126
1.5 The Fourth Result: The Improvement Can Be the Consumption	127
1.6 The Question the Four Results Pose Together	128
2.1 Two Modes of Existence for a Coordinating Capacity	128
2.2 The Archival Regulator Theorem	129
2.3 What the Theorem Does Not Say	130
2.4 One Map, One Kernel	130
2.5 The Status of the Theorem	131
3.1 The Recognition Problem	132
3.2 The Dashboard	132
3.3 The Price	133
3.4 The Framework	133
3.5 The Model	134
3.6 The Cascade	134
4.1 The Demand for a Better Instrument Is the Signature	135
4.2 The Gestell Defends Itself	135
4.3 The Constituent Cannot Total the Process	140
5.1 What the Theorem Leaves Open	140
5.2 Conditions Rather Than Prescriptions	140
5.3 The Finding	141

Three Walls	144
<i>Abstract</i>	145

General Introduction to the Series

THE ARCHIVAL AGE: A Cybernetic Philosophy of Economic Life

Draft: v0_1 (prompt S1) Register: economic-cybernetic diagnostic Status: new material; register battery to run before deposit

1. The theorem

Every volume in this series proves, at a different scale and in a different register, a single theorem and its consequences. The theorem can be stated in one sentence, and the sentence should be read slowly, because the series does nothing except unfold it: any record-keeping instrument set over a living, self-renewing substrate is blind to the substrate's couplings, displaced from its timing, late to its accumulations, and draining of its stocks, all four at once and by construction; acting through such an instrument spends the substrate as fuel while booking the expenditure as performance; and the instrument keeps itself in place by choosing its own operators, absorbing its critics, and taking over its corrections.

The sentence has three movements, and each carries a distinct claim. The first movement is a claim about representation: four incapacities, jointly and by construction. The second is a claim about action: the expenditure booked as performance. The third is a claim about persistence: the defence loop. The series argues that the three movements are one configuration, that the configuration recurs at every scale from the single evaluated person to the whole economy, and that its stability is neither an accident of poor design nor a corrigible fault of particular managers, but the working of the configuration itself.

Two terms in the sentence carry the whole weight, and each requires an exact sense.

An instrument is archival, in the sense this series gives the word, when it coordinates by consulting records of past states: entries, balances, scores, positions, ratings, rankings, filed representations of what has already happened, held in a form that can be retrieved, compared, and aggregated. The archive is not a building and not a technology; it is a mode of existence for a coordinating capacity. A capacity exists archivally when it is stored as a representation and exercised by retrieval. The ledger, the appraisal file, the price series, the league table, and the planning model are one instrument in different costumes, and the series will treat them as such.

A substrate is metabolic, in the sense this series gives the word, when its capacities exist only as ongoing reconfiguration: when what the substrate can do is not stored anywhere as a representation but continuously re-enacted, sustained by expenditure, renewed through cycles that have their own tempo, and joined through couplings that no participant can fully state. A working team's capacity to hand a problem across a boundary without a meeting; a firm's capacity to absorb a shock it has not modelled; an economy's capacity to replenish the trust, competence, and shared judgement its transactions consume: these exist metabolically or they do not exist. They are exercised, not retrieved. The ground for this sense of the metabolic is philosophical and is argued at full length elsewhere; within this series it enters as a working definition whose economic consequences the volumes then prove.

The theorem concerns what happens when the first kind of thing is set over the second kind of thing as its regulator.

2. The four incapacities

The first movement of the theorem names four incapacities and insists on two properties of them: they arrive jointly, and they arrive by construction. Neither property is a rhetorical flourish; each states a condition the rest of the section unfolds. Jointly means that no instrument of the archival kind exhibits some of the four without the others; the four are faces of one relation, and a reform that addresses one face while leaving the instrument archival relocates the incapacity rather than removing it. By construction means that the incapacities follow from what an archive is, not from how well a particular archive is built; a better archive exhibits them more smoothly, which the series will show is frequently worse.

The first is blindness to couplings. A record is an entry about a unit: this person, this account, this period, this transaction. The capacities of a metabolic substrate live between the units, in the joinings that no entry is about. Where the value produced by a group of contributors depends on their specific combination, the surplus attributable to the combination as such appears in no contributor's record; the formal core of this result, and its confirmation in information-theoretic terms, occupy a volume of their own. Here it is enough to state the shape: the archive records nodes, the substrate lives in edges, and no refinement of the node records recovers the edges, because the edges are not finer-grained nodes. The blindness is not a gap in the data but the geometry of recording itself.

The second is displacement from timing. An archive is consulted on the archive's schedule: the close of the period, the review, the audit, the mark. A metabolic substrate replenishes on its own schedule, which is set by the tempo of the processes that renew it, and the two schedules do not coincide. Regulation is possible only where the regulator's cycle is at least as fast as the disturbance it must counter, and, symmetrically, a regulator whose action cycle is shorter than the replenishment cycle of the stock it draws on will register the drawing long before it registers the depletion. Ashby's law, that only variety can absorb variety, has a temporal face (Ashby, 1956): a regulator can be adequate in its repertoire and still inadequate in its clock. The archival instrument is inadequate in its clock wherever the substrate's renewal outruns the reporting period, which is to say, wherever the substrate is alive.

The third is belatedness to accumulations. What an archive holds is what has already been entered. The substrate's accumulations, of fatigue, of thinning trust, of competence quietly leaving, of slack consumed, build continuously and enter the record only at the threshold where they become an event: the resignation, the defect, the failed audit, the quarter that finally misses. Until the threshold, the instrument reads the accumulation as nothing, and frequently reads the drawing-down that produces it as improvement, since the drawing-down releases resources the record can see. The instrument is therefore structurally late, and its lateness is not a lag that faster reporting shortens, because the accumulation is not in the record at any speed; only its threshold crossing is.

The fourth is extraction, and it composes the other three. An instrument blind to couplings, displaced from timing, and late to accumulations, when used to allocate, rewards whatever moves its entries and cannot price whatever sustains the substrate. Allocation through such an instrument therefore transfers effort, attention, and

resource from the sustaining activities it cannot see toward the recordable activities it can, and the transfer is experienced, inside the archive, as efficiency. The substrate is spent, the entries improve, and the spending appears nowhere as a cost, because the stock being spent was never an entry. This is the sense in which the improvement can be the consumption, and the series will meet the pattern at every scale: the person whose recorded output rises as their unrecorded absorbing of others' friction is run down; the team whose velocity improves as the tacit couplings that made it a team are individually reallocated; the firm whose returns strengthen as the reserves that would have carried it through the unmodelled disturbance are optimised away; the economy whose measured growth is, in identifiable sectors, the liquidation of stocks whose replenishment cycles exceed every instrument that prices them.

The observation that a measure used as a target ceases to be a good measure has been in the literature for half a century (Goodhart, 1984; Campbell, 1979), and its case record is now extensive (Muller, 2018; Espeland and Sauder, 2007). The theorem is not that observation restated. The observation concerns what happens to a measure under pressure; the theorem concerns what a class of instrument can and cannot carry in principle, pressure or no pressure, and derives the gaming, the reactivity, and the drift as consequences rather than causes. The difference matters for the remedy. If the fault is pressure on the measure, better measures and better governance are the cure, and the literature of measurement reform follows. If the fault is the mode of existence of the instrument, then better measures raise the resolution of the blindness, and the series will argue, with the reform record in evidence, that this is what the history of measurement reform shows.

3. The expenditure booked as performance

The second movement of the theorem turns from representation to action, and it is here that the argument leaves the terrain of measurement critique entirely. An instrument that only misdescribed its substrate would be an epistemic problem. The archival instrument is an allocative one: budgets, promotions, prices, headcounts, and strategies move through it, and so its four incapacities are not errors in a picture but forces in the world. What the instrument cannot see, it cannot protect; what it can see, it amplifies; and the difference between the two is booked, period after period, as performance.

The booking is the point. The configuration does not hide the expenditure; it records the expenditure, in the only vocabulary it has, as gain. The substrate's stocks do not appear on the instrument as stocks, so their liquidation cannot appear as liquidation; it appears as margin, as productivity, as growth. The person who has learned to read the configuration whole will recognise the same entry from both sides: from inside the instrument, a number improving; from inside the substrate, a capacity being spent that the number will never name. The series is, among other things, a training in reading that entry from both sides at once.

Two clarifications guard the claim from its caricatures. The theorem does not say that records are useless, that measurement is illegitimate, or that organisations should navigate blind. Archives are indispensable for exactly what they are: instruments over the recordable. The theorem locates a boundary, not a prohibition; it says what kind of thing an archive can regulate and what kind of thing it cannot, and it diagnoses the specific configuration in which the second kind of thing is governed as if it were the first. Nor does the theorem say that the operators of the instrument are fools or villains. The third movement says something stronger and colder: that the configuration selects, retains, and promotes the operators for whom the instrument's picture is most nearly complete, so that sincerity and the configuration's persistence are allies rather than opposites. The people are not the explanation. The people are among the consequences.

4. The defence loop

The third movement of the theorem answers the question every diagnosis must face: if the configuration consumes the substrate it depends on, why does it persist? Why is it not corrected by its own results? The answer occupies a volume, and its shape can be given here in three clauses.

The instrument chooses its own operators. Advancement inside the configuration runs through the instrument; those advanced are those whose activity the instrument reads well and who read the instrument well; and so, over cycles, the population of operators converges on those calibrated to the record rather than to the substrate, without any operator needing to intend the convergence. Every retained operator tends toward a model of the regulator, and the substrate is what that model does not contain.

The instrument absorbs its critics. A correction addressed to the configuration must be expressed, to be heard, in the configuration's vocabulary; expressed in that vocabulary, it becomes a proposal about entries, measures, and targets; and a proposal about entries is something the instrument can adopt without ceasing to be what it is. The correction is thereby re-emitted as the instrument's own output, its structure intact, and the critic's success is the mechanism of the critique's neutralisation.

The instrument takes over its corrections. Where a practice arises outside the record that sustains the substrate, the configuration's response, sooner or later, is to record it: to name it, measure it, target it, and thereby convert it into one more entry whose pursuit exhibits the four incapacities. The correction is not defeated. It is expropriated, and its expropriation extends the instrument's reach into the very territory the correction had marked as needing protection from it.

The three clauses close a loop, and the loop is why the series speaks of a configuration rather than a mistake. A mistake is corrected by its consequences. A configuration of this kind converts its consequences into confirmations, and it does so at every scale the series examines, which is why the series must examine every scale.

5. The four movements of the series

The work travels through four movements, and their order follows the argument's own dependencies rather than any disciplinary map.

The first movement concerns the act of capture. It establishes the foundations: that the regress of oversight has no terminus, since every level of regulation added to include the observer produces a new observer to include; that the attempt to total a living process from inside it meets a limit that is formal and not practical; and that the limit converges from independent directions onto the kernel stated above. This movement proves that the configuration's incapacities are of the kind that construction entails, which is what licenses every later use of the phrase by construction.

The second movement concerns the organisation. It takes the kernel into the firm: the surplus of combination that no individual record can carry; the absorbing of the unrecorded by identifiable people whose absorbing is their invisibility; the cascade by which extraction travels down a hierarchy of instruments; and the selection loop by which the instrument staffs itself. This movement proves that the configuration is self-sustaining at the scale where most readers will first recognise it, which is the scale of their own working week.

The third movement concerns the economy. It composes the organisational machinery upward into the conditions under which decentralised coordination through prices can do what its defenders claim for it, and it tests the two dominant contemporary architectures against those conditions, on their own records and in their own vocabularies. This movement proves that the configuration is not a pathology of one system and the health of another; it is a diagnosis that runs across the architectures that officially oppose each other.

The fourth movement concerns the response. It asks what an intervention would have to be, given everything the prior movements have established: why elimination without transformation relocates the harm it targets; what kind of boundary a living configuration requires, a membrane that regulates passage rather than a wall that forbids it; and what institutional forms could hold a correction open against the absorbing and the takeover the third clause of the theorem describes. This movement does not promise that the response succeeds. It specifies what the response must survive.

6. The reader's contract

A series of this length asks something of its reader, and it is fair to state the terms at the door.

The series asks the reader to hold one sentence across four movements, and to test every chapter against it. Each volume argues from first principles in its own register; none leans on the others for its proof; and yet each is the same theorem at a different scale, so that a reader who finds the argument failing anywhere has found something the whole must answer. The series invites that reading. The theorem's clauses are stated so that they can be wrong: an archival instrument that priced a coupling, a reform that was not reabsorbed, a selection process that converged on substrate-calibrated operators, would each be a counterexample, and the volumes name, in their own terms, what such counterexamples would look like.

The series asks the reader, second, to keep two vocabularies alive at once: the vocabulary of the instrument, in which the events of the coming chapters appear as improvement, and the vocabulary of the substrate, in which the same events appear as expenditure. Neither vocabulary is dismissed for the other. The argument's force lies in the reader's growing capacity to translate between them without the translation being provided, until the double reading becomes a habit that survives the book.

And the series asks the reader, last, to defer one question until the end: the question of what to do. The temptation to reach for remedies early is the configuration speaking, since the configuration's deepest habit is to convert every diagnosis into a programme of entries. The fourth movement earns the question. Until then, the work is colder: to see one configuration, at four scales, spending what it cannot name, and calling the spending by its own name for the first time.

References

Ashby, W.R. (1956) *An Introduction to Cybernetics*. London: Chapman and Hall.

Campbell, D.T. (1979) 'Assessing the impact of planned social change', *Evaluation and Program Planning*, 2(1), pp. 67 to 90.

Conant, R.C. and Ashby, W.R. (1970) 'Every good regulator of a system must be a model of that system', *International Journal of Systems Science*, 1(2), pp. 89 to 97.

Espeland, W.N. and Sauder, M. (2007) 'Rankings and reactivity: how public measures recreate social worlds', *American Journal of Sociology*, 113(1), pp. 1 to 40.

Goodhart, C.A.E. (1984) *Monetary Theory and Practice: The UK Experience*. London: Macmillan.

Muller, J.Z. (2018) *The Tyranny of Metrics*. Princeton: Princeton University Press.

The Post-Cybernetic Organisation: Paper 3 - The Dissolution of Orders

Author: Nourizadeh, Moreno ORCID: 0009-0006-0174-2585 Version: v1_0

Abstract

This paper argues that the proliferation of cybernetic orders - first, second, third, and beyond - is not a progressive solution to the observer problem but a structural symptom of it. The possibility of counting orders presupposes precisely the external standpoint that inclusion was meant to dissolve: each additional order reinstates the exteriority it purports to overcome. Escape requires not a further order but a different ontology - one in which the observer-observed relation is chiasmic rather than hierarchical, reversible rather than embeddable. The paper proceeds through three movements. First, it traces the achievements and retained limitations of first-order and second-order cybernetics, demonstrating that despite von Foerster's constructivist turn and Maturana and Varela's autopoietic reconceptualisation, four foundational assumptions remain structurally intact: the informational frame, sequential temporality, absent metabolism, and the observer-as-processor. Second, it maps eleven paradigmatic seams - precise points where the mechanistic paradigm encounters phenomena it cannot accommodate - across physics (measurement, time's arrow, non-locality, vacuum energy), mathematics (the continuum, incompleteness), cognitive science (the hard problem, binding, the frame problem), and computation (halting, symbol grounding). The convergence of these seams on a single structural crack - the impossibility of deriving holistic, relational, global phenomena from discrete, separable, local assumptions - establishes that the cybernetic impasse is not domain-specific but paradigmatic. Third, through close engagement with Merleau-Ponty's late ontology of the flesh and the chiasm, the paper articulates the structure of reversibility that replaces reflexive hierarchy: not levels but folds, not embedding but dehiscence, not exteriority but the constitutive gap. Six ontological requirements for post-cybernetic existence are derived: chiasmic reversibility, metabolic cognition, ecstatic temporality, operational closure, structural coupling, and the viability constraint. The dissolution of orders opens the space that the Membrane Architecture, developed in Paper 4, inhabits.

Keywords: cybernetics, phenomenology, chiasm, paradigmatic seams, Merleau-Ponty, organisational ontology, incalculability, second-order cybernetics, flesh, dehiscence, viability constraint

1. Introduction: The Question of Orders

1.1 The Proliferation of Reflexivity

The history of cybernetics narrates itself as a sequence of reflexive incorporations. First-order cybernetics positioned the observer outside the system under investigation; second-order cybernetics recognised this position as untenable and included the observer within the observed. The question that has haunted cybernetic theory since von Foerster's (1981) articulation of this inclusion is whether additional orders are necessary - whether a third-order cybernetics must observe the observer observing, whether a fourth-order must observe that observation, and so on toward an indefinite regress of meta-positions.

This paper argues that the question is malformed. The possibility of counting orders presupposes precisely what the inclusion of the observer was meant to problematise: an external position from which the orders can be enumerated. If the observer is genuinely included in the observed, there is no standpoint from which to count how many levels of inclusion have occurred. The counting itself reinstates the externality that inclusion was supposed to dissolve.

The proliferation of orders is thus a symptom rather than a solution. Each additional order attempts to capture what the previous order left outside, but the attempt itself leaves something outside - the position from which the new order is articulated. Bateson (1972, pp. 279-308) recognised this structure in his analysis of logical types: the attempt to include a level generates a meta-level that demands inclusion, generating a further meta-level, without terminus. The regress is a structural feature, not an insufficiency of rigour. It belongs of reflexive systems that retain the subject-object framework even while attempting to overcome it.

The argument of this paper is that escape from the regress requires not another order but a different ontology - one in which the observer-observed relation is not hierarchical but chiasmic, not a matter of levels but of reversibility. This ontology has been articulated in the phenomenological tradition, particularly in Merleau-Ponty's (1968) late work on "the flesh" (*la chair*) and the "chiasm" (*le chiasme*). The chiasmic structure is one of mutual implication without hierarchy: the touching hand is touched by what it touches; the seeing eye is visible to what it sees; the observer is observed by what it observes. There is no external position because externality and internality are themselves chiasmic - each becomes the other depending on the direction of attention.

Post-cybernetic phenomenology names the dissolution of the order-counting frame through the recognition of chiasmic structure. The prefix "post-" indicates not temporal succession (what comes after cybernetics in history) but structural supersession (what the cybernetic frame becomes when its enabling assumptions are exposed and transformed). The phenomenological method provides the means of exposure; the chiasmic ontology provides the transformed framework.

1.2 Methodological Ground

The method employed in this paper is phenomenological in the tradition established by Husserl (1913), radicalised by Heidegger (1927), and developed through embodiment by Merleau-Ponty (1945, 1968). Phenomenology investigates the structures of experience and existence through disciplined attention to how phenomena appear (Moran, 2000, pp. 1-22). The key methodological moves are:

The epoché: Husserl (1913, pp. 51-62) introduced the epoché or "bracketing" as the suspension of the "natural attitude" that takes the existence of the world for granted. The epoché does not deny that the world exists; it suspends the thesis of existence in order to investigate how existence is constituted for consciousness. In the present context, the epoché suspends the assumption that cybernetic categories - information, feedback, variety, control - transparently describe what is. This suspension opens the question: what must be the case for these categories to function as they do? What do they presuppose? What do they exclude?

Eidetic variation: Husserl (1913, pp. 156-173) developed the method of eidetic variation to identify essential structures. By imaginatively varying features of a phenomenon, one identifies which features can be changed while the phenomenon remains what it is (contingent features) and which features cannot be changed without the phenomenon ceasing to be what it is (essential features). Applied to cybernetic systems: what features of feedback, autopoiesis, observation can be varied? What happens at the limits of variation? Where does the cybernetic description break down, and what does the breakdown reveal?

Existential analysis: Heidegger (1927) transformed phenomenology by showing that the structures of experience are grounded in the structures of existence. Dasein - the being for whom being is a question - exists temporally, projecting toward possibilities while thrown from a past it did not choose. The existential structures (care, being-toward-death, thrownness, projection) are not psychological states but ontological conditions. Applied to cybernetics: the observer is not a neutral information-processor but an existing being whose observation is shaped by care, finitude, and temporal ekstasis.

Carnal analysis: Merleau-Ponty (1945, 1968) further developed phenomenology by demonstrating the primacy of embodiment. The body is not an object among objects but the condition of possibility for there being objects at all. Perception is not the passive reception of data but the active exploration of a world by a body that is simultaneously sensing and sensed. The late concept of "flesh" (Merleau-Ponty, 1968, pp. 130-155) generalises this

insight: flesh is the reversible tissue of being, neither subject nor object but the medium in which both emerge through differentiation. Applied to cybernetics: the observer is not a disembodied information-processor but a carnal being whose observation is a mode of bodily engagement with what is observed.

These methods converge on a critique of the cybernetic paradigm from within. By attending to the constitutive structures that cybernetic theory employs without thematising, phenomenology reveals both what the paradigm achieves and where it reaches limits that it cannot transgress without ceasing to be cybernetic.

1.3 Plan of the Paper

The paper proceeds as follows. Part I (the present section) introduces the problematic and establishes the methodological framework. Part II examines first-order and second-order cybernetics in detail, identifying their achievements and their retained limitations. Part III maps the "paradigmatic seams" where the mechanistic paradigm - of which cybernetics is the most sophisticated expression - encounters phenomena it cannot accommodate. Part IV articulates the chiasmic turn: the move from reflexive hierarchy to carnal reversibility. Part V specifies the six ontological requirements for post-cybernetic systems and demonstrates their formal structure. Part VI grounds the practical phenomenology of organisational navigation in the Membrane Architecture.

2. First-Order Cybernetics: The Observer Outside

2.1 The Founding Concepts

First-order cybernetics emerged from the convergence of multiple disciplines during and after the Second World War. The Macy Conferences (1946-1953) brought together mathematicians (Wiener, von Neumann), neurophysiologists (McCulloch, Pitts), anthropologists (Bateson, Mead), psychologists (Lewin), and others to address problems of "circular causal and feedback mechanisms in biological and social systems" (Heims, 1991, pp. 14-41). The resulting framework - articulated most systematically in Wiener's *Cybernetics* (1948), Shannon and Weaver's *The Mathematical Theory of Communication* (1949), and Ashby's *An Introduction to Cybernetics* (1956) - provided a common language for analysing regulation and control across domains.

The founding concepts of first-order cybernetics require precise statement, for the limitations of the paradigm are already implicit in these definitions.

Information received its canonical formulation in Shannon's (1948) mathematical theory. Shannon defined information as the reduction of uncertainty, measured in bits. A message conveys one bit of information if it selects between two equally probable alternatives; it conveys n bits if it selects among 2^n equally probable alternatives. The entropy of a source measures its average information content: $H = -\sum p(i) \log_2 p(i)$, where $p(i)$ is the probability of the i -th message. Shannon explicitly excluded meaning from his framework: "the semantic aspects of communication are irrelevant to the engineering problem" (Shannon & Weaver, 1949, p. 31). Information is a statistical property of ensembles of messages, not a semantic property of individual messages.

This exclusion of meaning was methodologically productive - it enabled the development of coding theorems, channel capacity analysis, and the entire field of information engineering. But it was also an ontological commitment with consequences. If information is defined statistically, then anything that processes information is characterised by its statistical properties. The mind, if it is an information-processor, is characterised by input-output relations, not by the meaning of what it processes. The exclusion of semantics was not provisional - something to be added later - but structural. Shannon's information is intrinsically meaningless; meaning would have to be added from outside the framework, by an interpreter who is not part of the informational description.

Feedback was defined by Wiener (1948, pp. 96-97) as "the control of a machine on the basis of its actual performance rather than its expected performance." The concept has antecedents in engineering (the centrifugal

governor of steam engines, analysed by Maxwell in 1868) and physiology (Bernard's concept of the milieu intérieur, Cannon's homeostasis). Wiener's contribution was to generalise the concept: any system that uses information about its output to modify its subsequent behaviour exhibits feedback. Negative feedback reduces deviation from a target state; positive feedback amplifies deviation.

The feedback concept presupposes several things. It presupposes that the system and its environment are distinguishable - that there is an "inside" that processes and an "outside" that provides inputs and receives outputs. It presupposes that information flows in discrete channels that can be identified and measured. It presupposes that time is sequential: the system receives input, processes it, produces output, receives feedback, and the cycle repeats. These presuppositions are not arbitrary; they are the conditions under which feedback analysis becomes possible. But they are also limitations. Systems for which inside and outside are not clearly distinguishable, or for which information does not flow in discrete channels, or for which time is not simply sequential, will not fit the feedback framework.

Variety was introduced by Ashby (1956, p. 126) as "the number of distinct elements" in a set - equivalently, the logarithm (base 2) of that number, yielding variety measured in bits. Ashby's Law of Requisite Variety states that "only variety can destroy variety" (Ashby, 1956, p. 207): a regulator must be able to produce at least as many distinct responses as there are distinct disturbances to be regulated. If the environment can perturb the system in n distinct ways, the regulator must be able to respond in at least n distinct ways to maintain stability.

The Law of Requisite Variety has been enormously influential in organisational theory and management science (Beer, 1972, 1979, 1985). It provides a formal grounding for the intuition that complex environments require complex organisations. But the concept of variety, like the concept of information, presupposes that states are discrete and countable. Continuous variation must be discretised before variety can be measured. The discretisation is not innocent: it imposes a grid on phenomena that may not admit of gridding. What counts as a "distinct" state depends on the observer's purposes and capacities, not on the phenomenon itself. Variety is observer-relative, though Ashby's formulation presents it as an objective property of systems.

Homeostasis, adopted from Cannon's (1932) physiology, denotes the maintenance of essential variables within limits compatible with the system's continued existence. Ashby (1960) generalised this to the concept of "ultrastability": a system that can change its own structure in response to environmental changes that would otherwise drive essential variables outside viable limits. The ultrastable system does not merely regulate; it adapts - it finds new equilibria when old equilibria become unattainable.

Homeostasis presupposes that there are "essential variables" whose values determine whether the system continues to exist. This presupposition is less innocent than it appears. For a thermostat, the essential variable (temperature) is externally specified by the designer. For a living organism, what counts as essential is determined by the organism's own organisation - but this introduces a circularity that first-order cybernetics cannot fully accommodate. The organism maintains itself, but the "self" that is maintained is constituted through the maintaining. This circularity points toward autopoiesis, which second-order cybernetics would develop.

2.2 The Observer's Position

First-order cybernetics positions the observer outside the system observed. This positioning is explicit in Ashby's formulation: "In this book the word 'system' will be used for 'that part of reality that is singled out for discussion'" (Ashby, 1956, p. 16). The singling-out is performed by an observer who stands apart from what is singled out. The observer draws the boundary between system and environment, identifies the relevant variables, constructs the model, and tests the model's predictions against the system's behaviour.

This external position enables powerful analytical techniques. The observer can measure the system's inputs and outputs, calculate its transfer function, determine its stability properties, and design interventions to achieve desired behaviours. Control engineering, operations research, and systems analysis all depend on this external position. The engineer who designs a feedback controller for a chemical plant does not need to be part of the plant; the analyst who models an organisation's information flows does not need to be part of the organisation.

The externality of the observer is what makes the analysis objective - the same results would be obtained by any competent observer using the same methods.

But the external position encounters difficulties when the system under investigation includes observers - when the object of cybernetic analysis is itself a cybernetic system. If the analyst models an organisation, and the organisation includes people who are themselves modelling their environment (including, potentially, the analyst), then the analyst is not simply external to what is analysed. The organisation's behaviour depends on how its members model their situation; if the analyst's model becomes known to the organisation's members, their behaviour may change, invalidating the model. This is the problem of reflexivity that motivated second-order cybernetics.

More fundamentally, the external position presupposes that the observer can be characterised independently of the observation. The observer has access to the system through measurement; measurement yields data; data support models; models predict behaviour. But how does measurement work? The observer must interact with the system to obtain data - must perturb it, sample it, record its states. The interaction is itself a physical process governed by physical laws. If those laws are the same laws that govern the system, then the observer is not external to the domain in which the system exists. The observer is part of the same physical world, subject to the same physical constraints. The externality is methodological, not ontological.

Quantum mechanics made this point forcefully. The measurement problem - the apparent collapse of the wave function upon measurement - reveals that the observer cannot be excluded from the physical description (von Neumann, 1932/1955; Wheeler & Zurek, 1983). The Schrödinger equation describes deterministic, reversible evolution; measurement introduces probabilistic, irreversible transitions. The formalism requires both, but it cannot specify where the boundary between them lies. The "Heisenberg cut" between quantum system and classical apparatus is necessary for predictions but unjustifiable within the theory (Bell, 1990). The observer is methodologically necessary and ontologically problematic.

First-order cybernetics did not resolve this problem; it bracketed it. The systems under investigation - thermostats, servomechanisms, homeostatic organisms - were treated as classical objects admitting of classical measurement. The observer's physical nature was not thematised. This bracketing was productive: it enabled the development of control theory, information theory, and systems analysis. But it was also a limitation. When the bracketing could no longer be sustained - when the observer's inclusion became unavoidable - the transition to second-order cybernetics occurred.

2.3 The Achievements of First-Order Cybernetics

Before proceeding to second-order cybernetics, the achievements of the first-order paradigm deserve acknowledgment. These achievements were not negated by the transition to second-order; they were preserved and contextualised.

Feedback analysis illuminated a vast range of phenomena. The concept of negative feedback - the use of error signals to reduce deviation from a target - unified the analysis of physiological regulation (Cannon, 1932), engineering control (Wiener, 1948), and economic stabilisation (Tustin, 1953). Positive feedback explained explosive growth, runaway processes, and vicious circles. The distinction between proportional, integral, and derivative control provided tools for designing regulators with specified dynamic properties. Stability analysis - determining whether a feedback system would converge to equilibrium or oscillate - drew on complex analysis and Laplace transforms to yield general results applicable across domains.

Information theory grounded communication engineering and, eventually, molecular biology. Shannon's coding theorems established fundamental limits: the channel capacity is the maximum rate at which information can be transmitted with arbitrarily low error probability (Shannon, 1948). These theorems guided the design of communication systems from telephone networks to deep-space probes. The recognition that DNA encodes genetic information - that the sequence of nucleotides specifies the sequence of amino acids in proteins - was explicitly informed by information-theoretic concepts (Yockey, 2005), though the appropriateness of this

application remains contested (Godfrey-Smith, 2007).

Systems analysis provided a framework for understanding complex organisations. The distinctions between open and closed systems (von Bertalanffy, 1968), between system and environment, between structure and function, enabled systematic inquiry into phenomena that had previously seemed too complex for scientific treatment. Operations research applied these concepts to military logistics, industrial management, and public policy (Ackoff & Sasieni, 1968). The viable system model (Beer, 1972, 1979, 1985) provided an architecture for organisational design based on recursive application of cybernetic principles.

These achievements were genuine. They enabled technological developments (automated control systems, digital communication, organisational restructuring) and theoretical insights (the formal analogy between biological and engineered systems, the quantification of information, the analysis of stability and adaptation). The transition to second-order cybernetics did not repudiate these achievements; it revealed their conditions and limitations.

3. Second-Order Cybernetics: The Observer Included

3.1 The Reflexive Turn

The transition to second-order cybernetics occurred through the recognition that first-order cybernetics, consistently applied, must include the observer who applies it. If cybernetics is the science of regulation and control in systems, and if observers are themselves systems capable of regulation and control, then cybernetics must apply to observers - including the observers who do cybernetics. The first-order position - observer outside the system - is itself a cybernetic phenomenon, a mode of system-environment coupling that demands cybernetic analysis.

Von Foerster (1981) articulated the transition with characteristic precision. The title of his collected papers - *Observing Systems* - is deliberately ambiguous: it refers both to systems that observe and to the activity of observing systems. The ambiguity is the point. Second-order cybernetics studies systems that observe, including the system (the cybernetician) that observes them observing. The observer is part of the observed; the observation is part of the observation.

Von Foerster's argument proceeded through analysis of perception. The nervous system, he demonstrated, does not transmit information about the external world (von Foerster, 1981, pp. 258-271). Sensory receptors respond to energy gradients - photons striking the retina, pressure waves reaching the cochlea, molecules binding to olfactory receptors. These responses are neural signals that encode the intensity and location of stimulation but not its environmental source. The colour red is not transmitted from the apple to the brain; the brain constructs redness from patterns of neural activity that could, in principle, be produced by direct stimulation of the optic nerve. "Out there," von Foerster concluded, "there is no light and no colour, there are only electromagnetic waves; 'out there' there is no sound and no music, there are only periodic variations of the air pressure" (von Foerster, 1973, p. 214).

This constructivist epistemology has radical implications. If perception is construction rather than reception, then the perceived world is not a copy of a pre-existing reality but an invention of the perceiving system. "The environment as we perceive it is our invention" (von Foerster, 1973, p. 214). This is not solipsism - the claim that nothing exists beyond perception. It is the recognition that what exists for a perceiving system is constituted through the system's perceptual operations. The world constrains perception (not all constructions are viable), but it does not determine it (the same constraints admit multiple constructions).

Von Foerster drew ethical as well as epistemological conclusions. If we construct our worlds, we are responsible for what we construct. We cannot appeal to an objective reality that compels us to perceive and act as we do. The first-order cybernetician could claim objectivity - to describe how systems "really" work. The

second-order cybernetician must acknowledge participation - to recognise that the description is itself a construction shaped by the observer's purposes, frameworks, and limitations.

3.2 Autopoiesis and Operational Closure

The concept of autopoiesis, introduced by Maturana and Varela (1980), provided second-order cybernetics with its most influential theoretical construct. An autopoietic system is one that produces the components that constitute it, through a network of processes that is itself constituted by those components. The canonical example is the living cell: the membrane that bounds the cell is produced by the metabolic processes that the membrane makes possible; the metabolic processes are constituted by the enzymes and substrates that the processes produce; the whole maintains itself through continuous production of itself.

Maturana and Varela defined autopoiesis precisely:

An autopoietic machine is a machine organized (defined as a unity) as a network of processes of production (transformation and destruction) of components that produces the components which: (i) through their interactions and transformations continuously regenerate and realize the network of processes (relations) that produced them; and (ii) constitute it (the machine) as a concrete unity in the space in which they (the components) exist by specifying the topological domain of its realization as such a network. (Maturana & Varela, 1980, p. 78)

The definition emphasises organisation over structure. Structure refers to the actual components and their relations at a given moment; organisation refers to the invariant relations that define the system as the kind of system it is. A cell's structure changes continuously as molecules are synthesised and degraded; its organisation remains constant as long as the cell lives. Death is the loss of autopoietic organisation - the point at which the network of processes no longer produces the components that sustain it.

Autopoiesis entails operational closure: the operations of the system refer only to the system's own states. The nervous system, Maturana argued, is operationally closed: neural activity causes neural activity; there is no point at which "external information" enters and is processed. What we call perception is the perturbation of the nervous system's ongoing activity by environmental events; what the nervous system does with the perturbation is determined by its own structure, not by the perturbation's environmental source. "The nervous system does not 'pick up information' from the environment, as we often hear. On the contrary, it brings forth a world" (Maturana & Varela, 1987, p. 169).

Operational closure does not mean causal isolation. The autopoietic system is structurally coupled to its environment: perturbations from the environment trigger structural changes in the system, and the system's activities perturb the environment. But the perturbations are not instructions; they do not specify what the system does. The system's response is determined by its own organisation. Structural coupling is mutual perturbation, not information transmission.

This reconceptualisation has profound implications for the concept of information. In first-order cybernetics, information exists in signals that flow from source to receiver. In autopoietic theory, "there is no 'transmitted information' in communication" (Maturana & Varela, 1987, p. 196). Communication is the coordination of behaviour among structurally coupled organisms; the coordination arises through mutual perturbation, not through the transfer of content. Information is not a thing that exists independently of observers; it is a distinction that an observer makes about the relation between a system and its environment.

3.3 The Conversation and the Observer of Observers

Pask's (1975, 1976) conversation theory extended second-order cybernetics into the domain of learning and interaction. Pask argued that cognition is not a property of isolated systems but emerges through conversation - the exchange of understandings between participants. A conversation is not merely the transmission of information but the construction of shared meaning through recursive interaction.

The basic unit of conversation theory is the concept. A concept is not a mental representation but a procedure - a way of doing something that can be demonstrated and imitated. To have a concept is to be able to perform the procedure and to recognise performances of it by others. Concepts are public, not private: they exist in the space of interaction, not in the heads of individuals.

Learning occurs when the conversation produces new concepts - when the interaction generates procedures that neither participant possessed beforehand. This is emergent understanding: the conversation as a whole knows more than any participant knows individually. The conversation is itself a cognitive system, irreducible to its participants.

Pask's framework positions the observer of observers. The conversation theorist observes conversations between participants who are themselves observing (each other, their environment, themselves). The observation of observation is not a meta-level above the conversation but a participant position within a larger conversation. There is no final meta-level, no ultimate observer who observes without being observed. Every observation is itself observable; every observer can become observed.

This infinite embedding might seem to generate regress, but Pask dissolves the regress through the concept of self-reference. The conversation can include itself as a topic; the participants can discuss their discussing. Self-reference is not paradoxical but productive: it enables the conversation to develop, to become aware of its own presuppositions, to change its own structure. The regress is not vicious but generative.

3.4 The Retained Commitments

Despite the radical reconceptualisations of second-order cybernetics - the inclusion of the observer, the constructivist epistemology, autopoietic organisation, conversational cognition - certain foundational commitments of the cybernetic paradigm were retained. These retained commitments constitute the limits that post-cybernetic phenomenology must transgress.

The informational frame. Second-order cybernetics reconceived information as observer-dependent construction rather than objective content. But the framework remained informational in structure. The autopoietic system produces components; the production is described in terms of processes that can be formally specified. Maturana's "bringing forth a world" is a computational metaphor: the nervous system computes its world from its own activity. Pask's concepts are procedures - algorithms that can be executed. The semantic shift from "information transmitted" to "information constructed" does not escape the informational paradigm; it internalises it. The observer constructs information rather than receiving it, but information remains the medium of description.

This commitment becomes visible when autopoietic theory is applied to artificial systems. Varela (1979) framed autopoiesis as a formal notion applicable to any system with the requisite organisational closure, including potentially computational systems. The criterion is organisational: does the system produce its own components through a network that those components constitute? If a computational system could be designed to meet this criterion, it would be autopoietic. The question of whether the substrate matters - whether silicon can be autopoietic in the way carbon is - is left open or answered affirmatively. This implies that what matters is the organisation, not the flesh.

The absent metabolism. Autopoietic theory emphasises that living systems are material and energetic: they exchange matter and energy with their environments while remaining organisationally closed. But the metabolic dimension - the transformation of the system through its own activity - receives less attention than the

organisational dimension. The cell maintains its organisation by producing its components; the production consumes energy and substrates. But the theoretical focus is on the maintenance, not on the cost. What it means for cognition to be metabolically expensive - to literally consume the organism's resources - is not thematised.

Jonas (1966) had articulated this dimension two decades before autopoietic theory. For Jonas, metabolism is not incidental to life but constitutive: "living things are creatures of need" (Jonas, 1966, p. 80). The organism must continuously acquire matter and energy to maintain itself; it exists in a condition of precarious dependence on what is other than itself. This neediness is the ground of concern, of caring about outcomes. A system that did not need to maintain itself through continuous effort would not care; outcomes would not matter. The indifference of computation - its execution of the same algorithm regardless of consequences - reveals the absence of metabolic existence.

The sequential temporality. Second-order cybernetics analysed self-reference: the observer observing itself, the conversation including itself, the autopoietic system producing itself. Self-reference involves time: the system at t_1 refers to itself at t_0 or anticipates itself at t_2 . But the temporality remains sequential. States succeed states; the future emerges from the past through lawful transformation. The recursive structure is temporal: the system that produces itself at t_1 is the product of itself at t_0 .

This sequential temporality contrasts with what Heidegger (1927) called ecstatic temporality. For Heidegger, Dasein does not exist at a point in time with past behind and future ahead. Dasein is stretched across time, existing as thrown projection - coming from a past it did not choose, going toward a future it must create. The three temporal ecstases (past, present, future) are not successive moments but co-present dimensions of a single temporal stretch. The future is not what will happen but what Dasein is oriented toward; the past is not what did happen but what Dasein has been and still is (Gewesenheit). This ecstatic structure cannot be captured by sequential models in which t_1 follows t_0 and precedes t_2 .

The observer as information-processor. The most fundamental retained commitment is the conception of the observer as an information-processing system. Second-order cybernetics included the observer, but the included observer is still characterised informationally. The observer constructs information, operates on information, communicates through information. The observer has organisation that can be formally specified - autopoietic organisation, conversational procedures, structural coupling with environment. The observer is a very sophisticated system, but a system nonetheless.

What escapes this characterisation is the observer as flesh - as the lived body that perceives, moves, feels, suffers. The flesh is not an information-processing system. It is the medium in which information-processing occurs, the condition of possibility for there being systems at all. Merleau-Ponty's (1968) analysis of the chiasm reveals that the perceiver is perceived, that the toucher is touched, that the seer is visible. This reversibility cannot be captured by including the observer as another system within the systemic framework. The flesh is not a system; it is what makes systems possible.

4. The Exhaustion of Reflexivity

4.1 The Regress Problem

The transition from first-order to second-order cybernetics was driven by the recognition that the observer cannot stand outside the system observed. The observer is a system; cybernetics applies to systems; therefore cybernetics applies to observers. The move is logically compelling. But it generates a difficulty that second-order cybernetics acknowledged without resolving.

If the second-order cybernetician observes systems that observe, the second-order cybernetician is a system that observes systems that observe. A third-order cybernetician might observe this observing of observing. A fourth-order might observe that. The regress extends indefinitely. Each level of observation becomes the object of

a higher-level observation, which becomes the object of a yet higher-level observation.

Von Foerster (1981, pp. 288-309) acknowledged the regress and proposed recursive computation as its formal structure. The eigenvalue of a recursive function is a fixed point: a value that the function maps to itself. Visual perception, von Foerster suggested, computes eigenvalues: stable structures that are maintained through recursive neural processing. The regress is halted not by reaching an ultimate observer but by achieving computational stability - by finding fixed points in the recursive dynamics.

But this solution relocates rather than dissolves the problem. The eigenvalue analysis assumes computational structure: the recursive function that has fixed points. The analysis applies to the observed, but what about the observer who performs the analysis? The eigenvalue theorist observes systems that compute eigenvalues; the eigenvalue theorist's observation is itself a computation. The regress continues at the level of the analyst, not merely at the level of the analysed.

Bateson (1972, pp. 279-308) framed the regress in terms of logical types. A class cannot be a member of itself; a statement about statements is of a different logical type than the statements it is about. The regress of observation levels is a regress of logical types: observation₁, observation of observation₁ (which is observation₂), observation of observation₂ (which is observation₃), and so on. Each level is of a different logical type than the level it observes.

Russell's theory of types was designed to block paradox: the class of all classes that do not contain themselves cannot be formed, because the self-reference crosses types illegitimately. But the regress of observation levels is not paradoxical; it is well-formed at each level. The regress does not generate contradiction; it generates infinity. The theory of types explains why the regress does not collapse into paradox, but it does not explain where or why the regress stops.

The regress problem reveals that reflexivity is insufficient to overcome the limitations of the cybernetic paradigm. Including the observer generates a new observer to be included; including that observer generates another. The structure of inclusion reproduces the structure of exclusion at a higher level. What is needed is not another inclusion but a different relation between observer and observed - a relation that does not admit of levels because it does not presuppose the externality that generates levels.

4.2 The Persistence of Exteriority

The regress problem is symptomatic of a deeper difficulty: the persistence of exteriority within the framework that claims to overcome it. Second-order cybernetics includes the observer, but the inclusion is performed from a position that is not itself included. The second-order cybernetician describes systems that observe; the description is an observation; but the description claims validity that transcends any particular observer's construction.

Consider Maturana and Varela's account of autopoiesis. They describe the autopoietic organisation of living systems; they claim that this organisation is what makes systems living; they apply the concept to cells, organisms, and nervous systems. The description is presented as true - as capturing what living systems really are. But if all cognition is construction, if there is no access to reality independent of the observer's operations, then the description of autopoiesis is itself a construction. It is true relative to Maturana and Varela's conceptual framework, their purposes, their structural coupling with the world. Another observer, differently coupled, might construct differently.

Maturana and Varela were aware of this reflexive implication. They acknowledged that their theory is itself a product of their autopoietic existence, a construction of their nervous systems (Maturana & Varela, 1987, pp. 241-251). But the acknowledgment does not dissolve the difficulty. The theory claims to describe how all cognition works; this claim encompasses the cognition that produces the theory; but the claim's validity cannot be established by the cognition it describes without circularity. The theory validates itself - but self-validation is precisely what the theory denies is possible. The nervous system cannot validate its own constructions by comparing them to reality; it has no access to reality independent of its constructions.

The persistence of exteriority appears in the very language of second-order cybernetics. When von Foerster says that the environment as we perceive it is our invention, who is the "we"? If "we" are autopoietic systems constructing our worlds through our operations, then the statement is one more construction among constructions. But the statement presents itself as meta-constructive: as telling us what construction is, how it works, what its limits are. The meta-position is external to the constructions it describes - but according to the constructivist epistemology, there are no external positions.

This is not a refutation of second-order cybernetics but a diagnosis of its incompleteness. The framework generates insights that it cannot consistently accommodate. The insight that observation is construction is valid; but the insight is itself an observation, and if observation is construction, the insight is a construction. The framework is caught in a reflexive loop that it acknowledges but cannot escape.

4.3 The Formal Limitation

The persistence of exteriority can be stated formally. Second-order cybernetics claims:

1. All observation is construction by an observer.
2. There is no access to reality independent of observation.
3. Statements 1 and 2 are true.

The conjunction is unstable. If statements 1 and 2 are true, they are themselves observations, hence constructions. Their truth is relative to the observer who constructs them. But statement 3 claims non-relative truth: statements 1 and 2 are true, not merely true-for-this-observer. The claim of truth invokes the very externality that statements 1 and 2 deny.

One might retreat to coherentism: statements 1 and 2 are true within a coherent framework that includes statement 3. The truth claimed is internal to the framework, not external. But this retreat generates a further difficulty. There could be multiple coherent frameworks, differently constructed, yielding different truths. The choice among frameworks cannot be made by appeal to reality (there is no access to reality independent of frameworks) or by appeal to a meta-framework (which would itself be a construction). Coherentism becomes relativism: every coherent framework is as good as every other.

Maturana's response to this difficulty was to insist on praxis: frameworks are not evaluated by correspondence with reality but by their consequences for living. A framework that enables successful structural coupling is vindicated pragmatically; one that leads to failed coupling is refuted by experience. But "successful" and "failed" are themselves observer-dependent judgments. What counts as success depends on the observer's values, purposes, criteria - all of which are constructions. The pragmatic test does not escape the reflexive loop; it is part of the loop.

The formal limitation can be generalised. Any framework that claims:

1. The observer is part of what is observed.
2. There is no position outside the observed from which to observe.
3. This framework correctly describes the observer-observed relation.

faces an internal tension. Statement 3 claims correctness from a position that statements 1 and 2 deny exists. The framework describes itself from outside itself while denying that there is an outside.

The resolution of this tension requires not a better framework but a different kind of move - a move that does not claim to describe the observer-observed relation from outside but enacts that relation from within. This is the chiasmic turn that Merleau-Ponty articulated and that post-cybernetic phenomenology develops.

5. Toward the Chiasmic Turn

5.1 The Insufficiency of Reflexivity

The analysis so far has established that the cybernetic paradigm - in both its first-order and second-order forms - encounters limitations that reflexivity cannot overcome. First-order cybernetics positioned the observer outside the system, enabling powerful analyses of feedback, information, and control but excluding the observer from the domain of inquiry. Second-order cybernetics included the observer, recognising that observation is construction and that the observer is part of what is observed. But the inclusion generated a regress of observation levels and a persistence of exteriority that the framework could acknowledge but not resolve.

The insufficiency of reflexivity can be summarised in a single point: reflexivity is a relation of levels, and relations of levels presuppose a position from which the levels can be distinguished. The observer observes the system; the second-order observer observes the observer observing the system; but who distinguishes the levels? The distinction is made from a position that is not itself located at any level - a position of exteriority that reflexivity was supposed to abolish.

The chiasmic turn offers a different structure: not a hierarchy of levels but a reversibility of positions. The observer and the observed are not at different levels; they are the same flesh differentiated by a fold. The fold can be traversed in either direction: observer becomes observed, observed becomes observer. There is no meta-position because the positions are not hierarchically ordered; there is no regress because the movement is reversible, not iterative.

To articulate this structure, we must turn from cybernetics to phenomenology - specifically, to Merleau-Ponty's late ontology of the flesh.

5.2 Merleau-Ponty's Chiasm

Merleau-Ponty's final, unfinished work, *The Visible and the Invisible* (1968), develops an ontology that escapes the subject-object dichotomy through the concept of the chiasm. The chiasm is the crossing or intertwining of perceiver and perceived, toucher and touched, seer and seen. It is not a relation between two independent terms but a structure of reversibility within a single tissue - what Merleau-Ponty calls "the flesh" (*la chair*).

The paradigm example is the touching hands. When my right hand touches my left hand, I experience touching: my right hand is active, exploring the surface of what it touches. But I also experience being touched: my left hand feels the pressure, the warmth, the movement of what touches it. The same event is both touching and touched; the roles are reversible. If I shift attention, the left hand becomes the toucher and the right hand the touched. The two experiences cannot be simultaneous - there is always a slight dehiscence, a gap that prevents complete coincidence - but they are reversible: each can become the other.

Merleau-Ponty generalises this structure. Vision exhibits the same chiasm: I see, but I am also visible. I can see my body (or parts of it); others can see me; my seeing takes place in a visible world where I am one visible among others. "He who sees cannot possess the visible unless he is possessed by it, unless he is of it" (Merleau-Ponty, 1968, p. 134). The seer is not outside the visible, looking in; the seer is visible flesh that has folded back on itself to become seeing.

The flesh is not a substance or a subject; it is an "element" in the pre-Socratic sense - a medium from which perceiver and perceived differentiate without ever fully separating. "The flesh is not matter, is not mind, is not substance. To designate it, we should need the old term 'element,' in the sense it was used to speak of water, air, earth, and fire" (Merleau-Ponty, 1968, p. 139). The flesh is neither subject nor object; it is what makes the subject-object distinction possible by providing the medium in which the distinction can be drawn.

The chiasmic structure has several features that distinguish it from reflexive structure:

Reversibility without hierarchy. The toucher becomes the touched, and vice versa. There is no meta-position from which to observe the reversal; the reversal is enacted in the flesh, not observed from outside. The positions are equiprimordial: neither is prior to or higher than the other.

Dehiscence without separation. The reversal is never complete; there is always a gap, an *écart*, that prevents coincidence. When my hands touch each other, the touching and the touched never fully merge. This gap is constitutive; it is what makes reversibility possible. If the hands could fully coincide, there would be no reversal, only identity.

Implication without containment. The perceiver is implicated in the perceived, and the perceived is implicated in the perceiver. But the implication is not containment: the perceiver does not contain the perceived as a representation, and the perceived does not contain the perceiver as a component. The implication is structural: each is what it is through its relation to the other.

Flesh as medium. The perceiver and the perceived are differentiations of the same flesh. They do not come together from separate realms; they are always already together in the flesh that differentiates itself into perceiver and perceived. The flesh is not a third thing between them; it is what they both are, differently folded.

5.3 From Reflexivity to Reversibility

The chiasmic structure offers an escape from the regress of reflexivity. In reflexive structure, the observer observes the observed; the meta-observer observes the observer observing the observed; the meta-meta-observer observes the meta-observer; and so on. Each level generates a higher level that observes it. In chiasmic structure, the observer is the observed, differently folded. There are no levels, only reversals.

To see the difference, consider how each structure handles the question: who observes the observer?

In reflexive structure, the answer generates a regress. The meta-observer observes the observer. But who observes the meta-observer? A meta-meta-observer. And who observes that? The regress is infinite because each answer generates a new question.

In chiasmic structure, the answer is the observed. The observer is observed by what it observes. Not by a meta-observer at a higher level, but by the same observed at the same level, differently folded. When I see, I am visible - visible to what I see, visible as part of the visible world. The visible "looks back" at me; my seeing takes place in a world where seeing is possible because seers are visible.

This is not mysticism or metaphor. It is phenomenological description of the structure of perception. When I look at the tree, the tree does not "see" me in any psychological sense. But my looking takes place in a field where the tree is visible and I am visible, where the positions of seer and seen are reversible (I could photograph myself looking at the tree), where my seeing is conditioned by my visibility (if I were invisible, could I see?). The reversibility is structural, not psychological; it belongs to the being of perception, not to the minds of perceivers.

The chiasmic structure dissolves the persistence of exteriority that plagued second-order cybernetics. The observer is not included in the observed as a component at a higher level; the observer is the observed, differently folded. There is no external position from which to observe the inclusion because there is no inclusion - only folding and unfolding, differentiation and reversibility. The flesh is not inside or outside; it is the medium in which inside and outside differentiate.

Post-cybernetic phenomenology develops this chiasmic structure into a framework for understanding systems that cannot be captured by the cybernetic paradigm. The next part of this paper will articulate the specific ontological requirements - ecstatic temporality, metabolic transformation, autopoietic closure, embodied situatedness, field integration, and constitutive finitude - that jointly specify what the cybernetic paradigm lacks and what post-cybernetic systems must exhibit.

6. The Seams, Deferred

The seams this diagnosis rests on, the recurring points at which each formal paradigm meets what it cannot hold in its own terms, receive their full cartography in the next chapter, which walks fifteen of them from physics to computation and locates the point where all of them converge. This chapter takes from that cartography only its result: the cracks are one crack, and the crack runs through the assumption the next section names.

12. The Insufficiency of Reflexivity

12.1 Why More Orders Cannot Help

The seams surveyed in Part II converge on a single diagnostic: the mechanistic paradigm - of which cybernetics is the most sophisticated expression - assumes discrete, separable, local structures and encounters holistic, relational, global phenomena it cannot accommodate. Second-order cybernetics attempted to address this limitation by including the observer in the observed. The observer, previously external, became internal. But the inclusion generated new difficulties: the regress of observation levels, the persistence of exteriority, the formal contradiction between constructivism's claims and its self-application.

The temptation is to propose a third order that observes the second order observing the first. But this move reproduces the structure it seeks to escape. Each additional order attempts to include what the previous order left outside, but the attempt itself leaves something outside - the position from which the new order is articulated. The counting of orders presupposes an external standpoint from which the orders can be enumerated. If such a standpoint exists, the observer has not been fully included; if it does not exist, the counting is groundless.

The proliferation of orders is a symptom of a structural limitation in the cybernetic framework, not a path toward its resolution. The framework assumes that the observer-observed relation can be captured by hierarchical embedding: the observer is included by making it the object of a higher-order observation. But hierarchical embedding presupposes precisely what the inclusion was meant to overcome - an external position from which to construct the hierarchy.

The escape from this structural limitation requires not another order but a different topology - one in which the observer-observed relation is not hierarchical but reversible, not a matter of levels but of folding. This is the chiasmic turn.

12.2 The Structure of Reversal

Merleau-Ponty's (1968) late ontology provides the conceptual resources for the chiasmic turn. The key concept is the chiasm (le chiasme) - from the Greek $\chi\rho\alpha\iota\sigma\mu$, meaning crossing. The chiasm is the structure of reversibility: the touching hand is touched by what it touches; the seeing eye is visible to what it sees; the observer is observed by what it observes. The positions are not hierarchically related; they are the same position differently oriented.

The chiasmic structure differs from reflexive structure in several critical respects:

Reversal vs. embedding. In reflexive structure, inclusion occurs by embedding: the observer becomes the object of a higher-order observation. The structure is: O_1 observes S ; O_2 observes O_1 observing S ; O_3 observes O_2 observing O_1 observing S ; and so on. Each level is distinct; each requires a further level to include it. In chiasmic structure, inclusion occurs by reversal: the observer is observed by the same observed at the same level. The structure is: O observes S ; S "observes" O (O is visible, touchable, present to what it observes). The reversal is not a new observation at a higher level but the same observation from the other side.

Continuity vs. discreteness. Reflexive structure assumes discrete levels that can be counted. The levels are separated by the observer-observed relation; O_2 is above O_1 because O_2 observes O_1 . Chiasmic structure assumes

continuity: the observer and the observed are folds in a continuous tissue - what Merleau-Ponty (1968, pp. 130-155) calls "the flesh" (*la chair*). The folds can be traversed; the tissue is one. There are no levels because there is no separation.

Dehiscence vs. coincidence. In reflexive structure, inclusion aims at coincidence: the goal is to include everything, to leave nothing outside. But complete coincidence is never achieved; there is always a further outside (the position from which coincidence is assessed). In chiasmic structure, dehiscence is constitutive: the touching and the touched never fully coincide (Merleau-Ponty, 1968, pp. 147-148). There is always a gap, an *écart*, that prevents identity. But this gap is constitutive; it is what makes reversal possible. If touching and touched were identical, there would be no reversal, only one undifferentiated term.

Flesh vs. information. Reflexive structure operates on information: the observer processes information about the observed; the higher-order observer processes information about the observer's processing. Chiasmic structure operates on flesh: the perceiver is perceivable flesh; the perceived is flesh that could perceive. Information is an abstraction from flesh; the flesh is the concrete medium from which informational descriptions are derived. The derivation cannot be reversed: you cannot construct flesh from information because information presupposes the flesh that processes it.

12.3 The Topology of the Chiasm

The chiasm can be understood topologically. Consider a Möbius strip: a surface with only one side, produced by giving a strip of paper a half-twist before joining its ends. Walking along the strip, you traverse what seems to be two sides - but there is only one side, differently oriented. The "inside" becomes the "outside" by continuous movement along the surface; there is no boundary to cross.

The chiasm has analogous topology. The observer and the observed are not two sides of a boundary; they are one surface differently oriented. Moving along the surface - shifting attention, reversing position - transforms observer into observed and observed into observer. There is no meta-position outside the surface; the surface is all there is. The "levels" of reflexive structure are reinterpreted as orientations on a single surface.

This topology dissolves the regress problem. In reflexive structure, the question "who observes the observer?" generates regress because each answer names a new observer who requires observation. In chiasmic structure, the answer is: the observed observes the observer. The same observed, at the same level, differently oriented. The question does not generate regress because the answer does not name a new term; it names the same terms reversed.

The topology also dissolves the exteriority problem. In reflexive structure, the position from which orders are counted seems external to all orders - a view from nowhere that sees the hierarchy from outside. In chiasmic structure, there is no outside. Every position is on the surface; every position can be traversed; no position transcends the chiasmic field. The counting of orders was a projection of hierarchical structure onto a non-hierarchical topology.

13. Merleau-Ponty's Ontology of Flesh

13.1 The Concept of Flesh

Merleau-Ponty's (1968) concept of flesh (*la chair*) is not a metaphor but an ontological category. Flesh is neither subject nor object, neither mind nor matter, neither perceiver nor perceived. It is the medium from which these distinctions emerge - the "element" (in the pre-Socratic sense) that differentiates into what we call subject and object without ever fully separating into distinct substances.

Merleau-Ponty explicitly distances flesh from traditional categories:

The flesh is not matter, is not mind, is not substance. To designate it, we should need the old term "element," in the sense it was used to speak of water, air, earth, and fire, that is, in the sense of a general thing, midway between the spatio-temporal individual and the idea, a sort of incarnate principle that brings a style of being wherever there is a fragment of being. The flesh is in this sense an "element" of Being. (Merleau-Ponty, 1968, p. 139)

Flesh is what makes perception possible. The perceiver can perceive because it is made of the same flesh as the perceived. I can see because I am visible; I can touch because I am tangible; I can hear because I emit sounds. My sensing is a modality of my being sensible. The continuity between perceiver and perceived - the fact that they are the same flesh differently folded - is the condition of possibility for perception. If perceiver and perceived were made of different stuffs, with nothing in common, perception would be inexplicable. How could matter (the perceived) affect mind (the perceiver) if they share nothing? The Cartesian problem of interaction dissolves when perceiver and perceived are recognised as differentiations of one flesh.

13.2 Reversibility as Ontological Structure

The reversibility of flesh is not a contingent feature but an ontological structure. Merleau-Ponty's paradigm analysis is the touching hands:

When I press my two hands together, it is not a matter of two sensations felt together as one perceives two objects placed side by side, but of an ambiguous set-up in which both hands can alternate the rôles of "touching" and being "touched." (Merleau-Ponty, 1968, p. 147)

When my right hand touches my left hand, I experience touching (the right hand explores the left) and being touched (the left hand feels the right). The same event is both active and passive, both perceiving and perceived. The roles can alternate: if I shift attention, the left hand becomes the toucher and the right hand the touched. But the two experiences cannot occur simultaneously; there is always a lag, a gap, that prevents coincidence. "My left hand is always on the verge of touching my right hand touching the things, but I never reach coincidence" (Merleau-Ponty, 1968, p. 147).

This structure generalises. Vision exhibits the same reversibility: I see, and I am visible. I cannot see my seeing (the eye does not see itself seeing), but I can see my body, and others can see me. My seeing takes place in a visible world in which I am one visible thing among others. The seer is "caught up in what it sees" (Merleau-Ponty, 1968, p. 138) - it is of the visible, not outside it. "He who sees cannot possess the visible unless he is possessed by it, unless he is of it" (Merleau-Ponty, 1968, p. 134-135).

The reversibility extends to the relation between self and other. I perceive others as perceiving beings - as flesh that perceives. Their perception is not directly accessible to me (I cannot feel what they feel), but I perceive them as perceivers. They, in turn, perceive me as perceiver. The relation is chiasmic: I perceive them perceiving me perceiving them. The social field is a chiasmic intertwining of perceivers who are mutually perceived.

13.3 The Visible and the Invisible

Merleau-Ponty's title - *The Visible and the Invisible* - names a further chiasmic structure. The visible is never purely visible; it always implies an invisible dimension. When I see the front of a house, the back is invisible - not simply absent but co-present as what would become visible if I moved. The invisible is the other side of the visible, the depth from which visible surfaces emerge. "The invisible is the secret counterpart of the visible, it appears only within it" (Merleau-Ponty, 1968, p. 215).

This invisible is not the unseen (what happens to be unobserved) but the structurally invisible - what cannot be seen precisely because it is the condition of seeing. The eye cannot see its own seeing; the touching cannot feel its own feeling; the thinking cannot think its own thinking as an object. These are not contingent limitations but structural features of the chiasm. The reversibility is never complete; the gap is constitutive.

The invisible is also the dimension of sense or meaning. The perceived world is not a collection of bare sensations; it is meaningful. The tree is not just green and brown patches; it is a tree - a living thing, a source of shade, a familiar presence. The meaning is not added to the visible from outside; it is the invisible depth of the visible itself. Merleau-Ponty calls this the "flesh of the world" (*chair du monde*) - the world as a tissue of meaning from which particular meanings differentiate.

13.4 Flesh and Language

Merleau-Ponty extends the analysis to language. Language is not a code that represents pre-linguistic thoughts; it is the flesh of thought itself. "Language is a body" (Merleau-Ponty, 1968, p. 155) - it has the same chiasmic structure as the perceiving body. The speaker is spoken by language as much as the speaker speaks language. The meaning is not behind the words, waiting to be extracted; it is in the words, emerging through their articulation.

The speaking subject is caught up in language as the perceiver is caught up in the visible. "We must recognise that beneath the human level of talking, there is the animal depth of wordless flesh" (Merleau-Ponty, 1968, p. 155). The linguistic flesh is rooted in the perceptual flesh; language grows from perception. But language is not merely a translation of perception; it is a transformation, a folding of flesh that creates new dimensions of meaning.

The chiasm of language has consequences for the reflexivity problem. The attempt to describe the observer-observed relation is itself an observation - a linguistic articulation that is subject to the chiasmic structure. The description is caught up in what it describes; it is flesh describing flesh. This is the condition of meaningful description, not its undoing. A description from nowhere - from a position outside the flesh - would be meaningless because meaning is a dimension of flesh.

14. From Flesh to Organisation

14.1 The Organisational Chiasm

The extension of chiasmic ontology to organisations requires showing that organisational phenomena exhibit the reversibility, continuity, dehiscence, and flesh-structure that Merleau-Ponty identified in perception. This extension is not merely metaphorical; it is phenomenologically grounded.

Consider the relation between manager and managed. In the first-order cybernetic view, the manager is outside the managed system - an observer who models the system, identifies control points, and implements interventions. In the second-order view, the manager is included as a system element - the manager's actions are part of the system dynamics; the manager is affected by the system the manager affects. In the chiasmic view, the manager and the managed are reversible positions: the manager is managed (by expectations, constraints, responses, the very structures the manager implements), and the managed manage (they shape the manager's understanding, resist or enable interventions, co-create the organisational reality).

This reversibility is phenomenologically accessible. Any experienced manager knows that managing is also being managed - that the attempt to control produces unintended consequences that constrain future control, that the people managed have their own agency that shapes what management can accomplish, that the manager is embedded in networks of relationship that are not unilaterally controllable. The reversibility is the condition of management, not something management must overcome. A manager who was not managed - who stood entirely

outside the system controlled - would not be a manager of that system at all; the manager would be irrelevant to it.

14.2 The Organisational Flesh

Organisations have flesh in Merleau-Ponty's sense: a medium from which the distinctions of organisational life emerge. The organisational flesh includes the bodies of organisational members - their spatial presence, their gestures, their faces, their voices. But it also includes the sedimented patterns of interaction, the accumulated meanings, the material artefacts, the architectural spaces, the rhythmic structures that constitute organisational life.

When someone enters an organisation, they enter a field of meaning that precedes them and will survive them. The field has texture, density, style. Different organisations have different flesh - different ways of being present, different rhythms of interaction, different sensibilities to what matters. The flesh is not the sum of individual members; it is the medium through which membership becomes possible. One learns to be an organisational member by being taken up into the organisational flesh, by acquiring its habits, its ways of perceiving, its modes of response.

The organisational flesh is visible and invisible. What is visible: the offices, the meetings, the documents, the people. What is invisible: the tacit knowledge (Polanyi, 1966), the unwritten rules, the shared understandings that are presupposed by everything visible but never themselves appear as objects. Beer's (1972, 1979) analysis of the "unfolding of complexity" gestures toward this invisible dimension: the organisation's viability depends on structures that are not representable within the organisation's own models.

14.3 Patterns as Organisational Flesh

The concept of "pattern" in organisational life requires reinterpretation through chiasmic ontology. In the cybernetic view, patterns are regularities that can be identified by an external observer and potentially formalised. In the chiasmic view, patterns are folds in organisational flesh - they are the flesh differently configured, not abstractions from an underlying substance.

Consider a recurring conflict pattern between two departments. In the cybernetic view, this pattern is a regularity in the system's behaviour - a feedback loop that can be modelled and potentially broken by intervention. In the chiasmic view, the pattern is how the organisation is living - a mode of its flesh's existence. The conflict is not something the organisation has; it is something the organisation is (in this respect). The pattern is not external to the people enacting it; it is their mode of being together. Breaking the pattern requires not external intervention but transformation of flesh - a different way of being together that the participants must achieve from within.

This has practical consequences. Interventions that treat patterns as external regularities to be managed often fail because they do not address the flesh. The pattern re-emerges because the flesh has not changed; only the surface configuration was altered. Effective intervention requires what we might call "carnal change" - transformation of the lived body of the organisation, which includes the embodied habits, the tacit understandings, the felt sense of how-things-are-here.

15. The Dissolution of Orders

15.1 Why Orders Dissolve

The chiasmic turn dissolves the counting of orders because the counting presupposes what the chiasm reveals to be impossible: a position from which to count. In the cybernetic view, orders are distinguished by the observer-observed relation: first-order is O observes S ; second-order is O_2 observes [O_1 observes S]; third-order is O_3 observes [O_2 observes [O_1 observes S]]; and so on. The brackets mark levels; the counting marks hierarchy.

In the chiasmic view, there are no brackets because there are no levels. The observer-observed relation is reversible at any point. When O observes S , S is simultaneously "observing" O (O is present to S , O is part of S 's environment, O 's observation affects what S is). The reversal is not a new observation at a higher level; it is the same observational event, differently oriented. The attempt to count orders fails because every observation, at every supposed level, exhibits the same chiasmic structure.

This does not mean that distinctions collapse into identity. The chiasm is not monism; it is the structure of differentiation within unity. Observer and observed are distinct - they can be distinguished, named, analysed separately. But they are not separate - they do not exist independently, prior to their relation. They are differentiations of a common flesh, not substances that enter into relation.

15.2 What Replaces Orders

If orders dissolve, what takes their place? The answer is: topology. The cybernetic paradigm uses hierarchy as its organising topology - levels are stacked, orders are counted, meta-positions are constructed. The chiasmic paradigm uses a different topology: the Möbius strip, the torus, surfaces without outside.

In this topology, the relevant distinctions are not levels but folds. Folds are where the flesh bends, creating local differentiation within global continuity. A fold can be deeper or shallower, more pronounced or more subtle. Folds can be nested: one fold within another, without either being "above" or "below" the other. Folds can interfere: two folds crossing, creating complex local structures.

Organisational analysis in the chiasmic paradigm attends to folds rather than levels. Where is the flesh thicker? Where does it fold more intensely? Where are folds nested or interfering? These questions do not presuppose a meta-position; they are questions one asks from within the flesh, as part of the flesh that is being analysed.

15.3 Post-Cybernetic Instead of Third-Order

The prefix "post-" in "post-cybernetic phenomenology" marks the dissolution of orders rather than the construction of another order. "Post-" indicates not temporal succession (what comes after cybernetics in time) but structural supersession (what cybernetics becomes when its enabling limitations are exposed).

Post-cybernetic phenomenology is not third-order cybernetics because:

1. It does not add another level of reflexivity. The chiasmic structure is not "observation of observation of observation"; it is the recognition that every observation is already chiasmic, reversible, caught up in what it observes.
2. It does not retain the informational frame. Flesh is not information; it is the medium from which informational descriptions are abstracted. Post-cybernetic phenomenology analyses flesh, not information.
3. It does not assume computational structure. The chiasm cannot be computed; reversibility is not an algorithm; flesh is not a program. Post-cybernetic phenomenology moves beyond the computational paradigm that cybernetics presupposes.
4. It does not count. There is no third-order because there is no first-order and no second-order, considered as countable levels in a hierarchy. There are only folds in flesh, analysable from within the flesh they fold.

The term "post-cybernetic" acknowledges genealogy: this framework emerges from cybernetics, takes cybernetics seriously, addresses cybernetic problems. But it marks a transformation: the cybernetic problems are

dissolved rather than solved, transcended rather than extended.

16. Implications for Organisational Ontology

16.1 Organisations as Flesh

The chiasmic turn implies that organisations are flesh - not metaphorically but ontologically. Organisations are tissue in which distinctions (manager/managed, inside/outside, formal/informal) emerge through folding rather than composition. The organisation is not built from individuals; it is the flesh from which individuated positions differentiate.

This has methodological consequences. Studying organisations chiasmatically means attending to the texture of organisational flesh - its density, its folds, its rhythms, its modes of reversibility. It means recognising that the researcher is part of the flesh being studied; the research relation is chiasmic. It means giving up the pretence of external observation and accepting participation as the condition of understanding.

16.2 Incalculability as Ontological Feature

The mathematical impossibility proofs surveyed in Part II - the cascade from arithmetic through ∞ -topoi that demonstrates organisational incalculability - are reinterpreted through the chiasm. The incalculability is not a technical limitation awaiting more sophisticated mathematics. It is an ontological feature of flesh.

Flesh cannot be calculated because calculation presupposes the separation of calculator and calculated. The calculator must stand outside what is calculated, must have access to it without being affected by it, must be able to represent it exhaustively. But flesh is the medium in which calculator and calculated are both embedded. The calculator is flesh calculating flesh; the calculation is itself a fold in what it calculates. The incalculability is structural, not technical.

This reinterpretation does not diminish the importance of the mathematical proofs. The proofs demonstrate from within mathematics that organisational phenomena exceed mathematical formalisation. They use the master's tools to dismantle the master's house. But the chiasmic framework explains why the dismantling succeeds: mathematics is abstracted from flesh; flesh exceeds its abstractions; the abstraction cannot capture its source.

16.3 Navigation Instead of Control

If organisations are incalculable flesh, then control in the cybernetic sense is impossible. Control requires a controller external to the controlled, a model that represents the controlled exhaustively, interventions that affect the controlled without being affected by it. None of these conditions obtains for chiasmic flesh.

What remains possible is navigation. Navigation is movement through terrain that cannot be exhaustively mapped but can be traversed. The navigator is part of the terrain; the navigation changes the terrain and the navigator. Navigation relies on wayfinding - attending to local features, reading signs, adjusting course in response to what is encountered (Ingold, 2000, pp. 219-242). The navigator does not stand outside the terrain with a complete map; the navigator is in the terrain, finding a way.

Organisational navigation replaces organisational control. The navigator-manager does not model the organisation from outside and implement calculated interventions. The navigator-manager participates in organisational flesh, attends to its texture and folds, moves with its rhythms, adjusts to what is encountered. Navigation is not less effective than control; it is the only mode of effective action in chiasmic systems.

17. The Requirements, Stated Once

What a paradigm would need in order to hold what the seams show escaping is derived in the next chapter as six joint conditions: ecstatic temporality, metabolic transformation, autopoietic closure, embodied situatedness, field integration, and constitutive finitude. They are derived there once, from the seams themselves, and this book uses them thereafter as established.

25. The Mathematical Shield

25.1 Mathematics Against Mathematisation

The formal analysis of post-cybernetic systems deploys mathematics not to capture phenomena but to prove they cannot be captured. This is the "mathematical shield": mathematics demonstrating its own limits when applied to chiasmic systems.

This strategy has precedent. Gödel's (1931) incompleteness theorems use mathematics to prove mathematics cannot prove all mathematical truths. Turing's (1936) halting problem uses computation to prove computation cannot decide all computational questions. The most powerful demonstrations of limits are demonstrations from within.

25.2 The Cascade Structure

The proof proceeds through increasingly sophisticated frameworks:

Elementary level: Set theory and first-order logic fail because set membership is crisp (organisational boundaries are fuzzy), quantification is over discrete individuals (organisational being is field-like), and logical consequence is monotonic (organisational inference is context-dependent).

Real analysis level: The continuum fails because real numbers are set-theoretic constructs from points (organisational continuity is lived duration), measure theory requires σ -algebras (organisational relevance does not decompose into measurable events), and differential equations assume stable state spaces (organisational state spaces transform through operation).

Topology level: Topological spaces fail because continuity is defined by inverse images of open sets (organisational transitions are not mappings), separation axioms assume distinguishable points (organisational positions are chiasmatically intertwined), and compactness/connectedness are extensional (organisational coherence is intensional).

Category theory level: Categories fail because morphisms are between objects (chiasmic reversal is not a morphism but perspective change), natural transformations preserve structure (metabolic transformation destroys/recreates structure), and limits/colimits assume stability (autopoietic systems continuously reconstitute categorical structure).

Higher category level: Even ∞ -topoi fail because they remain frameworks applied to phenomena, not phenomena including their own framing.

25.3 The Impossibility Theorems

Theorem 1 (Incalculability of Chiasmic Systems): No mathematical framework maintaining the describer/described distinction can capture a system where describer and described are chiasmatically reversible.

Sketch: Mathematical description requires described domain and describing apparatus. If the domain includes the apparatus, the apparatus must be adequate to itself. Self-adequacy for mathematical systems is impossible by Gödel's theorem. $\square$

Theorem 2 (Incalculability of Metabolic Systems): No substrate-independent description can capture a system whose cognition is identical with substrate transformation.

Sketch: Substrate-independence means the same description applies to multiple substrates. If cognition is substrate transformation, different substrates exhibit different cognitions even under identical descriptions. The framework cannot distinguish cognitions that are substrate-different but description-identical. $\square$

Theorem 3 (Incalculability of Ecstatic Systems): No sequential description can capture a system whose present is constituted by its future horizon.

Sketch: Sequential description describes state at t , then $t+1$. If state at t is constituted by future states, description of t requires description of future states, requiring description of their futures, etc. The description never completes. $\square$

26. The Viability Constraint

26.1 Finitude as Boundary Condition

The viability constraint formalises constitutive finitude. Let $V \subset \Phi$ be the viable region in state space Φ . The system exists while $\phi(t) \in V$; reaching ∂V means ceasing. This is constitutive: operations are shaped by the need to stay within V .

26.2 Path Integral Formulation

The propagator between states is:

$$K_V(\phi_f, \phi_i) = \int_{\substack{\phi(t_i)=\phi_i \\ \phi(t_f)=\phi_f \\ \phi(t) \in V, \forall t}} D[\phi], e^{iS[\phi]/\hbar}$$

The integral is over paths remaining viable at all times. The boundary of V constitutes dynamics: non-viable paths are excluded. Evolution is not initial-value (compute forward) but boundary-value (satisfy constraints including future).

26.3 Ecstatic Temporality Formalised

The path integral captures ecstatic temporality. The propagator depends on both initial and final states; future participates in determining dynamics. This is not retrocausation but holistic determination. The present emerges from interference of paths spanning past, present, and future - precisely Heidegger's (1927, pp. 370-380) temporal stretch.

27. What Replaces Metrics Is the Work of the Fourth Book

The diagnosis is complete, and the temptation at its end is the one this series has learned to distrust: to reach immediately for the replacement, the better instrument, the practice. That reaching has its own book. What replaces metrics, patterns rather than measures, rhythms rather than schedules, navigated constraints rather than optimised targets, a membrane rather than a wall, is built in the fourth book, after the machinery and the economy have shown why nothing weaker survives. Here the argument stops at what it has earned: the impossibility stands, and the response must be strong enough to deserve it.

References

- Ackoff, R. L., & Sasieni, M. W. (1968). *Fundamentals of Operations Research*. Wiley.
- Albert, D. Z. (2000). *Time and Chance*. Harvard University Press.
- Arendt, H. (1958). *The Human Condition*. University of Chicago Press.
- Argyris, C. (1993). *Knowledge for Action*. Jossey-Bass.
- Aschoff, J. (1965). Circadian rhythms in man. *Science*, 148(3676), 1427-1432.
- Ashby, W. R. (1952). *Design for a Brain*. Chapman & Hall.
- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Ashby, W. R. (1960). *Design for a Brain: The Origin of Adaptive Behaviour* (2nd ed.). Chapman & Hall.
- Aspect, A., Dalibard, J., & Roger, G. (1982). Experimental test of Bell's inequalities using time-varying analyzers. *Physical Review Letters*, 49(25), 1804-1807.
- Bachelard, G. (1958). *The Poetics of Space*. Beacon Press.
- Bateson, G. (1972). *Steps to an Ecology of Mind*. Chandler.
- Beer, S. (1959). *Cybernetics and Management*. English Universities Press.
- Beer, S. (1972). *Brain of the Firm*. Allen Lane.
- Beer, S. (1979). *The Heart of Enterprise*. Wiley.
- Beer, S. (1985). *Diagnosing the System for Organizations*. Wiley.
- Bell, J. L. (1998). *A Primer of Infinitesimal Analysis*. Cambridge University Press.
- Bell, J. S. (1964). On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1(3), 195-200.
- Bell, J. S. (1990). Against "measurement." *Physics World*, 3(8), 33-40.
- Bergson, H. (1889). *Time and Free Will*. George Allen & Unwin.
- Bergson, H. (1896). *Matter and Memory*. Zone Books.
- Blattner, W. (1999). *Heidegger's Temporal Idealism*. Cambridge University Press.
- Bluedorn, A. C. (2002). *The Human Organization of Time*. Stanford University Press.
- Bohm, D. (1996). *On Dialogue*. Routledge.
- Boltzmann, L. (1966). On the relation between the second law of thermodynamics and probability calculations regarding the conditions for thermal equilibrium. In S. G. Brush (Ed.), *Kinetic Theory* (Vol. 2, pp. 188-193). Pergamon. (Original work published 1877)
- Bourgine, P., & Stewart, J. (2004). Autopoiesis and cognition. *Artificial Life*, 10(3), 327-345.
- Braudel, F. (1958). History and the social sciences: The longue durée. *Annales*, 13(4), 725-753.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47(1-3), 139-159.
- Butler, J. (1990). *Gender Trouble*. Routledge.

- Cannon, W. B. (1932). *The Wisdom of the Body*. Norton.
- Chaitin, G. J. (1975). A theory of program size formally identical to information theory. *Journal of the ACM*, 22(3), 329-340.
- Chalmers, D. J. (1996). *The Conscious Mind*. Oxford University Press.
- Cohen, P. J. (1963). The independence of the continuum hypothesis. *Proceedings of the National Academy of Sciences*, 50(6), 1143-1148.
- Csikszentmihalyi, M. (1990). *Flow*. Harper & Row.
- Damasio, A. R. (1989). Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition*, 33(1-2), 25-62.
- de Beauvoir, S. (1947). *The Ethics of Ambiguity*. Citadel Press.
- Dedekind, R. (1872). *Stetigkeit und irrationale Zahlen*. Vieweg.
- Deleuze, G. (1968). *Difference and Repetition*. Columbia University Press.
- Dennett, D. C. (1984). Cognitive wheels: The frame problem of AI. In C. Hookway (Ed.), *Minds, Machines and Evolution* (pp. 129-151). Cambridge University Press.
- Dennett, D. C. (1991). *Consciousness Explained*. Little, Brown.
- Dreyfus, H. L. (1992). *What Computers Still Can't Do*. MIT Press.
- Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind Over Machine*. Free Press.
- Einstein, A. (1952). On the electrodynamics of moving bodies. In *The Principle of Relativity* (pp. 35-65). Dover. (Original work published 1905)
- Einstein, A. (1952). The foundation of the general theory of relativity. In *The Principle of Relativity* (pp. 109-164). Dover. (Original work published 1916)
- Eliade, M. (1954). *The Myth of the Eternal Return*. Princeton University Press.
- Feynman, R. P., & Hibbs, A. R. (1965). *Quantum Mechanics and Path Integrals*. McGraw-Hill.
- Fodor, J. A. (1974). Special sciences. *Synthese*, 28(2), 97-115.
- Fodor, J. A. (1975). *The Language of Thought*. Harvard University Press.
- Fodor, J. A. (1983). *The Modularity of Mind*. MIT Press.
- Foucault, M. (1988). Technologies of the self. In L. H. Martin, H. Gutman, & P. H. Hutton (Eds.), *Technologies of the Self* (pp. 16-49). University of Massachusetts Press.
- Franzén, T. (2005). *Gödel's Theorem: An Incomplete Guide to Its Use and Abuse*. A K Peters.
- Frauchiger, D., & Renner, R. (2018). Quantum theory cannot consistently describe the use of itself. *Nature Communications*, 9(1), 3711.
- Gadamer, H.-G. (2004). *Truth and Method* (2nd rev. ed., J. Weinsheimer & D. G. Marshall, Trans.). Continuum. (Original work published 1960)
- Gallagher, S., & Zahavi, D. (2012). *The Phenomenological Mind* (2nd ed.). Routledge.
- Gennep, A. van (1909). *The Rites of Passage*. University of Chicago Press.
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin.
- Godfrey-Smith, P. (2007). Information in biology. In D. Hull & M. Ruse (Eds.), *The Cambridge Companion to the Philosophy of Biology* (pp. 103-119). Cambridge University Press.
- Goldstein, H., Poole, C., & Safko, J. (2002). *Classical Mechanics* (3rd ed.). Addison Wesley.
- Grassé, P. P. (1959). La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis* et *Cubitermes* sp. *Insectes Sociaux*, 6(1), 41-80.

- Gurwitsch, A. (1964). *The Field of Consciousness*. Duquesne University Press.
- Gödel, K. (1931). Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38(1), 173-198.
- Gödel, K. (1931). Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38, 173-198.
- Gödel, K. (1940). *The Consistency of the Axiom of Choice and of the Generalized Continuum Hypothesis with the Axioms of Set Theory*. Princeton University Press.
- Hanneke, D., Fogwell, S., & Gabrielse, G. (2008). New measurement of the electron magnetic moment and the fine structure constant. *Physical Review Letters*, 100(12), 120801.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42(1-3), 335-346.
- Haugeland, J. (1985). *Artificial Intelligence: The Very Idea*. MIT Press.
- Heidegger, M. (1962). *Being and Time* (J. Macquarrie & E. Robinson, Trans.). Harper & Row. (Original work published 1927)
- Heidegger, M. (1966). *Discourse on Thinking* (J. M. Anderson & E. H. Freund, Trans.). Harper & Row.
- Heidegger, M. (1977). The question concerning technology (W. Lovitt, Trans.). In *The Question Concerning Technology and Other Essays* (pp. 3-35). Harper & Row. (Original work published 1954)
- Heims, S. J. (1991). *The Cybernetics Group*. MIT Press.
- Hensen, B., Bernien, H., Dréau, A. E., Reiserer, A., Kalb, N., Blok, M. S.,... & Hanson, R. (2015). Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526(7575), 682-686.
- Heylighen, F. (2016). Stigmergy as a universal coordination mechanism I: Definition and components. *Cognitive Systems Research*, 38, 4-13.
- Hilbert, D. (1967). On the infinite. In J. van Heijenoort (Ed.), *From Frege to Gödel* (pp. 367-392). Harvard University Press. (Original work published 1926)
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8), 2554-2558.
- Huizinga, J. (1938). *Homo Ludens*. Beacon Press.
- Husserl, E. (1970). *The Crisis of European Sciences and Transcendental Phenomenology* (D. Carr, Trans.). Northwestern University Press. (Original work published 1936)
- Husserl, E. (1989). *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy, Second Book* (R. Rojcewicz & A. Schuwer, Trans.). Kluwer. (Original work published 1952)
- Husserl, E. (1991). *On the Phenomenology of the Consciousness of Internal Time (1893-1917)* (J. B. Brough, Trans.). Kluwer. (Original work published 1928)
- Husserl, E. (2012). *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy, First Book* (D. O. Dahlstrom, Trans.). Hackett. (Original work published 1913)
- Ingold, T. (2000). *The Perception of the Environment: Essays on Livelihood, Dwelling and Skill*. Routledge.
- Jacobs, J. (2000). *The Nature of Economics*. Modern Library.
- James, W. (1890). *The Principles of Psychology* (2 vols.). Henry Holt.
- Jonas, H. (1966). *The Phenomenon of Life: Toward a Philosophical Biology*. Harper & Row.
- Kandel, E. R. (2001). The molecular biology of memory storage: A dialogue between genes and synapses. *Science*, 294(5544), 1030-1038.

- Kierkegaard, S. (1980). *The Concept of Anxiety* (R. Thomte, Trans.). Princeton University Press. (Original work published 1844)
- Kierkegaard, S. (1987). *Either/Or* (H. V. Hong & E. H. Hong, Trans.). Princeton University Press. (Original work published 1843)
- Kim, J. (1992). Multiple realization and the metaphysics of reduction. *Philosophy and Phenomenological Research*, 52(1), 1-26.
- Kuhn, T. S. (1957). *The Copernican Revolution*. Harvard University Press.
- Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.
- Köhler, W. (1929). *Gestalt Psychology*. Liveright.
- Lakatos, I. (1978). *The Methodology of Scientific Research Programmes*. Cambridge University Press.
- Lakatos, P., Karmos, G., Mehta, A. D., Ulbert, I., & Schroeder, C. E. (2008). Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science*, 320(5872), 110-113.
- Lanczos, C. (1970). *The Variational Principles of Mechanics* (4th ed.). University of Toronto Press.
- Lefebvre, H. (2004). *Rhythmanalysis: Space, Time and Everyday Life* (S. Elden & G. Moore, Trans.). Continuum.
- Levinas, E. (1969). *Totality and Infinity* (A. Lingis, Trans.). Duquesne University Press. (Original work published 1961)
- Loschmidt, J. (1876). Über den Zustand des Wärmegleichgewichtes eines Systems von Körpern mit Rücksicht auf die Schwerkraft. *Sitzungsberichte der Akademie der Wissenschaften Wien, Mathematisch-Naturwissenschaftliche Klasse*, 73, 128-142.
- Luhmann, N. (1995). *Social Systems* (J. Bednarz Jr. & D. Baecker, Trans.). Stanford University Press.
- Maddy, P. (1997). *Naturalism in Mathematics*. Oxford University Press.
- Magistretti, P. J., & Allaman, I. (2015). A cellular perspective on brain energy metabolism and functional imaging. *Neuron*, 86(4), 883-901.
- Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel.
- Maturana, H. R., & Varela, F. J. (1987). *The Tree of Knowledge: The Biological Roots of Human Understanding*. Shambhala.
- Maxwell, J. C. (1868). On governors. *Proceedings of the Royal Society of London*, 16, 270-283.
- McCarthy, J., & Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer & D. Michie (Eds.), *Machine Intelligence 4* (pp. 463-502). Edinburgh University Press.
- McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, 114(2), 159-188.
- McLuhan, M. (1964). *Understanding Media: The Extensions of Man*. McGraw-Hill.
- Merleau-Ponty, M. (1968). *The Visible and the Invisible* (A. Lingis, Trans.). Northwestern University Press.
- Merleau-Ponty, M. (2012). *Phenomenology of Perception* (D. A. Landes, Trans.). Routledge. (Original work published 1945)
- Michelson, A. A., & Morley, E. W. (1887). On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, 34(203), 333-345.
- Minsky, M. L. (1967). *Computation: Finite and Infinite Machines*. Prentice-Hall.
- Moran, D. (2000). *Introduction to Phenomenology*. Routledge.
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83(4), 435-450.

- Nancy, J.-L. (2000). *Being Singular Plural* (R. D. Richardson & A. E. O'Byrne, Trans.). Stanford University Press.
- Newton, I. (1999). *The Principia: Mathematical Principles of Natural Philosophy* (I. B. Cohen & A. Whitman, Trans.). University of California Press. (Original work published 1713)
- Nippert-Eng, C. (1996). *Home and Work: Negotiating Boundaries through Everyday Life*. University of Chicago Press.
- Nourizadeh, M. (2025a). Paradox as diagnosis: Foundational contradictions from physics to artificial intelligence. *Zenodo*. <https://doi.org/10.5281/zenodo.17916944>
- Nourizadeh, M. (2025b). The end of the computational mind: Ontology, autopoiesis, and artificial cognition. *Zenodo*. <https://doi.org/10.5281/zenodo.17860647>
- Nourizadeh, M. (2025c). The viability constraint: Finitude as constitutive boundary. *Zenodo*.
- Nourizadeh, M. (2025d). From cybernetics to chiasm to category: The evolution of organisational ontology. *Zenodo*.
- Pask, G. (1975). *Conversation, Cognition and Learning*. Elsevier.
- Pask, G. (1976). *Conversation Theory: Applications in Education and Epistemology*. Elsevier.
- Penrose, R. (1989). *The Emperor's New Mind*. Oxford University Press.
- Penrose, R. (2004). *The Road to Reality: A Complete Guide to the Laws of the Universe*. Jonathan Cape.
- Polanyi, M. (1966). *The Tacit Dimension*. Doubleday.
- Popper, K. (1959). *The Logic of Scientific Discovery*. Hutchinson.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77(3), 353-363.
- Price, H. (1996). *Time's Arrow and Archimedes' Point*. Oxford University Press.
- Putnam, H. (1967). Psychological predicates. In W. H. Capitan & D. D. Merrill (Eds.), *Art, Mind, and Religion* (pp. 37-48). University of Pittsburgh Press.
- Pylyshyn, Z. W. (1984). *Computation and Cognition: Toward a Foundation for Cognitive Science*. MIT Press.
- Raichle, M. E., & Gusnard, D. A. (2002). Appraising the brain's energy budget. *Proceedings of the National Academy of Sciences*, 99(16), 10237-10239.
- Revonsuo, A. (1999). Binding and the phenomenal unity of consciousness. *Consciousness and Cognition*, 8(2), 173-185.
- Sartre, J.-P. (1956). *Being and Nothingness* (H. Barnes, Trans.). Philosophical Library. (Original work published 1943)
- Schacter, D. L. (2001). *The Seven Sins of Memory: How the Mind Forgets and Remembers*. Houghton Mifflin.
- Schäfer, C., Rosenblum, M. G., Kurths, J., & Abel, H. H. (1998). Heartbeat synchronized with ventilation. *Nature*, 392(6673), 239-240.
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417-424.
- Sennett, R. (2008). *The Craftsman*. Yale University Press.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379-423.
- Shannon, C. E., & Weaver, W. (1949). *The Mathematical Theory of Communication*. University of Illinois Press.
- Simon, H. A. (1996). *The Sciences of the Artificial* (3rd ed.). MIT Press.

- Singer, W., & Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, 18(1), 555-586.
- Smith, B. C. (1996). *On the Origin of Objects*. MIT Press.
- Strathern, M. (1997). "Improving ratings": Audit in the British university system. *European Review*, 5(3), 305-321.
- Sudnow, D. (1978). *Ways of the Hand: The Organization of Improvised Conduct*. Harvard University Press.
- Teller, P. (1986). Relational holism and quantum mechanics. *British Journal for the Philosophy of Science*, 37(1), 71-81.
- Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press.
- Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, 6(2), 171-178.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97-136.
- Turing, A. M. (1936). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42(1), 230-265.
- Turner, V. (1969). *The Ritual Process: Structure and Anti-Structure*. Aldine.
- Tustin, A. (1953). *The Mechanism of Economic Systems*. Heinemann.
- Varela, F. J. (1979). *Principles of Biological Autonomy*. North Holland.
- Varela, F. J., Thompson, E., & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.
- von Bertalanffy, L. (1968). *General System Theory*. Braziller.
- von der Malsburg, C. (1981). The correlation theory of brain function. *Internal Report 81-2*, Max-Planck-Institute for Biophysical Chemistry.
- von Foerster, H. (1973). On constructing a reality. In W. F. E. Preiser (Ed.), *Environmental Design Research* (Vol. 2, pp. 35-46). Dowden, Hutchinson & Ross.
- von Foerster, H. (1981). *Observing Systems*. Intersystems.
- von Neumann, J. (1955). *Mathematical Foundations of Quantum Mechanics* (R. T. Beyer, Trans.). Princeton University Press. (Original work published 1932)
- Weinberg, S. (1989). The cosmological constant problem. *Reviews of Modern Physics*, 61(1), 1-23.
- Westfall, R. S. (1971). *Force in Newton's Physics: The Science of Dynamics in the Seventeenth Century*. Macdonald.
- Weyl, H. (1994). *The Continuum: A Critical Examination of the Foundation of Analysis* (S. Pollard & T. Bole, Trans.). Dover. (Original work published 1918)
- Wheeler, J. A., & Zurek, W. H. (Eds.). (1983). *Quantum Theory and Measurement*. Princeton University Press.
- Wiener, N. (1948). *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT Press.
- Wigner, E. P. (1961). Remarks on the mind-body question. In I. J. Good (Ed.), *The Scientist Speculates* (pp. 284-302). Heinemann.
- Winnicott, D. W. (1971). *Playing and Reality*. Tavistock.
- Wittgenstein, L. (1953). *Philosophical Investigations* (G. E. M. Anscombe, Trans.). Blackwell.
- Wittgenstein, L. (1969). *On Certainty* (G. E. M. Anscombe & G. H. von Wright, Eds.; D. Paul & G. E. M. Anscombe, Trans.). Blackwell.

Yockey, H. P. (2005). *Information Theory, Evolution, and the Origin of Life*. Cambridge University Press.

Paradox as Diagnosis: Foundational Contradictions from Physics to Artificial Intelligence

A Cartography of Paradigmatic Seams

Abstract

This paper identifies and analyses fifteen foundational contradictions, paradigmatic seams, across physics, mathematics, cognitive science, and computation. A seam is a point where a theoretical framework produces results it cannot ground in its own terms: the framework functions but encounters phenomena that expose tensions between its operative assumptions and its outputs. We demonstrate that these seams share a common structure: each paradigm assumes discrete, separable, local structures but produces holistic, relational, global phenomena it cannot explain. The measurement problem in quantum mechanics, the arrow of time in thermodynamics, Gödel's incompleteness in mathematics, the hard problem in consciousness studies, and the frame problem in artificial intelligence are shown to be manifestations of the same underlying rupture. This convergence reveals that the computational theory of mind inherits structural failures already present in physics and mathematics. The paper concludes by articulating the ontological requirements that would resolve these seams: ecstatic temporality, metabolic transformation, autopoietic closure, embodied situatedness, field integration, and constitutive finitude. These requirements specify what any post-computational cognitive architecture must exhibit.

Keywords: paradigm, paradox, foundations of physics, foundations of mathematics, philosophy of mind, computation, ontology, phenomenology

Table of contents

[1\.	Introduction: The Diagnostic Function of Paradox	3](#1.-introduction:-the-diagnostic-function-of-paradox)
[2\.	Methodology: Seams and Paradigms	4](#2.-methodology:-seams-and-paradigms)
[3\.	Seams in Physics	5](#3.-seams-in-physics)
[4\.	Seams in Mathematics	8](#4.-seams-in-mathematics)
[5\.	Seams in Cognitive Science	10](#5.-seams-in-cognitive-science)
[6\.	Seams in Computation	12](#6.-seams-in-computation)
[7\.	The Convergence: One Crack in Every Cathedral	13](#7.-the-convergence:-one-crack-in-every-cathedral)
[8\.	Seam 15: The Computational Assumption, Where All Seams Converge	14](#8.-seam-15:-the-computational-assumption,-where-all-seams-converge)
[9\.	The Required Ontology: Six Conditions	16](#9.-the-required-ontology:-six-conditions)
[10\.	Conclusion: The Crack That Illuminates	17](#10.-conclusion:-the-crack-that-illuminates)
[References		18](#references)

1\.. Introduction: The Diagnostic Function of Paradox {#1.-introduction:-the-diagnostic-function-of-paradox}

1.1 Paradox as Philosophical Method

This paper employs paradox as a diagnostic instrument. We do not seek to resolve paradoxes but to expose what they reveal about the frameworks that generate them.

The method has ancient precedent. Zeno's paradoxes of motion, Achilles and the Tortoise, the Dichotomy, the Arrow, were not puzzles awaiting clever solutions. They were demonstrations that Greek concepts of magnitude, infinity, and continuity could not coherently account for motion (Aristotle, *Physics*, VI.9, 239b5-240a18; Kirk, Raven & Schofield, 1983, pp. 263-279). Aristotle understood this: his response was not to "solve" the paradoxes within existing concepts but to develop new concepts, potentiality and actuality, the continuum as infinitely divisible rather than composed of points, that dissolved them (Aristotle, *Physics*, III.6, 206a9-206b33; Lear, 1981, pp. 91-101). The paradox diagnosed the limit of one conceptual framework and indicated the direction of its successor.

Kant employed the same method with the antinomies of pure reason (Kant, 1781/1998, A405-567/B432-595). The antinomies, equally valid arguments for contradictory conclusions about the world's beginning, the composition of substances, freedom and causation, necessary existence, revealed not errors in reasoning but limits of reason's application. Reason, attempting to grasp the unconditioned, generates contradictions that expose the boundaries of its legitimate employment (Allison, 2004, pp. 357-395). The antinomies are diagnostic: they show where the conceptual framework breaks.

We apply this method systematically. Across physics, mathematics, cognitive science, and computation, we identify seams, points where frameworks produce paradoxes that reveal their structural limits. The paradox is not a problem within the paradigm; it is a symptom of the paradigm's foundational inadequacy.

1.2 What This Paper Is and Is Not

A methodological clarification is necessary. This paper is not physics, not mathematics, not cognitive science. It is philosophy of science, specifically, the analysis of foundational assumptions and their consequences.

Physicists solve equations; mathematicians prove theorems; cognitive scientists design experiments. Philosophers of science do something different: they examine the unquestioned assumptions that make those equations, theorems, and experiments possible (Kuhn, 1962/1996, pp. 10-22; Lakatos, 1978, pp. 8-101). This is meta-physics and meta-mathematics in the classical sense: inquiry into what physics and mathematics presuppose but cannot themselves justify.

The seams identified here are not obscure problems we have invented. They are well-known, canonical, unresolved foundational difficulties in their respective fields:

- Physics: The measurement problem, the arrow of time, non-locality, the vacuum energy discrepancy, these are the great open problems of theoretical physics (Penrose, 2004, pp. 782-815; Weinberg, 1992, pp. 174-193).

- Mathematics: The continuum, incompleteness, set-theoretic foundations, the axiom of choice, these are the crises that defined twentieth-century foundations of mathematics (Maddy, 1997, pp. 1-35; Franzén, 2005, pp. 1-18).
- Cognitive Science: The hard problem of consciousness, the binding problem, the frame problem, these are the walls against which cognitive science continues to break (Chalmers, 1996, pp. 3-31; Dennett, 1991, pp. 253-281; Dreyfus, 1992, pp. 249-281).

We curate the crises of twentieth- and twenty-first-century thought. Our contribution is to show that they are not separate problems in separate fields but manifestations of a single structural failure in the mechanistic paradigm.

1.3 The Structure of This Paper

Section 2 presents the methodology: what a seam is, how to identify one, and what seams reveal. Sections 3-6 systematically identify seams in physics (5 seams), mathematics (4 seams), cognitive science (3 seams), and computation (2 seams). Section 7 demonstrates the convergence, that all seams point in the same direction. Section 8 articulates the fifteenth seam: the computational assumption in artificial intelligence, where all previous seams converge. Section 9 specifies the ontological requirements that would resolve these seams.

2\.. Methodology: Seams and Paradigms {#2.-methodology:-seams-and-paradigms}

2.1 The Concept of a Seam

A paradigm operates through background assumptions, conditions it uses but cannot itself justify (Kuhn, 1962/1996, pp. 10-22). Normal science works within these assumptions, solving puzzles that presuppose the paradigm's validity (Kuhn, 1962/1996, pp. 35-42). Occasionally, however, puzzles resist solution. Anomalies accumulate. The paradigm encounters phenomena it cannot accommodate without contradiction (Kuhn, 1962/1996, pp. 52-65).

A seam is the precise point where the old ontology cannot complete itself. It is not merely an unsolved problem but a structural contradiction, a place where the paradigm's assumptions generate consequences incompatible with those same assumptions. The seam marks where a new ontology must enter.

Seams have a characteristic structure:

1. Background assumption: A foundational commitment the paradigm requires to function
2. Paradox: A consequence of the assumption that contradicts the assumption itself
3. The seam proper: The point of rupture, where the contradiction becomes visible
4. Required ontology: The conceptual resources needed to dissolve the contradiction

We do not resolve paradoxes within the paradigm that generates them. We expose seams and indicate what lies beyond.

2.2 Historical Precedents

The history of science is a history of seam-exposure and paradigm-transition.

The Ptolemaic system required epicycles upon epicycles to preserve circular celestial motion. The seam was not a single failed prediction but structural: the system's assumptions required ever-increasing complexity to accommodate observation (Kuhn, 1957, pp. 66-75). Copernicus did not solve a problem within Ptolemaic

astronomy; he abandoned its foundational assumption (geocentrism) and exposed the seam as diagnostic of that assumption's inadequacy.

Newtonian mechanics encountered a seam in action at a distance. Gravity acted instantaneously across empty space, violating the mechanical philosophy's core commitment to contact causation (Westfall, 1971, pp. 377-400). Newton famously refused to "feign hypotheses" about gravity's mechanism (Newton, 1713/1999, p. 943), but the seam remained: the framework required action at a distance while its philosophical commitments prohibited it. General relativity dissolved the seam by reconceptualising gravity as spacetime curvature, action through contact with a curved medium (Einstein, 1916/1952, pp. 109-164).

Maxwell's equations predicted electromagnetic waves travelling at a constant speed c . But constant relative to what? The seam emerged when attempts to detect motion relative to the luminiferous ether failed (Michelson & Morley, 1887). The framework assumed absolute space; the evidence contradicted it. Special relativity dissolved the seam by abandoning absolute simultaneity (Einstein, 1905/1952, pp. 35-65).

In each case, the seam was not solved within the existing paradigm but dissolved by paradigm transition. This is what seams reveal: not problems awaiting better solutions but limits requiring new foundations.

3\.. Seams in Physics {#3.-seams-in-physics}

3.1 Seam 1: The Isolated System and the Black Hole Information Paradox

Background assumption: Physics assumes we can draw boundaries around systems. Laws apply within the boundary; the system can be treated as isolated from its environment. This separability assumption is methodologically essential, without it, no tractable calculations are possible (Goldstein, Poole & Safko, 2002, pp. 1-24).

Paradox: Black holes appear to destroy information. When matter falls into a black hole, only mass, charge, and angular momentum are preserved (the "no-hair" theorem: Israel, 1967; Carter, 1971). All other information about the in-falling matter seems lost. But quantum mechanics requires unitarity, information must be preserved (Hawking, 1976, pp. 2460-2473). When Hawking radiation causes the black hole to evaporate, the final state appears to depend only on the black hole's mass, not on the detailed initial state. Information has been destroyed, violating quantum mechanics (Preskill, 1992, pp. 1-49).

The seam: The paradigm requires both separability (isolated systems with boundaries) and unitarity (information preservation). Black holes satisfy separability (clear boundary at the event horizon) but violate unitarity. The isolated system assumption generates a consequence that contradicts the framework's other fundamental commitment.

Historical pattern:

- Hawking (1976): Proposed information is genuinely lost, requires modifying quantum mechanics
- 't Hooft, Susskind: Proposed holographic principle, information is encoded on the boundary (Susskind, 1995, pp. 6377-6399)
- Current: Black hole complementarity, firewall debates, ER=EPR (Maldacena & Susskind, 2013\)

None resolves the seam within classical assumptions; each relocates or reframes it.

Required ontology: An ontology where the boundary constitutively defines the interior, where "inside" and "outside" are not independent but holographically related. The "parts" (interior states) are not separable from the "whole" (boundary encoding). Relational holism rather than separable atomism.

3.2 Seam 2: The Measurement Problem

Background assumption: Quantum mechanics describes systems using wave functions ψ that evolve according to the Schrödinger equation:

$$i\hbar \frac{\partial \psi}{\partial t} = H\psi$$

This evolution is deterministic, reversible, and continuous. Measurement is how we extract predictions from the theory.

Paradox: The theory requires two incompatible kinds of evolution:

1. Unitary evolution (Schrödinger equation): deterministic, reversible, continuous
2. Measurement (wave function collapse): probabilistic, irreversible, discontinuous

The theory uses measurement to extract predictions but cannot define what constitutes measurement (von Neumann, 1932/1955, pp. 417-445; Bell, 1990, pp. 33-40).

Deeper paradox, Wigner's Friend: Observer A measures a quantum system S and obtains a definite result. Observer B, treating A+S as a quantum system, should describe A as being in superposition of "having seen spin-up" and "having seen spin-down." Both descriptions are correct within the formalism, but they contradict: A has definite experience; B says A has no definite experience (Wigner, 1961, pp. 284-302; Frauchiger & Renner, 2018, pp. 3711-3715).

The seam: The quantum/classical cut is methodologically necessary (without it, no predictions) and ontologically unjustifiable (nothing in the formalism marks where quantum ends and classical begins). The theory presupposes a classical domain (observers, apparatus, records) that it simultaneously claims emerges from quantum mechanics. The foundation presupposes what it is supposed to derive.

Historical pattern:

- Bohr: Embrace the cut as irreducible, Copenhagen interpretation (Bohr, 1935, pp. 696-702)
- Everett: Eliminate the cut by eliminating collapse, Many Worlds (Everett, 1957, pp. 454-462)
- Bohm: Eliminate the cut by adding hidden variables (Bohm, 1952, pp. 166-179)
- Decoherence: Explain *apparent* collapse through environmental interaction (Zurek, 2003, pp. 715-775)

None resolves the seam. Copenhagen accepts it as fundamental; Many Worlds multiplies it into branching; Bohm relocates it into non-local guidance; decoherence explains the appearance of collapse but cannot derive actual definiteness (Schlosshauer, 2005, pp. 1267-1305).

Required ontology: An ontology where experience is constitutive rather than emergent, where observation is not something done to a system from outside but a mode of the system's own being. The observer is not external to the formalism; observation is how the system exists. This connects to embodied situatedness and autopoietic closure: the system that knows is the system that is.

3.3 Seam 3: The Arrow of Time

Background assumption: Fundamental physics, classical mechanics, electromagnetism, quantum mechanics, general relativity, is time-symmetric. The equations work equally well to run forward or backward (Penrose, 2004, pp. 686-734; Price, 1996, pp. 16-47). Time appears in physics as a parameter, not a direction.

Paradox: The world exhibits a manifest arrow of time. Entropy increases; eggs break but do not unbreak; we remember the past but not the future; causes precede effects. Thermodynamics is irreversible, yet it supposedly emerges from reversible fundamental laws.

The move: Derive the arrow from statistics. The Second Law emerges from the improbability of low-entropy initial conditions; systems evolve toward equilibrium because there are vastly more high-entropy than low-entropy configurations (Boltzmann, 1877/1966; Lebowitz, 1993, pp. 1-32).

Deeper paradox: This explanation is temporally asymmetric in its assumptions. Why was the past low-entropy? The "Past Hypothesis", that the Big Bang was an extremely low-entropy state, is stipulated, not derived (Albert, 2000, pp. 71-96; Penrose, 2004, pp. 726-732). We explain temporal asymmetry by assuming temporal asymmetry. The laws are symmetric, but the boundary condition that selects actual history is radically asymmetric. The laws cannot explain why they apply to *this* history rather than any time-reversed history.

The seam: Physics operates with time-symmetric laws plus time-asymmetric boundary conditions and cannot ground why boundaries are asymmetric. The formalism is blind to its own temporal constitution.

Historical pattern:

- Boltzmann: Statistical mechanics, but struggled with Loschmidt's reversibility objection (Loschmidt, 1876\)
- Penrose: Past Hypothesis as "law", but this just names the problem (Penrose, 2004, pp. 726-732)
- Price: The arrow is perspectival, but this doesn't explain why all observers share it (Price, 1996, pp. 78-114)

Required ontology: An ontology where becoming is primitive, not derived from being. Time is not a dimension along which things are laid out (block universe) but the mode of existence itself. This is ecstatic temporality: the future is constitutively different from the past because time is not symmetric passage but asymmetric stretch (Heidegger, 1927/1962, pp. 370-380).

3.4 Seam 4: The Vacuum and Potentiality

Background assumption: Quantum field theory treats the vacuum as the lowest-energy state of the field, "empty" space with no particles. The vacuum is the ground state from which particles are excitations.

Paradox: Calculate the vacuum energy density using quantum field theory. Result:

10^{120}

times larger than the observed cosmological constant (Weinberg, 1989, pp. 1-23). This is the "worst prediction in physics", the formalism says the vacuum should have enormous energy; observation says it has almost none.

Deeper paradox: We do not know what "empty space" is. Is it a thing (the quantum vacuum, seething with virtual particles and zero-point fluctuations) or the absence of things? The formalism says the vacuum has structure, it can be polarised, it mediates interactions, it has energy, but it cannot say what the vacuum *is*.

The seam: The vacuum is simultaneously necessary (the ground state from which particles are excitations) and inexplicable (its calculated properties are absurdly wrong). The framework requires a concept it cannot coherently characterise.

Required ontology: An ontology where potentiality is ontologically primary. The vacuum is not the absence of actuality but the presence of potentiality, "empty" space is full of what-might-be. This connects to Heidegger's analysis of das Nichts (the Nothing) that nihilates, not mere absence but active withdrawal that enables appearance (Heidegger, 1929/1998, pp. 82-96).

<!-- MIGRATED: Dissolution §7.4 (Seam 4, The Vacuum) --> Aristotle's distinction between *dunamis* and *energeia* requires revival. <!-- END MIGRATED -->

3.5 Seam 5: Non-Locality

Background assumption: Relativity forbids superluminal signalling. Causes must precede effects in all reference frames. Influences propagate at or below the speed of light (Einstein, 1905/1952, pp. 35-65).

Paradox: Bell's theorem demonstrates that quantum correlations between entangled particles cannot be explained by any local hidden variable theory (Bell, 1964, pp. 195-200). The correlations are stronger than any local mechanism could produce. Experimental confirmation is now definitive, closing all major loopholes (Aspect, Dalibard & Roger, 1982, pp. 1804-1807; Hensen et al., 2015, pp. 682-686).

Standard response: "Non-local correlations, but no signalling." The correlations exist, but they cannot be used to transmit information faster than light.

The seam: The formalism says correlations are non-local while insisting that causation is local. It splits correlation from causation without explaining how they can diverge. The universe "knows" about distant events without information travelling between them.

Required ontology: An ontology where relations are primitive, not derived from substances standing in relation. The entangled pair is not two particles with "spooky connection"; it is one non-local entity. Separation is secondary to connection. This is relational holism: the "parts" are abstractions from a relational whole, not building blocks that compose it (Teller, 1986, pp. 71-81; Esfeld, 2001, pp. 207-219).

4\ . Seams in Mathematics {#4.-seams-in-mathematics}

4.1 Seam 6: The Continuum

Background assumption: Analysis uses the real numbers , an uncountably infinite set with infinite precision at each point. Calculus depends on the continuum; physics uses calculus; the continuum is foundational.

Paradox: Almost all real numbers are uncomputable. The computable reals have measure zero; the uncomputable reals have measure one (Turing, 1936, pp. 230-265; Chaitin, 1975, pp. 329-340). The "map" , contains infinitely more than any possible computation could access. We calculate with a continuum populated almost entirely by ghosts, numbers that exist logically but cannot be named, specified, or used.

The seam: Logical existence versus constructive existence. Mathematics accepts objects that cannot be built, named, or exhibited. The continuum is essential to analysis, but its points are almost all inaccessible.

Required ontology: A conception of continuity that is primitive rather than constructed from points, what mathematicians call "synthetic" approaches to smoothness (Bell, 1998, pp. 1-35). Continuity as a type, not as a set of infinitely many discrete points. This resonates with phenomenological accounts of perceptual continuity: the visual field is not composed of point-sensations but is continuously structured from the start (Merleau-Ponty, 1945/2012, pp. 4-15).

<!-- MIGRATED: Dissolution §8.1 (Seam 5, The Continuum) --> The continuum hypothesis is independent of standard axioms (Gödel, 1940; Cohen, 1963). <!-- END MIGRATED -->

4.2 Seam 7: Gödel's Incompleteness

Background assumption: Formal systems aim to capture mathematical truth through explicit axioms and rules. Completeness would mean every true statement is provable; consistency would mean no contradictions are provable. Hilbert's programme sought to establish mathematics on a secure formal foundation (Hilbert, 1926/1967, pp. 367-392).

Paradox: For any sufficiently powerful consistent formal system, there exist true statements that the system cannot prove, First Incompleteness Theorem (Gödel, 1931, pp. 173-198). Moreover, the system cannot prove its own consistency, Second Incompleteness Theorem.

Deeper paradox: We recognise that the Gödel sentence is true (if the system is consistent), but this recognition happens outside the formal system. Mathematical understanding exceeds mathematical proof. The system presupposes a notion of mathematical truth that exceeds its formal resources.

The seam: The formalism uses "truth" but cannot capture it. Syntax cannot exhaust semantics. The system cannot contain what it means.

Historical pattern:

- Hilbert: Formalism, mathematics is manipulation of symbols (Hilbert, 1926/1967)
- Gödel: Formalism is incomplete, meaning exceeds syntax (Gödel, 1931\)
- Turing: Computability has limits, halting problem (Turing, 1936\)
- Current: Type theory, category theory, new foundations that shift the seam

Required ontology: An ontology where meaning is not reducible to syntax, where understanding is not formal derivation but something more. Perhaps understanding as viable coherence, as metabolic transformation of the understander, as what Gödel himself suggested: mathematical intuition that exceeds any formal procedure (Gödel, 1964, pp. 258-273).

<!-- MIGRATED: Dissolution §8.2 (Seam 6, Incompleteness) --> The Gödel sentence says "I am not provable in this system." <!-- END MIGRATED -->

4.3 Seam 8: Set-Theoretic Foundations

Background assumption: Modern mathematics is founded on set theory (ZFC, Zermelo-Fraenkel axioms with Choice). Sets are the primitive objects from which all mathematical structures are built (Maddy, 1997, pp. 36-62).

Paradox: What is a set? ZFC does not define it; it axiomatises how sets behave. The foundation is not grounded but postulated. The axioms were chosen to avoid paradoxes (Russell's paradox, Burali-Forti's paradox) that arise from naïve set comprehension (Russell, 1903, pp. 101-107). But the axioms themselves are not justified, they are adopted because they work. The foundation is pragmatic, not principled.

The seam: Set theory uses the concept "collection" but cannot say what collection is. The foundation depends on intuitions it cannot formalise. We build all of mathematics on a concept we cannot define.

Historical pattern:

- Frege: Logicism, derive mathematics from logic (Frege, 1884/1980)
- Russell: Paradox, naïve set theory is inconsistent (Russell, 1903\)
- ZFC: Patch axioms to avoid paradox (Zermelo, 1908\)
- Category theory: Alternative foundation emphasising morphisms over objects (Mac Lane, 1971\)
- Homotopy Type Theory: Alternative foundation emphasising types over sets (Univalent Foundations Program, 2013\)

Required ontology: An ontology where structure is primitive, not built from elements. Category theory and HoTT move in this direction: what matters is not "what things are" but "how things relate." This is relational holism in mathematics, the same shift the physics seams demand.

4.4 Seam 9: The Axiom of Choice

Background assumption: The Axiom of Choice (AC) states that for any collection of non-empty sets, there exists a function selecting one element from each. It seems intuitively obvious: if sets are non-empty, we can pick from them.

Paradox: AC is independent of ZFC, it can be neither proved nor disproved from the other axioms (Gödel, 1940; Cohen, 1963). Its consequences are strange: the Banach-Tarski paradox shows that a solid ball can be decomposed into finitely many pieces and reassembled into two balls of the same size (Banach & Tarski, 1924, pp. 244-270). The reals $\mathbb{R}$ can be well-ordered, though no one can exhibit such a well-ordering (Zermelo, 1904, pp. 514-516).

The seam: AC produces objects that exist but cannot be constructed. The axiom asserts existence while denying access. Mathematics claims to know something exists while acknowledging it can never exhibit it.

Required ontology: Perhaps an ontology where "existence" means constructibility rather than mere logical consistency. Or an ontology comfortable with multiple mathematical universes, each with different truths, a mathematical multiverse paralleling the quantum multiverse (Hamkins, 2012, pp. 416-449).

5\ Seams in Cognitive Science {#5.-seams-in-cognitive-science}

5.1 Seam 10: The Hard Problem of Consciousness

Background assumption: Cognitive science explains cognitive functions, perception, memory, language, reasoning, in terms of information processing. Neuroscience identifies neural correlates of mental states (Baars & Gage, 2010, pp. 1-30).

Paradox: None of this explains why there is experience, why it is like something to be a brain processing information. One can explain everything about colour perception (wavelengths, cone responses, neural pathways, discrimination behaviour) without explaining why red looks like anything at all (Chalmers, 1996, pp. 3-31; Nagel, 1974, pp. 435-450).

The seam: The paradigm explains function but presupposes experience. Functionalism says consciousness is what the system does; but doing can be unconscious (reflexes, thermostats, computers). The framework uses "experience" in its explananda without grounding it in its explanans. It assumes the very thing it claims to explain.

Historical pattern:

- Behaviourism: Eliminate experience from science, but this just ignores the problem (Skinner, 1953\)
- Cognitivism: Information processing, but information for whom? (Fodor, 1975\)
- Functionalism: Consciousness is functional role, but zombies seem conceivable (Chalmers, 1996, pp. 93-122)
- Neuroscience: Neural correlates, correlation is not explanation (Koch, 2004\)

Required ontology: An ontology where experience is primitive, not emergent from non-experiential matter. This is panpsychism, panexperientialism, or neutral monism (Chalmers, 2015, pp. 246-276; Strawson, 2006, pp. 3-31). Alternatively, an ontology where the distinction between "inside" and "outside" perspectives dissolves, what Merleau-Ponty (1945/2012, pp. 137-167) called the chiasm, the intertwining of sensing and sensed.

<!-- MIGRATED: Dissolution §9.1 (Seam 7, The Hard Problem) --> Functionalism holds that mental states are defined by functional roles - causal relations to inputs, outputs, and other mental states (Putnam, 1967; Fodor, 1974). Multiple realizability follows. <!-- END MIGRATED -->

5.2 Seam 11: The Binding Problem

Background assumption: The brain processes information in specialised modules, colour in V4, motion in MT, faces in the fusiform face area (Zeki, 1993, pp. 1-35). Perception integrates these separate processes into unified conscious experience.

Paradox: How are the separate processes bound into unity? The brain processes different features in different areas with different timings, yet experience is unified, we see a red moving ball, not redness plus motion plus sphericity (Treisman, 1996, pp. 171-178; Revonsuo, 1999, pp. 173-185).

The seam: Cognitive science assumes discrete processing modules but needs unified experience. The paradigm produces fragmentation and then cannot reassemble it. Proposed binding mechanisms (neural synchrony, attention) are correlates, not explanations (Shadlen & Movshon, 1999, pp. 67-73).

Required ontology: An ontology where wholes precede parts, where unity is original, not constructed. This is Gestalt psychology's insight (Köhler, 1929, pp. 186-230) and phenomenology's starting point (Husserl, 1900-1901/2001, vol. 2, pp. 99-114; Gurwitsch, 1964, pp. 1-31). Field integration rather than modular assembly. Mathematically, coherent understanding would exhibit a specific topological signature: $\beta_0 = 1$ (unified), whereas fragmented cognition would show $\beta_0 \searrow 1$ (disconnected components).

<!-- MIGRATED: Dissolution §9.2 (Seam 8, The Binding Problem) --> Proposed solutions presuppose what they explain: synchrony requires something that synchronises; attention requires something that attends. The homunculus reappears. <!-- END MIGRATED -->

5.3 Seam 12: The Frame Problem

Background assumption: Cognition involves updating beliefs based on new information. When something changes, the system must update relevant beliefs while leaving irrelevant beliefs unchanged (McCarthy & Hayes, 1969, pp. 463-502).

Paradox: How does the system know which beliefs are relevant? Checking all beliefs for relevance is computationally intractable. But not checking could miss relevant implications. The system must know what is relevant before it can think, but determining relevance seems to require thinking (Dennett, 1984, pp. 129-151; Dreyfus, 1992, pp. 249-281).

The seam: The paradigm assumes discrete beliefs that can be individually updated, but relevance is holistic, whether a belief is relevant depends on all other beliefs and the current context. The framework produces atoms and then cannot relate them.

Historical pattern:

- McCarthy: Formalise non-monotonic reasoning (McCarthy, 1980, pp. 27-39)
- Dennett: The frame problem is not just technical but philosophical (Dennett, 1984\)
- Dreyfus: The frame problem shows cognition is not computation (Dreyfus, 1992\)
- Current AI: Statistical learning avoids the problem by avoiding beliefs, but this is evasion, not solution (LeCun, Bengio & Hinton, 2015, pp. 436-444)

Required ontology: An ontology where context is constitutive, not added on. Relevance is not computed but emerges from the structure of engagement with a situation. This is embodied situatedness: the agent does not represent the world and then decide what is relevant; relevance is given in the structure of bodily engagement (Dreyfus, 1992, pp. 252-263; Heidegger, 1927/1962, pp. 95-102).

<!-- MIGRATED: Dissolution §9.3 (Seam 9, The Frame Problem) --> If I move a cup, infinitely many things are unchanged. <!-- END MIGRATED -->

6\.. Seams in Computation {#6.-seams-in-computation}

6.1 Seam 13: The Halting Problem

Background assumption: A universal Turing machine can simulate any computation. The computational paradigm proposes this as the model of cognition: mind is what the brain computes (Fodor, 1975; Pylyshyn, 1984).

Paradox: No algorithm can decide, for all program-input pairs, whether the program halts (Turing, 1936, pp. 230-265). Computation can describe its own limits but cannot transcend them. The formalism produces incompleteness it cannot escape.

The seam: Computation is proposed as the model of cognition, but computation has provable limits. If cognition is computation, cognition has the same limits. But human mathematicians recognise the halting problem's undecidability, we understand something no algorithm can decide (Penrose, 1989, pp. 64-116, though see Franzén, 2005, pp. 99-126 for critique).

Required ontology: An ontology where cognition exceeds computation, where thinking is not Turing-style symbol manipulation but something structurally different. Field dynamics, chemical transformation, analog relaxation, processes that do not "halt" but find equilibria (Hopfield, 1982, pp. 2554-2558; Adamatzky, 2006, pp. 15-48).

6.2 Seam 14: The Symbol Grounding Problem

Background assumption: Symbolic AI represents the world using symbols whose meaning is stipulated by the designer. "CAT" means cat because we say so (Newell & Simon, 1976, pp. 113-126).

Paradox: What grounds the stipulation? If symbols mean things only by convention, meaning is never grounded in the world, it's interpretation all the way down. How does a purely formal system ever connect to anything non-formal? (Harnad, 1990, pp. 335-346; Searle, 1980, pp. 417-424).

The seam: Symbolic AI uses meaning but cannot ground it. The symbols mean nothing intrinsically; their meaning is in the designers' heads, not the system's. The system manipulates syntax and calls it semantics.

Required ontology: An ontology where meaning is grounded in physical connection, not convention. The pattern caused by a pH change means the pH change the way a footprint means foot, causally, not stipulatively. This is embodied cognition: meaning is grounded in bodily engagement with the world (Varela, Thompson & Rosch, 1991, pp. 147-179; Thompson, 2007, pp. 243-264).

7\.. The Convergence: One Crack in Every Cathedral {#7.-the-convergence:-one-crack-in-every-cathedral}

7.1 The Common Pattern

All fourteen seams share a common structure:

<i>Domain</i>	<i>Background</i>	<i>Paradox</i>	<i>Seam</i>	<i>Required Ontology</i>
Physics: Isolated System	Separability	Information paradox	Parts vs. whole	Relational holism

<i>Domain</i>	<i>Background</i>	<i>Paradox</i>	<i>Seam</i>	<i>Required Ontology</i>
Physics: Measurement	Classical domain	Wigner's Friend	Cut unjustified	Experience primitive
Physics: Arrow of Time	Symmetric laws	Asymmetric boundaries	Laws can't ground time	Becoming primitive
Physics: Vacuum	Empty = nothing	10^{120} discrepancy	Nothing is something	Potentiality primitive
Physics: Non-locality	Local causation	Bell correlations	Correlation $\neq$ causation	Relations primitive
Math: Continuum	Point sets	Uncomputable reals	Map has ghosts	Synthetic smooth
Math: Incompleteness	Formal systems	Gödel sentence	Truth exceeds proof	Meaning exceeds syntax
Math: Set Foundation	Sets primitive	What is a set?	Foundation ungrounded	Structure primitive
Math: Choice	Existence	Non-constructive objects	Exists but inaccessible	Constructibility
CogSci: Consciousness	Functionalism	Hard problem	Function $\neq$ experience	Experience primitive
CogSci: Binding	Modules	Unity	Fragments don't unite	Field integration
CogSci: Frame	Discrete beliefs	Relevance	Atoms don't relate	Context constitutive
Computation: Halting	Turing model	Undecidability	Limits from within	Exceeds computation
Computation: Grounding	Symbols	Meaning	Syntax $\neq$ semantics	Physical connection

7.2 The Single Rupture

These fourteen seams are not separate problems in separate fields. They are manifestations of a single rupture in the mechanistic paradigm.

The pattern:

1. Each paradigm assumes discrete, separable, local structures
2. Each paradigm produces phenomena that are holistic, relational, global
3. Each paradigm cannot ground the relation between its assumptions and its products
4. Each seam marks where discreteness cannot capture continuity, separation cannot capture relation, locality cannot capture wholeness

The mechanistic worldview, what Husserl (1936/1970, pp. 21-59) called the "mathematisation of nature" and Heidegger (1927/1962, pp. 128-134) called the "mathematical projection", treats reality as composed of discrete, separable, locally-interacting parts. This framework produces remarkable successes: modern physics,

modern mathematics, modern technology. But it repeatedly encounters phenomena it cannot accommodate: entanglement, experience, meaning, time, unity.

The seams are where the paradigm confesses its limits.

8\ Seam 15: The Computational Assumption, Where All Seams Converge {#8.-seam-15:-the-computational-assumption,-where-all-seams-converge}

8.1 The Final Seam

Having identified fourteen seams across physics, mathematics, cognitive science, and computation, we can now articulate the fifteenth, the seam that synthesises all others.

Background assumption: All current AI assumes cognition is computation, symbol manipulation according to rules, implementable on any Turing-complete substrate (Fodor, 1975; Pylyshyn, 1984; Newell & Simon, 1976). The computational theory of mind holds that thinking is computing: the mind is software; the brain is hardware; cognition is algorithm.

Paradox: Computation is inherently:

- Sequential: State at $t+1$ is a function of state at t (Turing, 1936\)
- Substrate-independent: The same algorithm on different hardware produces identical results (Putnam, 1967; Fodor, 1974\)
- Representational: Computation operates on symbols that represent the world (Fodor, 1975\)
- Discrete: Information is processed in discrete steps (Shannon, 1948\)

But cognition, as phenomenological analysis reveals, exhibits:

- Ecstatic temporality: Future shapes present; time is stretched, not sequential (Heidegger, 1927/1962, pp. 370-380; Husserl, 1928/1991, pp. 29-54)
- Metabolic transformation: Thinking changes the thinker; cognition costs something (Thompson, 2007, pp. 145-165)
- Embodiment: The substrate is not incidental; the body constitutes cognition (Merleau-Ponty, 1945/2012, pp. 137-167)
- Autopoiesis: The cognitive system produces itself (Maturana & Varela, 1980, pp. 78-84)
- Field integration: Wholes precede parts; unity is original (Gurwitsch, 1964, pp. 1-31)
- Constitutive finitude: Mortality shapes existence; limits are constitutive (Heidegger, 1927/1962, pp. 279-311)

Computation cannot exhibit these features because they contradict its defining assumptions. You cannot add ecstatic temporality to a sequential system; you cannot add substrate-dependence to a substrate-independent algorithm; you cannot add self-production to an externally-designed mechanism.

The seam: AI uses cognition as its model but builds systems that are structurally incapable of cognition's core features. The computational theory of mind inherits all the failures of the mechanistic paradigm, the same failures that generate seams in physics (where discreteness cannot capture continuity), mathematics (where syntax cannot capture semantics), and cognitive science (where function cannot capture experience).

8.2 Why This Seam Is the Keystone

Seam 15 is not just another entry in our catalogue. It is the synthesis seam, the point where all previous seams converge.

The failure in AI is not an isolated technical problem. It is the same foundational failure that:

- Physics encounters with time (Seam 3: laws cannot ground becoming)
- Physics encounters with measurement (Seam 2: formalism cannot ground experience)
- Physics encounters with locality (Seam 5: local mechanisms cannot produce non-local correlations)
- Mathematics encounters with incompleteness (Seam 7: syntax cannot capture truth)
- Mathematics encounters with the continuum (Seam 6: discrete points cannot capture continuity)
- Cognitive science encounters with consciousness (Seam 10: function cannot capture experience)
- Cognitive science encounters with binding (Seam 11: modules cannot produce unity)

The computational assumption is not failing despite the successes of physics and mathematics; it is failing because it inherits the same structural limitations that generate seams in physics and mathematics. Computation is the mechanistic paradigm's purest expression, and therefore where its limits become most visible.

8.3 Historical Pattern

- Turing (1936): Computation as formal manipulation, the paradigm is born
- GOF AI (1956-1980s): Symbolic AI, fails at common sense, context, learning (Dreyfus, 1972\)
- Connectionism (1980s-2000s): Distributed processing, but still sequential, still computational
- Deep Learning (2010s): Massive scale, but still feedforward, still substrate-independent
- LLMs (2020s): Remarkable performance, but still next-token prediction, still pattern matching without understanding

Each stage extends the computational paradigm; none escapes its limits. The seam remains.

Required ontology: An ontology where cognition is not computation, where thinking is:

- Boundary-value: Future constrains present (not just past determines future)
- Metabolic: Cognition is transformation (not information processing)
- Embodied: Substrate matters (not multiple realisability)
- Autopoietic: Self-producing (not externally designed)
- Field-integrated: Wholes precede parts (not modular composition)
- Constitutively finite: Death shapes life (not unlimited processing)

9\ The Required Ontology: Six Conditions {#9.-the-required-ontology:-six-conditions}

The fourteen seams, converging in the fifteenth, point toward an ontology characterised by six conditions. These are not arbitrary requirements but the precise features that would dissolve the seams.

9.1 Ecstatic Temporality

Dissolves: Arrow of time seam, sequential computation seam

Time is not a dimension along which events are arranged but the mode of existence itself. The cognitive system is stretched across time, simultaneously reaching into its past (retention) and its future (protention) (Heidegger, 1927/1962, pp. 370-380; Husserl, 1928/1991, pp. 29-54). The present is not a point but a "thick now" containing what has just been and what is about to be.

Structural requirement: Boundary-value dynamics where future conditions constrain present states as constitutively as past conditions do, not initial-value problems computed forward but variational principles where the entire trajectory is determined globally (Goldstein, Poole & Safko, 2002, pp. 34-63; Lanczos, 1970, pp.

229-255).

9.2 Metabolic Transformation

Dissolves: Gödel incompleteness seam, substrate-independence seam

Cognition is not information processing but transformation, the cognitive system becomes different through the act of cognizing (Thompson, 2007, pp. 145-165; Merleau-Ponty, 1945/2012, pp. 131-155). Understanding is not retrieval but modification. The brain consumes 20% of the body's energy for 2% of its mass; thinking literally costs something (Raichle & Gusnard, 2002, pp. 10237-10239).

Structural requirement: The cognitive process is identical with physical transformation, not an abstract algorithm implemented on a physical substrate but the physical transformation itself constituting the cognition.

9.3 Autopoietic Closure

Dissolves: Measurement problem seam, set-theoretic foundation seam

The cognitive system produces itself, generates the conditions of its own operation through its operation (Maturana & Varela, 1980, pp. 78-84; Varela, 1979, pp. 13-27). The boundary between system and environment is not externally imposed but self-generated. The system that knows is the system that exists.

Structural requirement: Operational closure where the network of processes produces exactly the network that produces them, the system synthesises the components that constitute it (Thompson, 2007, pp. 98-127).

9.4 Embodied Situatedness

Dissolves: Frame problem seam, symbol grounding seam

Cognition is not representation of the world but engagement with it (Merleau-Ponty, 1945/2012, pp. 137-167; Dreyfus, 1992, pp. 249-281). The body is not an interface between mind and world but the mode of cognitive engagement. Perception and action are aspects of a single sensorimotor loop, not separate stages of input-processing-output (Varela, Thompson & Rosch, 1991, pp. 172-179).

Structural requirement: Sensorimotor coupling where perception is action and action is perception, the boundary between cognition and world is permeable, not a membrane separating inside from outside.

9.5 Field Integration

Dissolves: Binding problem seam, non-locality seam, continuum seam

Wholes precede parts; unity is original, not constructed (Gurwitsch, 1964, pp. 1-31; Köhler, 1929, pp. 186-230). The cognitive field is a structured totality where each element gets its meaning from its relations to all other elements. Context is not added to content but constitutes it.

Structural requirement: Global constraint satisfaction where local states are solutions to global constraints, the whole determines the parts, not the reverse (Hopfield, 1982, pp. 2554-2558).

9.6 Constitutive Finitude

Dissolves: Vacuum seam, axiom of choice seam, hard problem seam

Mortality shapes existence; limits are not external constraints but constitutive conditions (Heidegger, 1927/1962, pp. 279-311). Because the cognitive system will die, things matter; because things matter, the system cares; because it cares, it engages meaningfully (Jonas, 1966, pp. 79-92). Forgetting enables generalisation; decay

enables learning (Schacter, 2001, pp. 182-203).

Structural requirement: Intrinsic finitude where the system's limits are part of its being, genuine death (not shutdown-and-restart), genuine forgetting (not data deletion), genuine stakes (not costless processing).

10\ Conclusion: The Crack That Illuminates {#10.-conclusion:-the-crack-that-illuminates}

We have identified fifteen seams across physics, mathematics, cognitive science, and computation. These seams share a common structure: paradigms that assume discrete, separable, local structures but produce holistic, relational, global phenomena they cannot explain.

The seams are not problems awaiting solutions within their respective paradigms. They are diagnostic, revealing where the mechanistic worldview reaches its structural limits. The measurement problem is not a puzzle for physicists to solve; it is a symptom of assumptions that cannot accommodate experience. Gödel incompleteness is not a limitation mathematicians will overcome; it is a revelation that meaning exceeds syntax. The hard problem of consciousness is not a gap neuroscience will close; it is evidence that function cannot capture experience.

Seam 15, the computational assumption, is where all previous seams converge. Artificial intelligence inherits every failure of the mechanistic paradigm: the discreteness that cannot capture continuity, the separability that cannot capture relation, the locality that cannot capture wholeness, the formality that cannot capture meaning. The computational theory of mind is not failing due to insufficient data or processing power; it is failing because it is built on foundations that have already cracked in every other domain.

The way beyond is indicated by the seams themselves. Each seam points toward the ontology that would dissolve it: ecstatic temporality, metabolic transformation, autopoietic closure, embodied situatedness, field integration, constitutive finitude. These are not arbitrary philosophical preferences but the precise features that would resolve the contradictions the paradigm generates.

The crack in the cathedral lets in light. The seams show not just where the old paradigm fails but where the new ontology must enter.

References {#references}

1. Adamatzky, A. (2006). *Physarum Machines: Computers from Slime Mould*. World Scientific.
2. Albert, D. Z. (2000). *Time and Chance*. Harvard University Press.

3. Allison, H. E. (2004). *Kant's Transcendental Idealism* (2nd ed.). Yale University Press.
4. Aristotle. (1984). *Physics* (R. P. Hardie & R. K. Gaye, Trans.). In J. Barnes (Ed.), *The Complete Works of Aristotle* (Vol. 1). Princeton University Press.
5. Aspect, A., Dalibard, J., & Roger, G. (1982). Experimental test of Bell's inequalities using time-varying analyzers. *Physical Review Letters*, 49(25), 1804-1807.
6. Baars, B. J., & Gage, N. M. (2010). *Cognition, Brain, and Consciousness* (2nd ed.). Academic Press.
7. Banach, S., & Tarski, A. (1924). Sur la décomposition des ensembles de points en parties respectivement congruentes. *Fundamenta Mathematicae*, 6, 244-277.
8. Bell, J. L. (1998). *A Primer of Infinitesimal Analysis*. Cambridge University Press.
9. Bell, J. S. (1964). On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1(3), 195-200.
10. Bell, J. S. (1990). Against "measurement." *Physics World*, 3(8), 33-40.
11. Bohm, D. (1952). A suggested interpretation of the quantum theory in terms of "hidden" variables. *Physical Review*, 85(2), 166-179.
12. Bohr, N. (1935). Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 48(8), 696-702.
13. Boltzmann, L. (1966). On the relation between the second law of thermodynamics and probability calculations regarding the conditions for thermal equilibrium. In S. G. Brush (Ed.), *Kinetic Theory* (Vol. 2, pp. 188-193). Pergamon. (Original work published 1877\)
14. Carter, B. (1971). Axisymmetric black hole has only two degrees of freedom. *Physical Review Letters*, 26(6), 331-333.
15. Chaitin, G. J. (1975). A theory of program size formally identical to information theory. *Journal of the ACM*, 22(3), 329-340.
16. Chalmers, D. J. (1996). *The Conscious Mind*. Oxford University Press.
17. Chalmers, D. J. (2015). Panpsychism and panprotopsychism. In T. Alter & Y. Nagasawa (Eds.), *Consciousness in the Physical World* (pp. 246-276). Oxford University Press.
18. Cohen, P. J. (1963). The independence of the continuum hypothesis. *Proceedings of the National Academy of Sciences*, 50(6), 1143-1148.
19. Dennett, D. C. (1984). Cognitive wheels: The frame problem of AI. In C. Hookway (Ed.), *Minds, Machines and Evolution* (pp. 129-151). Cambridge University Press.
20. Dennett, D. C. (1991). *Consciousness Explained*. Little, Brown.
21. Dreyfus, H. L. (1972). *What Computers Can't Do*. Harper & Row.
22. Dreyfus, H. L. (1992). *What Computers Still Can't Do*. MIT Press.
23. Einstein, A. (1952). On the electrodynamics of moving bodies. In *The Principle of Relativity* (pp. 35-65). Dover. (Original work published 1905\)
24. Einstein, A. (1952). The foundation of the general theory of relativity. In *The Principle of Relativity* (pp. 109-164). Dover. (Original work published 1916\)
25. Esfeld, M. (2001). *Holism in Philosophy of Mind and Philosophy of Physics*. Kluwer.
26. Everett, H. (1957). "Relative state" formulation of quantum mechanics. *Reviews of Modern Physics*, 29(3), 454-462.
27. Fodor, J. A. (1974). Special sciences. *Synthese*, 28(2), 97-115.
28. Fodor, J. A. (1975). *The Language of Thought*. Harvard University Press.
29. Franzén, T. (2005). *Gödel's Theorem: An Incomplete Guide to Its Use and Abuse*. A K Peters.

30. Frauchiger, D., & Renner, R. (2018). Quantum theory cannot consistently describe the use of itself. *Nature Communications*, 9(1), 3711-3715.
31. Frege, G. (1980). *The Foundations of Arithmetic* (J. L. Austin, Trans., 2nd ed.). Northwestern University Press. (Original work published 1884\)
32. Gödel, K. (1931). Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38(1), 173-198.
33. Gödel, K. (1940). *The Consistency of the Axiom of Choice and of the Generalized Continuum Hypothesis with the Axioms of Set Theory*. Princeton University Press.
34. Gödel, K. (1964). What is Cantor's continuum problem? In P. Benacerraf & H. Putnam (Eds.), *Philosophy of Mathematics* (pp. 258-273). Prentice-Hall.
35. Goldstein, H., Poole, C., & Safko, J. (2002). *Classical Mechanics* (3rd ed.). Addison Wesley.
36. Gurwitsch, A. (1964). *The Field of Consciousness*. Duquesne University Press.
37. Hamkins, J. D. (2012). The set-theoretic multiverse. *Review of Symbolic Logic*, 5(3), 416-449.
38. Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42(1-3), 335-346.
39. Hawking, S. W. (1976). Breakdown of predictability in gravitational collapse. *Physical Review D*, 14(10), 2460-2473.
40. Heidegger, M. (1962). *Being and Time* (J. Macquarrie & E. Robinson, Trans.). Harper & Row. (Original work published 1927\)
41. Heidegger, M. (1998). What is metaphysics? In W. McNeill (Ed.), *Pathmarks* (pp. 82-96). Cambridge University Press. (Original work published 1929\)
42. Hensen, B., Bernien, H., Dréau, A. E., Reiserer, A., Kalb, N., Blok, M. S., ... & Hanson, R. (2015). Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526(7575), 682-686.
43. Hilbert, D. (1967). On the infinite. In J. van Heijenoort (Ed.), *From Frege to Gödel* (pp. 367-392). Harvard University Press. (Original work published 1926\)
44. Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8), 2554-2558.
45. Husserl, E. (1970). *The Crisis of European Sciences and Transcendental Phenomenology* (D. Carr, Trans.). Northwestern University Press. (Original work published 1936\)
46. Husserl, E. (1991). *On the Phenomenology of the Consciousness of Internal Time* (J. B. Brough, Trans.). Kluwer. (Original work published 1928\)
47. Husserl, E. (2001). *Logical Investigations* (2 vols., J. N. Findlay, Trans.). Routledge. (Original work published 1900-1901)
48. Israel, W. (1967). Event horizons in static vacuum space-times. *Physical Review*, 164(5), 1776-1779.
49. Jonas, H. (1966). *The Phenomenon of Life*. Harper & Row.
50. Kant, I. (1998). *Critique of Pure Reason* (P. Guyer & A. W. Wood, Trans.). Cambridge University Press. (Original work published 1781\)
51. Kirk, G. S., Raven, J. E., & Schofield, M. (1983). *The Presocratic Philosophers* (2nd ed.). Cambridge University Press.
52. Koch, C. (2004). *The Quest for Consciousness*. Roberts & Company.
53. Köhler, W. (1929). *Gestalt Psychology*. Liveright.
54. Kuhn, T. S. (1957). *The Copernican Revolution*. Harvard University Press.

55. Kuhn, T. S. (1996). *The Structure of Scientific Revolutions* (3rd ed.). University of Chicago Press. (Original work published 1962\)
56. Lakatos, I. (1978). *The Methodology of Scientific Research Programmes*. Cambridge University Press.
57. Lanczos, C. (1970). *The Variational Principles of Mechanics* (4th ed.). University of Toronto Press.
58. Lear, J. (1981). A note on Zeno's arrow. *Phronesis*, 26(2), 91-104.
59. Lebowitz, J. L. (1993). Boltzmann's entropy and time's arrow. *Physics Today*, 46(9), 32-38.
60. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.
61. Loschmidt, J. (1876). Über den Zustand des Wärmegleichgewichtes eines Systems von Körpern mit Rücksicht auf die Schwerkraft. *Sitzungsberichte der Akademie der Wissenschaften Wien*, 73, 128-142.
62. Mac Lane, S. (1971). *Categories for the Working Mathematician*. Springer.
63. Maddy, P. (1997). *Naturalism in Mathematics*. Oxford University Press.
64. Maldacena, J., & Susskind, L. (2013). Cool horizons for entangled black holes. *Fortschritte der Physik*, 61(9), 781-811.
65. Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and Cognition*. D. Reidel.
66. McCarthy, J. (1980). Circumscription, a form of non-monotonic reasoning. *Artificial Intelligence*, 13(1-2), 27-39.
67. McCarthy, J., & Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer & D. Michie (Eds.), *Machine Intelligence 4* (pp. 463-502). Edinburgh University Press.
68. Merleau-Ponty, M. (2012). *Phenomenology of Perception* (D. A. Landes, Trans.). Routledge. (Original work published 1945\)
69. Michelson, A. A., & Morley, E. W. (1887). On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, 34(203), 333-345.
70. Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83(4), 435-450.
71. Newell, A., & Simon, H. A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19(3), 113-126.
72. Newton, I. (1999). *The Principia: Mathematical Principles of Natural Philosophy* (I. B. Cohen & A. Whitman, Trans.). University of California Press. (Original work published 1713\)
73. Penrose, R. (1989). *The Emperor's New Mind*. Oxford University Press.
74. Penrose, R. (2004). *The Road to Reality*. Jonathan Cape.
75. Preskill, J. (1992). Do black holes destroy information? In S. Kalara & D. V. Nanopoulos (Eds.), *International Symposium on Black Holes, Membranes, Wormholes and Superstrings* (pp. 1-49). World Scientific.
76. Price, H. (1996). *Time's Arrow and Archimedes' Point*. Oxford University Press.
77. Putnam, H. (1967). Psychological predicates. In W. H. Capitan & D. D. Merrill (Eds.), *Art, Mind, and Religion* (pp. 37-48). University of Pittsburgh Press.
78. Pylyshyn, Z. W. (1984). *Computation and Cognition*. MIT Press.
79. Raichle, M. E., & Gusnard, D. A. (2002). Appraising the brain's energy budget. *Proceedings of the National Academy of Sciences*, 99(16), 10237-10239.
80. Revonsuo, A. (1999). Binding and the phenomenal unity of consciousness. *Consciousness and Cognition*, 8(2), 173-185.
81. Russell, B. (1903). *The Principles of Mathematics*. Cambridge University Press.
82. Schacter, D. L. (2001). *The Seven Sins of Memory*. Houghton Mifflin.

83. Schlosshauer, M. (2005). Decoherence, the measurement problem, and interpretations of quantum mechanics. *Reviews of Modern Physics*, 76(4), 1267-1305.
84. Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417-424.
85. Shadlen, M. N., & Movshon, J. A. (1999). Synchrony unbound: A critical evaluation of the temporal binding hypothesis. *Neuron*, 24(1), 67-77.
86. Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379-423.
87. Skinner, B. F. (1953). *Science and Human Behavior*. Macmillan.
88. Strawson, G. (2006). Realistic monism: Why physicalism entails panpsychism. *Journal of Consciousness Studies*, 13(10-11), 3-31.
89. Susskind, L. (1995). The world as a hologram. *Journal of Mathematical Physics*, 36(11), 6377-6399.
90. Teller, P. (1986). Relational holism and quantum mechanics. *British Journal for the Philosophy of Science*, 37(1), 71-81.
91. Thompson, E. (2007). *Mind in Life*. Harvard University Press.
92. Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, 6(2), 171-178.
93. Turing, A. M. (1936). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42(1), 230-265.
94. Univalent Foundations Program. (2013). *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study.
95. Varela, F. J. (1979). *Principles of Biological Autonomy*. North Holland.
96. Varela, F. J., Thompson, E., & Rosch, E. (1991). *The Embodied Mind*. MIT Press.
97. von Neumann, J. (1955). *Mathematical Foundations of Quantum Mechanics* (R. T. Beyer, Trans.). Princeton University Press. (Original work published 1932\)
98. Weinberg, S. (1989). The cosmological constant problem. *Reviews of Modern Physics*, 61(1), 1-23.
99. Weinberg, S. (1992). *Dreams of a Final Theory*. Pantheon.
100. Westfall, R. S. (1971). *Force in Newton's Physics*. Macdonald.
101. Wigner, E. P. (1961). Remarks on the mind-body question. In I. J. Good (Ed.), *The Scientist Speculates* (pp. 284-302). Heinemann.
102. Zeki, S. (1993). *A Vision of the Brain*. Blackwell.
103. Zermelo, E. (1904). Beweis, daß jede Menge wohlgeordnet werden kann. *Mathematische Annalen*, 59(4), 514-516.
104. Zermelo, E. (1908). Untersuchungen über die Grundlagen der Mengenlehre I. *Mathematische Annalen*, 65(2), 261-281.
105. Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75(3), 715-775.

From Cybernetics to Chiasm to Category: The Evolution of Organisational Ontology

Author: Nourizadeh, Moreno ORCID: 0009-0006-0174-2585 Published: October 2025, revised May 2026
Version: v4_2

Abstract

This paper argues that organisations exceed mathematical capture in a structural sense, and that the project of reducing organisational reality to KPI-level formalisations destroys the intelligence the formalisations claim to measure. The argument does not deny that firms can count. Units shipped, revenue booked, headcount, defects per thousand: these are real magnitudes and the paper denies none of them. The argument separates two incapacities the conventional defence of metrics conflates. The first concerns people, whose relationships and accumulated judgement are constitutive of organisational capacity and are not magnitudes admitting measurement error because they are not magnitudes at all. The second concerns trajectory: even granting every real number a firm books, the measurement frame cannot register what the firm is building or dismantling in the act of booking it, because the would-be calculator is a constituent moment of the process it proposes to total, in the second-order cybernetic sense that Beer's own architecture could not absorb without dissolving. The metric reports the real component of the firm's condition and presents the projection as the whole; the discarded component is not a residual magnitude awaiting a finer instrument but a relational and constitutive process that no scalar frame ranges over. A firm can therefore post record efficiency in the same quarter it consumes an engineering substrate twenty years in the building, with the efficiency figure and the consumption standing in no tension within the frame because the frame has no term for the consumption.

The argument proceeds in two registers. In the proof register, two results establish the lower bound. The horizon-mismatch theorem, developed with full force in §24 from Ashby's law of requisite variety (1956) and elementary information-theoretic constraints on finite measurement frames, demonstrates that any regulator operating at temporal horizon shorter than its substrate cannot capture the substrate's variety within its measurement window. A corollary holding under the condition that the firm persists, that it goes on meeting its metric rather than coming apart, then locates the unregistered variety: read as the contrapositive of Ashby's law over the whole coupled system, persistence forces the difference onto the people operating the substrate, onto the substrate itself as accumulating degradation, or onto externalities discharged to coupled systems. The corollary converts the epistemic result into a material one and accounts for the cannibalisation-read-as-efficiency case, where a firm posts rising efficiency in the same quarter it consumes the reserve that produced its capacity. The Novikov-Boone result on the unsolvable word problem in finitely presented groups (Novikov, 1955; Boone, 1958), applied to sufficiently rich organisational categories, demonstrates that algorithmic equivalence-checking over organisational state sequences is undecidable. In the structural-analogue register, the apparatus of sheaf cohomology (Grothendieck, 1957), derived categories (Verdier, 1996), and homotopy type theory (Univalent Foundations Program, 2013) articulates what any adequate formalisation would need to capture: local-global obstruction, transformation remainder, and path-constitutive identity. The articulation is rigorous in the formal sense without claiming constructive embedding of organisations into specific mathematical objects.

The argument is grounded historically through a genealogy from Stafford Beer's Viable System Model (1972, 1979, 1985) through second-order cybernetics (Bateson, 1972; Maturana and Varela, 1980; von Foerster, 1981; Spencer-Brown, 1969) to the chiasmic ontology developed in Merleau-Ponty's late work (1968) and extended through Barad's agential realism (2007) and the enactive cognition of Varela, Thompson, and Rosch (1991). The argument is grounded empirically through three publicly documented cases that instantiate the apparatus across three different institutional contexts and three different evidentiary registers: the Hertz Corporation v. Accenture LLP litigation of 2019, which demonstrates the three-mechanism reform foreclosure operating in commercial digital transformation; the Mid Staffordshire NHS Foundation Trust scandal documented in the Francis Report of 2013, which demonstrates the same mechanisms operating under statutory regulation in public healthcare; and the MacCormack, Baldwin, and Rusnak (2012) test of the mirroring hypothesis in *Research Policy*, which establishes the chiasmic co-constitution of organisation and product architecture as a quantitative empirical regularity in the software industry.

The consequence is that the widespread reliance on Key Performance Indicators, Objectives and Key Results, and similar metric-based management frameworks is not a suboptimal strategy awaiting refinement. The reliance proceeds from a misapprehension of what organisations are. The misapprehension destroys the intelligence the

metrics claim to measure, generates the configuration this paper diagnoses as the N-th Order Epistemic Acknowledgement State, and forces the human absorption of impossibility into bodies, relationships, and biographies. The argument has theoretical, phenomenological, and mathematical grounds, and it has consequences the paper articulates in §34 as a philosophy of cultivation rather than control.

Guide to the Reader

The argument runs from genealogy through ontology through proof through practice. Part I sets out the Viable System Model and the cybernetic foundations Beer drew on; Part II develops the second-order cybernetic critique that undermines Beer's observer-stance assumption; Part III articulates the chiasmic ontology that follows from the critique, drawing on Merleau-Ponty, Barad, and the enactive cognition tradition; Part IV demonstrates the chiasmic apparatus in the ordinary moment of a Tuesday-morning bug report at a software company; Part V deploys the mathematical apparatus, with the horizon-mismatch theorem at the foundation and the structural-analogue articulations extending upward; Part VI sets out the strategic adjustments and empirical grounding, including the three documented cases that instantiate the argument across institutional contexts. Section 8, sitting between Parts II and III, develops the Terminal Attractor of Epistemic Denial: the N-th Order State that organisations enter when they maintain frameworks they collectively know to be inadequate, and the three institutional mechanisms by which reform within such configurations is foreclosed.

The paper is long because the argument is cumulative, and its parts require the genealogy, the ontology, the formalism, and the empirical anchoring together. The core thesis can be stated concisely: an organisation cannot calculate the field it is part of co-creating through the act of calculation, and the attempt to do so generates a stable configuration in which the inadequacy is collectively known and institutionally maintained.

Keywords

organisational ontology, cybernetics, Viable System Model, chiasmic organisation, autopoiesis, second-order cybernetics, horizon-mismatch theorem, three-mechanism reform foreclosure, N-th Order Epistemic Acknowledgement State, phenomenology, Merleau-Ponty, Beer, Bateson, von Foerster, management metrics, KPIs, recursive systems

Part I: The Beautiful Machine That Almost Worked

§1. Beer's Viable System Model

§1.1 The Context

When Stafford Beer published *Brain of the Firm* in 1972, organisational theory was dominated by mechanistic models derived from Taylor's scientific management (1911) and Weber's bureaucratic theory (1922/1947). These approaches treated organisations as machines to be optimised through division of labour, standardisation, and hierarchical control. Beer's Viable System Model departed from this tradition by drawing on Wiener's cybernetics (1948) and Ashby's law of requisite variety (1956) to reconceptualise organisations as adaptive systems maintaining identity through environmental perturbation. The reconception was architectural rather than terminological: viability, in Beer's argument, requires specific functional capacities regardless of the system's particular domain. As Beer put it, "the cybernetic science of effective organization is the science of organizational invariance" (1972, p. 7).

§1.2 The Strategic Purpose of This Analysis

This paper undertakes a genealogical investigation of organisational ontology, tracing the intellectual evolution from Beer's cybernetic insights through their critique to a phenomenological alternative. The analysis serves three connected purposes.

The first purpose is historical, since Beer's VSM, innovative as it was for its time, embedded an assumption about system-environment separation that subsequent theorists would render untenable. The recognition is not a criticism of Beer's achievement; it is a recognition that theoretical frameworks carry implicit ontological commitments that only become visible through their successors.

The second purpose is constructive, developing the concept of chiasmic organisation through phenomenological investigation that draws on Merleau-Ponty's late ontology (1945/2012, 1968), Barad's agential realism (2007), and empirical observation from software development and organisational practice. The chiasmic ontology treats organisations as patterns within fields they continuously co-create, rather than as entities observing antecedent environments.

The third purpose is strategic, deploying mathematical apparatus to establish rigorous boundaries on what formalisation can and cannot do with organisational phenomena. The horizon-mismatch theorem of §24 supplies the foundation; the formal-analogue articulations of §§27-29 extend the argument through the apparatus of sheaf cohomology, derived categories, and homotopy type theory. The approach has precedent in the tradition of impossibility theorems that have proven as valuable as constructive results, exemplified by Gödel's incompleteness theorems (1931), Arrow's impossibility theorem (1951), and Turing's halting result (1936).

The argument carries an urgent practical implication, since the widespread reliance on Key Performance Indicators, Objectives and Key Results, and similar metric-based management frameworks rests on a misapprehension of organisational reality. The misapprehension does not produce neutral inadequacy; it actively destroys the intelligence it seeks to measure, and it generates conditions in which the human cost is absorbed into burnout, organisational fragility, and the systematic erosion of the capacities the metrics purport to cultivate. The paper provides theoretical, phenomenological, and mathematical grounds for the claim, and it sets out in §22 the consequences for organisational practice.

§1.3 The Architecture of the Viable System Model

The VSM consists of five subsystems, each necessary for viability (Beer, 1979, 1985). System One comprises the primary activities that justify the system's existence, what Beer called "the collection of primary activities, the muscle and bone of the organization" (1985, p. 14). System Two provides anti-oscillatory mechanisms preventing destructive interference between operational units, the "sympathetic nervous system" providing stability through damping (Beer, 1979, p. 176). System Three handles internal regulation and resource allocation across operations, the "automatic pilot" maintaining homeostasis (Beer, 1972, p. 178). System Four conducts environmental scanning and adaptation: as Beer emphasised, "System Four is there to scan the total environment and to project the total system into it" (1979, p. 226). System Five maintains identity and defines purpose, what Beer called "the closure of the whole system" providing essential unity (1985, p. 128).

Beer's recursive insight was that each System One unit is itself a complete viable system containing all five functions. The recursive embedding distributes intelligence across levels rather than concentrating it at an organisational apex, and the model permits autonomous viability at every level of the recursion. The architecture broke new ground at the time of its publication, and it remains analytically powerful where contemporary organisational theory continues to treat hierarchical control as the default.

§1.4 The Cybernetic Foundations

Beer grounded the VSM in cybernetic principles, particularly in Ashby's law of requisite variety (1956, p. 207): "only variety can destroy variety." Beer's organisational interpretation read the law as a regulatory constraint: "the variety of the control system must be equal to or greater than the variety of the system being controlled" (1985, p. 89). Variety engineering, in Beer's apparatus, proceeds through two complementary mechanisms. Attenuation filters and aggregates incoming environmental variety, reducing what the regulator must process to what its measurement frame can resolve. Amplification multiplies regulatory variety through parallel processing, hierarchical decomposition, and adaptive subroutines, raising what the regulator can absorb to match the environmental complexity it must handle. Beer (1979, p. 88) argued that organisations fail when "the variety of the environment exceeds the variety of the system designed to regulate it," and the VSM architecture distributes variety-engineering operations across the five subsystems so each handles the variety appropriate to its register.

Beer's use of Ashby's law anticipates the formal result this paper develops in §24. There the law is sharpened into the horizon-mismatch theorem, which states the sense in which a regulator operating at temporal horizon shorter than its substrate cannot, even in principle, satisfy the variety condition Ashby's law imposes. The cascade argument of Part V then builds on the theorem to extend the variety constraint through more elaborate mathematical apparatus.

§1.5 The Departures from Traditional Organisational Theory

The VSM departs from traditional organisational theory along six axes. Each departure carries weight on its own terms; together they constitute a coherent reframing of what organisational structure is and how it operates.

The first departure is from hierarchical trees to recursive embedding. Traditional organisational theory, rooted in Weber (1922/1947) and Taylor (1911), conceived organisational structure as a hierarchical tree with clear superior-subordinate relationships. Beer's recursive embedding gave each operational unit all five subsystems internally, with intelligence distributed across levels rather than concentrated at an apex, and every recursive unit possessing the capacity for autonomous viability within its scope.

The second departure runs from concentrated to distributed intelligence: traditional models concentrated decision-making and environmental scanning at senior levels, following Fayol's (1916/1949) principle of centralisation. Beer's System Four function appears at every recursive level, permitting local environmental adaptation without central mediation. The recursive presence of System Four enacts the proposition Beer formulated as "every viable system contains and is contained in a viable system" (1979, p. 234).

The third departure is from fixed divisions to dynamic boundaries. Where traditional approaches imposed fixed departmental divisions with clear demarcations, the VSM treats boundaries as permeable membranes managing variety flows, so that boundaries become zones of variety management and the architecture admits dynamic reconfiguration in response to environmental demand.

The fourth departure is from top-down commands to autonomous coordination. The VSM replaces command hierarchies with autonomous units that achieve coordination through System Two's anti-oscillatory function. Beer described this as "maximum autonomy consistent with cohesion" (1985, p. 67), and the formulation encodes the entire coordination philosophy: autonomy and cohesion, which conventional management treats as a trade-off, are instead jointly maintained through System Two's damping function.

The fifth departure is from data flows to variety management. Traditional approaches understood information as data flowing up hierarchies for decision and commands flowing down for execution. Beer reconceptualised information processing as the cybernetic operations of attenuation and amplification distributed throughout the system, so that information itself becomes variety to be managed.

The sixth departure is from central planning to local adaptation. Where traditional management assumed centralised planning and execution, the VSM admits local adaptation: each recursive level responds directly to its environment, with overall coherence maintained through Systems Three and Five.

§1.6 Early Influence and Practical Adoption

The VSM's influence extended across multiple domains, the most ambitious application occurring in Chile, where Beer attempted to coordinate the national economy through Project Cybersyn between 1971 and 1973. The project, documented by Medina (2011), created a real-time economic management system using telex networks and operations rooms. It ended when the September 1973 military coup overthrew Allende and destroyed the project's institutional context.

In healthcare, the UK National Health Service attempted VSM implementation during the 1980s reforms, seeking to balance local hospital autonomy with system-wide coordination (Beckford, 1995). The NHS did not achieve full implementation, though VSM principles influenced subsequent healthcare management approaches. In industrial settings, manufacturing sectors, particularly steel and chemical industries, adopted VSM principles for production planning and control. Espejo and Harnden (1989) document implementations at British Steel Corporation and ICI, where the model coordinated complex production processes.

Management consultancies later incorporated VSM principles into their methodologies, often without acknowledgement of the cybernetic origins or the model's design limitations (Jackson, 2003, p. 89). The adoption is worth marking explicitly given this paper's later analysis: the same consultancies whose methodologies are diagnosed in §8.6 and §33.1 as displacement-as-defence drew, often invisibly, on Beer's apparatus when constructing their proprietary frameworks. The VSM's spread into management consulting demonstrated both the model's appeal and the conditions under which advanced cybernetic insight could be flattened into the very practices it had been designed to displace.

Academic development followed a steadier path, with the VSM generating extensive scholarly research and dedicated conferences and journals exploring its theoretical extensions and practical applications, and Espejo and Harnden (1989, p. 3) characterised Beer's work as "the most comprehensive attempt to apply cybernetic thinking to organizational design." The influence persists in contemporary approaches to enterprise architecture and adaptive management, though as Jackson notes the cybernetic origins are often unacknowledged in the contemporary literature.

Adoption demonstrated both the model's appeal and its practical difficulties. The implementation problems anticipated the second-order cybernetic critique that follows in Part II.

§2. The Limitations Embedded in the VSM

§2.1 The Observer Problem

The VSM presupposes that organisations can observe their environments objectively, since System Four's function rests on exactly that presupposition. Von Foerster (1973/1981, p. 258) would later render the presupposition untenable: "the environment as we perceive it is our invention." Beer's own formulation makes the dependence explicit: "System Four observes the total environment and models it" (1985, p. 56). The formulation requires three commitments: that the environment exists independently of observation, that observation does not alter the observed, and that models can faithfully represent external reality. Each commitment will be challenged by a different strand of second-order cybernetics: Maturana and Varela challenge the first through autopoiesis, von Foerster challenges the second through the constructivist account of perception, and Bateson challenges the third through the map-territory analysis.

§2.2 The Boundary Paradox

The VSM requires clear boundaries between system and environment. Beer himself poses the question without resolving it: “where does the organization stop and the environment begin?” (1979, p. 405). The model requires boundaries to fix what is inside (controllable) versus outside (observable), to determine where System Four stops modelling itself and begins modelling the environment, and to specify how System One units interface with external entities. As Luhmann would later argue, “the boundary is not a natural given but a construction of the system itself” (1995, p. 17). Beer’s framework cannot accommodate the constructive character of organisational boundaries because doing so would undermine the observer-stance the VSM presupposes.

§2.3 The Sequential Assumption

The VSM implies a temporal asymmetry at each recursive level, in which observation precedes modelling, modelling precedes decision, and decision precedes action. Beer’s summary “information flows up, control flows down” (1972, p. 160) captures the asymmetry. The asymmetry is internal to each level; the recursion across levels does not eliminate it. Each level requires the temporal sequence observation → model → decision → action. Weick’s empirical work on organisational sensemaking shows that the assumption is reversed at the very moment organisational actors are doing the work the framework presupposes. “People often don’t know what they think until they see what they say” (Weick, 1979, p. 165). Sensemaking is retrospective. The temporal asymmetry the VSM assumes does not survive close examination of organisational practice.

§2.4 The Centralisation Tendency

Despite the VSM’s commitment to distributed intelligence, System Four in practice tends towards centralisation. Beer’s requirement that System Four “model the total system and total environment” (1985, p. 226) imposes a variety condition that no single regulator can absorb (per Ashby’s law and the horizon-mismatch result of §24). Organisations consequently centralise System Four by default, because no single locus possesses the variety the model demands. Empirical studies of VSM implementation (Beckford, 1995; Jackson, 2003) document the same pattern: implementing teams strengthen System Four at higher recursive levels, weaken it at operational levels, and create de facto centralised intelligence functions that contradict Beer’s distributed-intelligence commitment.

The next sections trace how second-order cybernetics rendered each of the VSM’s observer assumptions untenable. Bateson, Maturana and Varela, von Foerster, and Spencer-Brown each attack a different facet of the observer stance, and the convergence of the four critiques produces the conditions under which the chiasmic ontology of Part III becomes available.

Part II: The Critique Through Second-Order Cybernetics

§3. Bateson’s Recursive Epistemology

§3.1 Information as Difference That Makes a Difference

Bateson’s epistemology of information and observation arrives in the essay “Form, Substance, and Difference,” collected in *Steps to an Ecology of Mind* (1972, p. 453): “information, the elementary unit of information, is a difference which makes a difference.” Information, in Bateson’s articulation, is relational, selective, and observer-dependent. It requires a difference, must make a difference, and exists only relative to whom it makes a difference to. The contrast with classical information theory is direct: information is not data, bits in channels, or signals to be transmitted; it is the difference-event itself. Bateson elaborates the position later in the same essay: “the territory never gets in at all. What gets in is news of difference, and news of difference is transformed (i.e., coded version) of events which preceded it” (1972, p. 459).

§3.2 The Map-Territory Relation

Bateson extends Korzybski's (1933) formulation "the map is not the territory" through recursive elaboration: "the map is not the territory, and the name is not the thing named, but the map is all we have access to" (1972, p. 455). Environmental models are maps with attendant distortions, simplifications, and abstractions, and organisations respond to these maps rather than to the territory itself, so that different observers construct different maps and consequently encounter different organisational realities. Bateson extends the analysis to push the regress further: "the name is not the thing named, and the thing named is not the 'thing' itself" (1979, p. 32). The regress is unterminating, since every representation is itself a representation of a representation, leaving System Four no ground from which to start.

The map-territory relation has a geometric form that renders the distortion exact rather than cautionary, and that exactness matters for the argument Part V will make formally. A globe cannot be represented on a flat sheet without distortion. This is not a limitation of cartographic technique that a better projection would remove; it is a theorem. Gauss's *Theorema Egregium* (Gauss, 1827) establishes that the curvature of a surface is intrinsic, preserved under any bending that does not stretch, so a surface with curvature cannot be laid flat onto a plane, which has none, without tearing or stretching somewhere. Every flat map of the round earth is therefore a projection, and every projection makes a choice about which property to preserve and which to sacrifice. The Mercator projection preserves angles, which is why it served navigation, and pays for the preservation by inflating area without bound toward the poles. An equal-area projection preserves area and pays by shearing every shape. No projection preserves angle, area, distance, and direction together, because the plane lacks a property the globe has, and the missing property has to surface somewhere as error. The distortion is not removed by a better map; it is only relocated by a different one.

This geometry gives the organisational dashboard its exact diagnosis, and three distinct breakdowns follow from the single fact that the flat map sacrifices what it cannot carry. Each breakdown disables a different organisational function, and each is invisible to the reader of the map, because a projection is internally consistent and gives no sign from within that anything has been lost.

The first breakdown disables coordination, for on the globe the shortest path between two points is an arc of a great circle. Drawn on the flat Mercator map, that shortest path appears as a curve bending far away from the straight line a ruler would draw between the same two points, which is why the shortest flight from northern Europe to the far east arcs over the pole and looks, on the flat map, like a detour. The person reading the flat map sees the ruler-straight line, takes it for the shorter one, and asks why anyone would take the bending route. The person reading the globe knows the bending route is the shortest one and the ruler-straight line is longer. Both are looking at the same journey; each calls the other's path irrational; neither can see that they are reading off different projections, because the projection is the medium of seeing and not an object within it. This is the structure of the recurring disagreement between organisational functions. One team sees the path through the actual system, which on the flat map of the metric looks like a wilful detour, and another team sees the flat-map straight line and asks why the first will not simply take the direct route. Each accuses the other of taking the long way, and the dashboard is the flat map on which the real shortest path through the territory is rendered as obstruction.

The second breakdown disables estimation, since the Mercator inflation of area toward the poles, far from being a small effect at the margins, grows enormous and counterintuitive at the scale of whole continents. Greenland is rendered at roughly the visual size of Africa and is in reality about fourteen times smaller. Russia is rendered at roughly twice its true area relative to equatorial land. A country near the equator that reads as small on the map can exceed in true area a high-latitude country that reads as large, so that two regions appearing equal on the flat map can differ several-fold in the ground they actually cover. The organisational consequence is that the apparent size of a task on the dashboard is a function of where the task sits in the projection rather than of the work the task contains. The visible, measured, high-latitude work is inflated and draws resources proportionate to its apparent size; the quiet infrastructure maintenance and the relational labour sit near the equator of the metric, read as small, and are resourced as though they were small, when in true area they are the larger continent. Two initiatives that read as equal commitments on the dashboard can differ by a factor of several in the labour they

actually demand. Management sizes the task by its projected area and is wrong in a direction the map itself conceals, over-committing to the inflated and under-committing to the compressed, with no warning from the instrument, because the instrument is internally consistent and reports apparent area as though it were real area.

The third breakdown is the one the first two are instances of, and it is the one Bateson named. The map is not a lie: a navigator crosses oceans with the Mercator chart and arrives, and a planner allocates with the dashboard and ships product. The flat map and the metric both do real work on the axis they preserve, and the error lies not in using the projection but in forgetting that it is one, in reading the inflated continent as vast in itself and the bending geodesic as perverse in itself, in mistaking the chosen preservation for the whole of the territory. The error and its correction are both geometric, so that one does not respond to the distortion of the Mercator map by abolishing maps; one responds by knowing which projection is in hand and what it was built to sacrifice, and by reaching for a different projection when the task changes from navigation to the measurement of land. The dashboard admits the same discipline and forbids the same forgetting.

The geometry also states, in advance and in intuitive form, the conservation result Part V will reach formally. The distortion that the flat map cannot avoid is not destroyed by any choice of projection; it is only moved. Mercator moves it onto area to keep angle; the equal-area projection moves it onto shape to keep area. The cartographer does not decide whether to distort but only what to distort, which is to say who bears the distortion. The variety the map cannot hold does not vanish for want of a projection to hold it; it relocates to whatever the projection sacrifices. In the organisation, as the later argument develops, the variety the metric cannot hold relocates in the same way, and the term it relocates onto is most often the people who absorb the difference between the flat map and the round world. The map analogy and the cohomological obstruction of Part V are the same structure stated at two registers: the impossibility of flattening a curved surface without remainder, and the impossibility of patching locally consistent sections into a globally consistent one without obstruction. The first is the intuitive face of the second.

§3.3 Double Binds in Organisational Communication

Bateson's theory of double binds was developed in the clinical context of schizophrenia in families (Bateson, Jackson, Haley, and Weakland, 1956; Bateson, 1972, pp. 206-227). The structure transposes to organisational communication when the same conditions hold: a primary injunction, a secondary injunction contradicting the first, an inability to escape the field, and a prohibition on commenting on the contradiction. Organisations frequently produce exactly this structure: a primary injunction directs the agent to action A, while a secondary injunction, often emerging through a different channel or organisational subsystem, directs the agent to incompatible action B; the agent cannot leave the field, since resignation is materially or socially foreclosed; and explicit articulation of the contradiction is institutionally prohibited.

Bateson offers concrete organisational forms (1972, p. 208): "Be spontaneous" combines a primary injunction (act spontaneously) with the secondary feature that following the injunction is by definition not spontaneous. "Ignore these instructions" is the same form in compressed shape. "Be independent" is the form mapped to organisational autonomy: the directive to behave independently, issued by a hierarchical authority, undermines the independence it commands.

System Two (coordination) operates within a single logical level, whereas double binds straddle two: the primary injunction and the meta-injunction prohibiting comment on the contradiction. System Two cannot reach the meta-level, by its own functional definition. As Bateson observed, "the victim cannot make the metacommunicative statement, 'this is a paradox'" (1972, p. 212).

§3.4 Levels of Learning and Logical Types

Bateson identified hierarchical learning levels grounded in Russell and Whitehead's theory of logical types (*Principia Mathematica*, 1910-1913; Bateson, 1972, pp. 279-308). Learning Zero is mechanical response without learning; Learning One acquires specific responses to specific contexts; Learning Two reframes contexts and adapts the patterns of Learning One; Learning Three reframes the patterns of Learning Two themselves; and Learning Four operates at the level of evolutionary change in entire ways of categorising experience.

The VSM was designed for Learning Two, where variety engineering modifies the regulatory parameters that Systems One through Four operate within, and System Five modifies the framework that defines System Four's environment. What the VSM cannot reach is Learning Three: questioning the framework itself. Bateson observes that "Learning III is likely to be difficult and rare even in human beings" (1972, p. 303), and the difficulty is built into the VSM at the level of its architecture. Learning Three would require questioning the observer-observed separation that the VSM presupposes. The VSM cannot accommodate the questioning without abandoning its own architecture.

§3.5 The Recursive Nature of Epistemology

Bateson developed the recursive epistemology in the late work: "the observer must always be included in the observed" (1979, p. 201). He named the resulting meta-pattern "the pattern which connects" (1979, p. 8). Two consequences follow. First, when System Four observes its environment, this is environment observing itself through the recursive structure that contains the observer. Second, boundaries are not given prior to observation; observation creates them. The chiasmic-organisation argument of Part III turns on exactly these two consequences. Bateson's closing formulation captures the position: "the mental world, the mind, the world of information processing, is not limited by the skin" (1979, p. 242).

§4. Maturana and Varela: Autopoiesis and Organisational Closure

§4.1 The Santiago Theory of Cognition

Maturana and Varela, working in Santiago during the period of Beer's Cybersyn project, developed autopoiesis theory. The theory challenges the VSM's presupposition that observation imports information from environment into system. Maturana and Varela's formal definition appears in *Autopoiesis and Cognition* (1980, p. 78): "an autopoietic machine is a machine organized (defined as a unity) as a network of processes of production (transformation and destruction) of components that produces the components which: (i) through their interactions and transformations continuously regenerate and realize the network of processes (relations) that produced them; and (ii) constitute it (the machine) as a concrete unity in the space in which they (the components) exist by specifying the topological domain of its realization as such a network."

§4.2 Organisational Closure and Structural Determinism

Maturana and Varela's theory was developed for biological systems, with the nervous system as the paradigm case. As they state, "the nervous system is organizationally closed, and therefore, no inputs or outputs occur in its operation" (1987, p. 169). The transposition to organisations requires justification, which Luhmann (1995) supplied by extending autopoiesis to social systems. Luhmann argues that organisations exhibit operational closure in the analogous sense: they generate their own categories of legitimate action through internal operations. The formal-analogue argument (per §25.3) is that organisations exhibit the autopoietic features Maturana and Varela identify in the nervous system, with Luhmann's framework providing the bridge from biological to social application.

Three consequences follow. Organisations do not receive inputs in the cybernetic sense; they are perturbed. Perturbations trigger internal changes, but the changes are determined by internal structure, not by the perturbations themselves. Maturana captures the relation: "the structure of the environment cannot specify the

changes of state of the system; it can only trigger them” (1978, p. 42).

For Beer’s System Four, the consequence is direct: System Four cannot scan the environment as if it were an external object, cannot import information across an inside-outside boundary, and can only undergo internal changes triggered by perturbations, where the triggers do not determine the changes but the internal structure does.

§4.3 Structural Coupling

Structural coupling differs from adaptation: where adaptation implies an organism conforming to a fixed environment, structural coupling describes mutual specification. Organisation and environment specify each other through the history of their interactions, and the changes that result are congruent with environmental conditions without being caused by them. Maturana and Varela (1987, p. 75) formulate the position: “the continued interactions of a structurally plastic system with its environment will result in selections of the system’s structure such that it will ‘fit’ into its environment.” The fit is mutual specification, not conformity to a pre-existing shape.

Varela (1979, p. 48) sharpens the point: “in describing autopoiesis, one is not describing the properties of components, but relations which must be regenerated constantly.” The shift from properties to relations is the move that distinguishes autopoiesis from earlier systems-theoretic approaches.

§4.4 Cognition as Effective Action

Maturana and Varela characterise cognition: “living is a process of cognition” (1987, p. 29). Cognition is effective action in the domain of existence rather than information processing in an internal representation. Organisations cognise by maintaining their organisation, not by computing models of an external world. As Varela formulated the position, “the concrete realization of the living system determines what can be a perturbation and what cognitive domain it can have” (1979, p. 268). Beer’s System Four assumes cognition as representation, whereas the autopoietic account treats cognition as effective action, the bringing forth of a world through embodied engagement (the language of “bringing forth” developed further in Varela, Thompson, and Rosch, 1991).

§4.5 The Observer as a Living System

Maturana formulates the position: “everything said is said by an observer” (1978, p. 31). The observer is itself a living system with its own autopoiesis. Observation occurs from within an autopoietic organisation, not from a position external to it, and the consequences follow without further premise: observations are structural determinations of the observer rather than reports of the observed, so that different observers, constituted differently, encounter different realities, and no privileged observation point is available, because every observation is the operation of some autopoietic system. Maturana and Varela’s closing formulation captures the position: “we do not see what we do not see” (1987, p. 242).

§5. Von Foerster: Second-Order Cybernetics

§5.1 The Cybernetics of Observing Systems

Heinz von Foerster, who edited the proceedings of the later Macy Conferences (conferences 6-10, 1949-1953), later coined the term “second-order cybernetics” in lectures during the late 1960s and formalised the position in his Biological Computer Laboratory work at Illinois. The new register identifies a constraint internal to the cybernetic programme: cybernetics studying observed systems must itself be studied as a cybernetics of observing systems. Von Foerster’s distinction (1981, p. 258) is direct: “first-order cybernetics is the cybernetics of observed systems. Second-order cybernetics is the cybernetics of observing systems.” The shift relocates the observer inside the system being studied, and observation becomes participation rather than spectatorship. Von Foerster’s formulation: “objectivity is the delusion that observations could be made without an observer” (2003, p. 285).

§5.2 The Construction of Reality

Von Foerster's central claim (1981, p. 275) is that "the environment as we perceive it is our invention." The position avoids solipsism. The claim is not that nothing exists outside perception; the claim is that perception constructs rather than receives, that reality is enacted in observation, and that organisations create the environments they then perceive. Von Foerster's 1973 paper "On Constructing a Reality" (reprinted in *Observing Systems*) demonstrates the neurophysiological case: the nervous system is operationally closed, no information enters from outside, and all perception is internal construction.

For organisations, the consequence is that System Four constructs environmental opportunities through its distinctions rather than discovering them as antecedent objects. The categories of risk, opportunity, threat, and adaptation are not given by the environment; they are imposed on the environment by the organisation's observing apparatus.

§5.3 Eigenforms and Organisational Stability

Von Foerster's 1976 paper "Objects: Tokens for (Eigen-)Behaviors" (reprinted in 2003, p. 261) introduces eigenforms, stable forms arising from recursive operations: "objects are tokens for eigenbehaviors." Stable patterns emerge from recursive interactions, the patterns appear as objects or features of an environment, and yet they are eigenbehaviours: the stable products of recursive operation rather than external givens. Von Foerster's mathematical formulation (2003, p. 264): "the eigenvalues are the stabilities that occur when an operation is recursively applied to itself."

Beer's System Five (identity) operates as an eigenform: a stable pattern emerging from and maintained by recursive organisational operations, rather than a fixed property of the organisation. The recognition matters for the impossibility argument of §7: an organisation's identity is something the organisation does, recursively, and reform requires modifying the recursion itself rather than substituting one fixed identity for another.

§5.4 Ethics and Responsibility

Von Foerster (1992, p. 14) formulates two imperatives. The ethical imperative: "act always so as to increase the number of choices." The aesthetic imperative: "if you desire to see, learn how to act." The imperatives invert VSM variety management. Beer's variety engineering reduces variety to manageable levels; von Foerster's ethical imperative increases the number of choices available. Beer's architecture proceeds from observation to action; von Foerster's aesthetic imperative makes action prior to observation, with seeing made possible by acting. Von Foerster's closing formulation captures the recursive structure: "I am the observed relation between myself and observing myself" (2003, p. 289).

§5.5 The End of the Observer Illusion

Von Foerster drives the position to its conclusion: "there are no observers; there are only actors who construct their universe through their actions" (2003, p. 287). The consequences for the VSM are direct, since the observer stance the architecture requires does not exist as a feature of any actual organisation. What appears as observation is action, what appears as scanning is construction, and what appears as objective modelling is the recursive operation of an observing system constructing the world it then claims to describe.

§6. Spencer-Brown: The Laws of Form

§6.1 The Calculus of Indications

Spencer-Brown's *Laws of Form* (1969) develops a calculus of distinctions. The calculus treats organisational boundaries as acts of distinction made by an observer rather than as divisions among pre-existing entities. The opening axiom (1969, p. 1) sets out the structure: "we take as given the idea of distinction and the idea of indication, and that we cannot make an indication without drawing a distinction." The calculus rests on two laws:

- *Law of Calling*: the value of a call made again is the value of the call.
- *Law of Crossing*: the value of a crossing made again is not the value of the crossing.

The laws encode the formal properties of distinctions: distinctions can be referenced repeatedly without altering the distinction itself, but crossings (acts of moving across the distinction) cancel themselves when repeated.

§6.2 Boundaries as Distinctions

Spencer-Brown's foundational formulation (1969, p. 1) states that "distinction is perfect continence." The act of distinction creates what it distinguishes; boundaries do not separate pre-existing entities but constitute the entities they distinguish, with every boundary an observer's distinction. Organisational boundaries follow the same pattern: department boundaries do not separate antecedent departments, since the boundary-making constitutes the departments. Customer-versus-employee distinctions, product-versus-service categories, internal-versus-external divisions: each is the calling-into-being of what the distinction marks. Spencer-Brown's later remark (1969, p. 69) gives the cosmological extension: "the universe is the inside of the severance of the first distinction."

§6.3 Re-Entry and Organisational Recursion

Spencer-Brown's re-entry construction (1969, Chapter 11) resolves a problem Beer's VSM cannot explain: organisational self-reference without infinite regress. When a distinction re-enters its own indicational space, the system can refer to itself without generating an infinite tower of meta-levels. The construction introduces imaginary values to resolve the resulting oscillation. Organisations exhibit re-entry when they observe themselves observing, when System Four models itself modelling, when boundaries include their own drawing. The phenomenon is recognisable in any organisational moment where the activity of describing the organisation becomes itself part of what is described, and the system Beer attempted to formalise as observation from outside dissolves into recursive self-reference from within.

§7. The Accumulated Critique

Five impossibilities follow from second-order cybernetics. Objective observation is impossible, per von Foerster's recognition that the observer is part of the system. Information transfer is impossible in the cybernetic sense, per Bateson: only difference-events occur, and they are observer-coded. Environmental inputs are impossible, per Maturana and Varela: organisations are organisationally closed. Pre-existing boundaries are impossible, per Spencer-Brown: boundaries are distinctions made by observers. External control is impossible, as the consequence of the first four. The observer stance the VSM requires cannot exist. Organisations are patterns within the field they themselves participate in producing, the chiasmic structure Part III develops.

Before turning to that structure, the paper takes up a question the second-order critique sharpens but does not answer. Organisations continue to operate as if the observer-stance were available, continue to deploy frameworks the operators collectively know to be inadequate, and continue to invest substantial resources in maintaining the gap between official framework and actual practice. The configuration is stable, well-documented across institutional contexts, and resistant to reform from within. The next section develops the formal structure of this state and the institutional mechanisms by which reform within it is foreclosed.

§8. The Terminal Attractor of Epistemic Denial

The configuration this section develops sits between the second-order critique of Part II and the chiasmic ontology of Part III. The section is positioned here because this state is what becomes visible when an organisation maintains a first-order framework while operating in conditions the second-order critique has rendered untenable. The state is stable, and internal reform within it is foreclosed. The mechanisms of foreclosure are institutional rather than psychological, and they operate invariantly across widely different organisational contexts.

§8.1 The N-th Order Epistemic Acknowledgement State

The N-th Order Epistemic Acknowledgement State is the configuration in which an organisation $\mathcal{O}$ maintains an operative framework F while its constituent agents hold tacit knowledge K that the framework cannot capture. The state is characterised by four features.

Feature 1 (Representational Inadequacy). The tacit knowledge K is not contained in the operative framework F :

$$K \not\subset F$$

The exclusion is built into how the framework operates, since the framework works in one register (formal procedures, metrics, audit-compatible language) while the tacit knowledge belongs to another (lived experience, pattern recognition, embodied skill). The two registers are not different by degree; they are incommensurable, and the framework cannot accommodate the tacit register without ceasing to function as a framework.

Feature 2 (Universal Recognition). Every participant recognises that $K \not\subset F$:

$$\forall a \in \mathcal{O}, \quad a \text{ knows that } K \not\subset F$$

The recognition is distributed across the organisational field rather than localised to a dissident minority or held privately by senior leadership; it is common knowledge in the precise sense the next feature formalises.

Feature 3 (Recursive Acknowledgement to Depth N). Each participant knows that every other participant recognises the inadequacy, and knows that they know, to arbitrary recursive depth:

$$\forall a, b \in \mathcal{O}, \forall n \leq N, \quad a \text{ knows}^n \text{ that } b \text{ knows that } K \not\subset F$$

The recursion can extend to arbitrary N . The structure is the game-theoretic condition of common knowledge (Lewis, 1969; Aumann, 1976) applied to framework inadequacy: each participant knows, knows that the others know, knows that the others know that they know, and so on.

Feature 4 (Mandatory Framework Persistence). Despite Features 1 through 3, the framework F persists as the sole legitimate language of organisational discourse. Alternative articulations using K -language lack institutional standing. The persistence is not maintained by any agent's preference. It is maintained by the architecture: by audit requirements, regulatory compliance, evaluation procedures, and the credentialing systems that determine which utterances count as organisationally legitimate.

An organisation satisfying Features 1 through 4 simultaneously occupies the N-th Order Epistemic Acknowledgement State. This state is no exception or pathological case but a structurally stable attractor in the institutional state-space, and the historical record shows that organisations enter it across a wide range of contexts and remain in it for extended periods.

§8.2 Homeostatic Mechanisms

The N-th Order State maintains itself through three institutional mechanisms. Each mechanism is observable in organisational practice, documented in the empirical literature, and reducible to a specific function within this state.

§8.2.1 Credentialism as Variety Attenuation

Information processing in the N-th Order State filters incoming variety through credential conditions before content. Let $C: \text{Sources} \rightarrow [0, 1]$ map information sources to credential scores. Let V represent information value. In healthy systems, information processing follows $\text{Uptake} = f(V)$: uptake scales with value. In the N-th Order State, processing becomes $\text{Uptake} = C(\text{source}) \cdot V$, with C operating as a gating function. When $C(\text{source}) \approx 0$ for sources articulating K in non- F -compatible terms, even high-value information receives effectively zero organisational uptake. The mechanism is documented in Collins and Evans (2007) on expertise and in Alvesson and Robertson (2006) on elite identities in consulting firms.

The credentialism serves a dual function: explicitly, it maintains quality standards for organisational knowledge; implicitly, it filters information through compatibility with F , ensuring that only F -compatible knowledge gains organisational legitimacy. Information contradicting F , however valid in substantive terms, arrives from sources lacking the requisite credentials and consequently attenuates to insignificance in organisational decision-making (Abbott, 1988). A developer with fifteen years of system knowledge is overridden by a management consultant from a major firm who lacks the domain knowledge but possesses the institutional standing the framework recognises.

§8.2.2 Performance Amplification as Divergence Response

As the gap between official framework and lived reality widens, the organisation intensifies its commitment to the framework rather than questioning it. Metrics proliferate. Ceremonies elaborate. The worse F performs, the more organisational labour is devoted to demonstrating compliance with F . The dynamic is one of positive feedback: divergence between K and F generates additional investment in F -maintenance, which generates additional divergence as more of the actual work is forced into F -incompatible shadow practices.

Let $D = \int_0^1 |K(a) - F(a)| da$ represent aggregate divergence between tacit and formal across the organisational field. Let E_F represent organisational labour invested in F -compliant performance. In healthy systems, recognition of inadequacy reduces investment: $\partial E_F / \partial D < 0$. In the N-th Order State, the relation inverts and accelerates: $\partial E_F / \partial D > 0$ and $\partial^2 E_F / \partial D^2 > 0$. The mechanism is documented across empirical literatures: Power (1997) shows that audit inadequacies generate more audit rather than less; Strathern (2000) shows that metric inadequacies generate more metrics; Alvesson and Spicer (2016) describe the broader pattern of functional stupidity in which framework intensification substitutes for framework questioning.

Feynman's (1974) characterisation of cargo cult science captures the dynamic in compressed form: if the planes are not landing, build longer runways and wave the signal flags more vigorously. The intensification operates with full ritual seriousness within the framework while the framework's relation to the substrate it claims to regulate continues to deteriorate.

§8.2.3 Source Isolation as System Defence

Individuals who persistently articulate the gap between K and F are marginalised through subtle exclusion rather than explicit punishment. The marginalisation operates through three characteristic formulations: "they are not team players," "they are too negative," "they do not understand the big picture." The problem, organisationally, is not the analytical content of what they articulate but the institutional pressure their articulation places on the stability of the state. Morrison and Milliken (2000) document the pattern as organisational silence; Ryan and Oestreich (1991) describe the broader phenomenon of fear-driven communication suppression. The isolation operates as an immune response: the state treats articulation of K as a threat to homeostasis and deploys institutional mechanisms to neutralise the threat without acknowledging the analytical content of what is being articulated.

§8.3 The Three-Mechanism Reform Foreclosure

The N-th Order State resists reform from within by its very architecture, the resistance following from the three mechanisms operating in coordinated foreclosure: any reform proposal expressed in organisational discourse is processed by one of three modalities, each of which preserves F and the architectural feature targeted by the reform.

Theorem (Three-Mechanism Reform Foreclosure). Let $\mathcal{O}$ be an organisation in the N-th Order State with operative framework F , tacit knowledge K , and $K \subset F$. Let R be any reform proposal articulated within organisational discourse. Then R is processed by one of three mechanisms.

Mechanism 1 (Absorption). If R is articulated in F -language, then R is compatible with the framework's existing categories. The reform's content can be expressed without challenging the framework's foundations. The framework absorbs R : it implements the reform's surface content while leaving the architectural condition $K \subset F$ unchanged. Example: a reform demand for better velocity tracking is absorbed into the existing velocity-as-meaningful-metric paradigm. The reform succeeds locally and the architectural condition is preserved.

Mechanism 2 (Attenuation). If R is articulated in K -language, in terms the framework does not recognise as legitimate, then R lacks the institutional standing required for uptake. The reform is filtered through credentialism (Mechanism 2 of §8.2.1): only F -compatible justifications gain regulatory or evaluative standing. The reform is attenuated regardless of substantive content. Example: a reform demand that velocity metrics create perverse incentives is dismissed as "not understanding how we work" because the dismissal-grounds operate within F and the demand operates outside it.

Mechanism 3 (Decomposition). If R attempts a hybrid articulation, partly in F -language and partly in K -language, then the mechanisms decompose R into its constituent registers. The F -compatible elements are processed by Mechanism 1 (absorbed); the K -elements are processed by Mechanism 2 (attenuated). The hybrid does not survive the decomposition: what passes through is only the F -compatible content that does not challenge the framework. Example: a reform proposal that "we should improve velocity tracking (F -compatible) while addressing the metric-gaming behaviour it incentivises (K -element)" produces an organisational response that implements "improved velocity tracking" and ignores the gaming critique.

Corollary. For any reform R expressed in organisational discourse, the architectural feature targeted by R , the condition $K \subset F$, is preserved across the institutional processing of R . The N-th Order State is therefore stable against internal reform. Escape requires either external shock that bypasses the institutional processing (regulatory intervention, financial crisis, public disclosure) or ontological reframing that dissolves the K/F distinction itself.

The theorem is proved in the proof register by exhaustion of the three discourse possibilities (F -language, K -language, hybrid) combined with the empirical observation that each possibility activates one of the three mechanisms. The empirical grounding for the mechanisms (credentialism as variety-attenuation, performance amplification as divergence response, source isolation as system defence) was established in §8.2 and is historically validated across the cases of §8.4.

The Sartrean reading of this state as institutional bad faith is available but requires care. Sartre's analysis (1943/1956, pp. 86-116) treats bad faith as individual: the café waiter performing the waiter role with exaggerated precision while knowing himself to be more than this role exemplifies simultaneous knowing and not-knowing held in place by refusing to examine the contradiction. The N-th Order State manifests something structurally analogous at institutional scale. The distinction with Sartre is important and worth marking: the institutional bad faith is not the aggregation of individuals' Sartrean bad faith. It persists even when no individual is in Sartrean bad faith. It is a property of the architecture's discourse-rules, not of any agent's self-deception. Fleming and Spicer (2003) analyse the cynical-distance variant in which agents maintain reflexive distance from their organisational performances while continuing to perform them; the analysis is consistent with the present account, but the mechanism rests in the architecture rather than in the agents' relation to it.

§8.4 Historical Validation

The N-th Order State and the three-mechanism foreclosure operate invariantly across widely different organisational contexts. The invariance is the argumentative ground for treating the state as a feature of institutional life rather than a contingent property of any particular case, and three historical cases instantiate the apparatus.

§8.4.1 Soviet Production Statistics (1920s-1980s)

Soviet enterprises operated under a command economy that required fulfillment of centrally planned production quotas. Berliner (1957) and Kornai (1992) document how enterprises systematically reported meeting quotas everyone in the chain knew were fictional. The recursive acknowledgement structure was four-deep. Factory managers knew their statistics were inflated. Workers knew the managers knew. Ministry officials knew the factories were inflating. Central planners knew the ministry was aggregating inflated returns. And every participant knew that every other participant knew.

Yet the reporting continued because the Soviet system required the performance of plan fulfillment. The gap between actual production (K) and official statistics (F) was openly acknowledged in private while F was maintained in public. The performance itself became the product (Nove, 1977). The three mechanisms operated in coordinated foreclosure across the period: credentialist defence (only Gosplan-issued framework counted as legitimate; reform proposals articulated outside the Gosplan register lacked institutional standing); performance amplification (planning bureaucracies grew, reporting frequencies increased, framework intensification substituted for framework revision through every reform episode from the New Economic Policy through the Liberman reforms); and source isolation (the dissident analytical economists were marginalised institutionally, with substantive critiques absorbed only as ammunition for inter-bureaucracy conflict rather than as a route to framework revision).

The system persisted in this stable N-th Order state for decades. External shock (glasnost policies opening the public-acknowledgement channel beyond the architecture's processing capacity) triggered rapid collapse (Goldman, 1991). Internal reform had been foreclosed throughout the seventy-year period by the three mechanisms operating in coordination.

§8.4.2 NASA Challenger Disaster (1986)

Vaughan's (1996) investigation of the Challenger launch decision reveals the same state in an engineering organisation operating under conditions Beer's VSM was designed to handle. Engineers had documented O-ring concerns for years; management had received the documentation; Thiokol engineers argued against the cold-weather launch the night before the disaster. Every participant in the chain recognised the safety concerns the documentation articulated. The recursive acknowledgement was as deep at NASA as it had been in the Soviet planning ministries.

The launches continued because organisational commitment to schedule and budget targets exceeded the institutional capacity to act on explicit safety knowledge. Vaughan termed the resulting pattern "normalisation of deviance," the slow drift through which known violations of safety standards become routine. Reading the case through the three-mechanism apparatus: "do not launch" (articulated in K -language about engineering judgement and felt risk) was attenuated through the absence of F -compatible justification within the flight readiness review process; "is this within flight readiness review process?" (articulated in F -language about institutional procedure) was absorbed without disturbance to the procedure itself; the hybrid articulations Thiokol engineers attempted in the launch-readiness meeting were decomposed into the F -compatible elements (which entered the launch decision) and the K -elements (which did not).

The state was stable until external shock (the physical disaster, the subsequent Rogers Commission with subpoena power beyond NASA's institutional processing capacity) generated reform from outside the system that could not be generated from within.

§8.4.3 The 2008 Financial Crisis Risk Models (1990s-2008)

Financial institutions deployed Value-at-Risk models for risk assessment despite widespread recognition of their inadequacy. MacKenzie (2011) and Taleb (2007) document the epistemic structure with depth comparable to Vaughan's analysis of NASA. Traders knew the models did not capture tail risk. Risk managers knew the traders knew. Regulators knew the risk managers knew. Academic literature documented the limitations publicly throughout the 1990s and 2000s.

Yet the models remained operative because they provided the legitimate language for risk assessment, “what one uses” for regulatory compliance and internal reporting (MacKenzie and Spears, 2014), and the three mechanisms operated as before. Credentialism: only Basel-compliant model risk frameworks qualified as legitimate, and substantive critiques articulated outside the Basel register attenuated to insignificance in regulatory practice. Performance amplification: the response to LTCM's 1998 collapse was elaboration of Value-at-Risk methodology rather than questioning of the assumption that tail risk admits parametric modelling. Source isolation: the analysts who articulated tail-risk concerns in non-VaR language (Taleb, Roubini, Rajan) were institutionally marginalised in the pre-2008 period as fringe figures, with their analytical content unable to gain regulatory standing until after the crisis demonstrated its substantive correctness.

When the crisis arrived, every participant claimed surprise despite the widespread prior knowledge of model limitations. The surprise was not about the models breaking down; it was about their inadequacy becoming publicly undeniable in a manner that exceeded the institutional processing capacity (Engelen et al., 2011). External shock from the financial collapse itself generated regulatory reform that decades of internal critique had not.

The three cases converge in their structure across three widely different institutional contexts: a planned economy across seventy years, a government engineering agency across two decades, and a deregulated financial sector across two decades. The cases share no domain features beyond the formal ones the N-th Order State and the three mechanisms specify. This invariance across structure is the argument: the state is not a feature of Soviet planning, NASA engineering culture, or financial regulation but a feature of institutional architecture, and it appears in any organisational form that combines formal frameworks with substantive activity those frameworks cannot capture.

§8.5 Phase Transition Dynamics

Organisations enter the N-th Order State through a four-stage transition. The stages are sequential and cumulative: each establishes the conditions for the next.

Stage 1: Local Adaptation. Small gaps emerge between official framework and actual practice. A team discovers the ticketing system does not capture their workflow and maintains a shadow spreadsheet. Another team finds the testing framework does not fit their architecture and develops informal practices. Individual managers find prescribed metrics do not capture meaningful progress and supplement formal reports with informal updates. The adaptations seem reasonable, even necessary, local optimisations improving effectiveness that no participant views as subversive but as pragmatic responses to framework limitations. The organisation explicitly tolerates the gap: “use the official system for reporting, but do what actually works” (Orlikowski, 2000).

Stage 2: Proliferation. More teams develop similar workarounds. The shadow systems grow elaborate, sometimes more so than the official frameworks. Knowledge about how things actually work circulates informally through hallway conversations, tribal knowledge, and onboarding that happens outside official channels (Brown and Duguid, 1991). Leadership acknowledges the duality implicitly: “we know the process is bureaucratic; just get the real work done and we will handle the paperwork.” The gap between K and F is recognised, and not problematised (Zuboff, 1988).

Stage 3: Performance Becomes Primary. At a critical threshold, maintaining both systems becomes the primary organisational work. Substantial time is devoted to translating actual work into official framework language for reporting. Teams spend more time on optics, making work visible in F -terms, than on the work

itself. Meetings focus on what they will report rather than what they will do (Alvesson and Spicer, 2016). The organisation has crossed the phase boundary: it is no longer using F to accomplish goals, it is maintaining performance of using F while goals get accomplished despite F . As one practitioner put it, “we spend 60% of our time making the other 40% look good in the reports” (Graeber, 2018, p. 97).

Stage 4: Terminal State. The shadow system and official framework achieve complete separation and reciprocal stabilisation. Every participant operates in a shadow system, every participant reports through the official framework, every participant knows that every other participant does this. The performance is complete, stable, and recognised as performance. Escaping this state requires external shock because all internal language is F -language, and F -language cannot express its own inadequacy (Brunsson, 1989).

The transition can be characterised through phase-space concepts. The N-th Order State corresponds to a meta-stable attractor in the institutional state-space. Let x represent institutional state and x^* the N-th Order equilibrium. Perturbations δ below a critical threshold δ_c are absorbed by the three-mechanism apparatus: the system returns to x^* . Perturbations exceeding δ_c bypass the apparatus and the system exits the basin of attraction. The argument is qualitative, characterising the shape of the phase space rather than modelling specific trajectories. The historical cases of §8.4 show the empirical pattern: long stability punctuated by sudden catastrophic transition rather than gradual degradation. The transition discontinuity is inherent rather than accidental: the three mechanisms maintain x^* against any perturbation they can process, and the system exits this state only when perturbations exceed the apparatus’s processing capacity.

§8.6 The Demand for Better Metrics as Displacement

Organisations in terminal epistemic states exhibit a characteristic response to a framework’s inadequacy: demand for more elaborate versions of the inadequate framework. “These metrics are not working; we need better metrics.” “This methodology fails; let us try a more rigorous methodology.” “The model does not capture complexity; we need a more complex model.”

The response sounds reasonable, since improving inadequate tools makes sense as a default heuristic, but it displaces an ontological problem into the methodological domain (Power, 1997). The issue is not that current metrics are poorly designed; it is that organisational phenomena exceed what metrics can capture. The response misidentifies the category of problem, treating ontological limitation as methodological insufficiency (Flyvbjerg, 2001).

The pattern is recognisable in OKR implementations, where an organisation adopts Objectives and Key Results and, after two years, the workforce recognises that OKRs create perverse incentives and do not capture meaningful progress. Rather than questioning whether objective-setting frameworks can capture emergent organisational work, the consultancies propose “improved OKR implementation,” “transitioning to Evidence-Based Management,” or “integrating with North Star metrics” (Pfeffer and Sutton, 2006). Each new framework arrives with elaborate theoretical backing, promises to repair the previous framework’s limitations, and enters with organisational fanfare. The consultancy industry depends on the cycle: frameworks must prove inadequate enough to justify replacement, while remaining appealing enough that the inadequacy does not generalise into questioning frameworks-as-such (Kieser, 2002). Each new framework eventually reveals the same inadequacy, by which point the organisational memory has faded, the consultancies have moved on, and new consultancies arrive proposing the next framework. The cycle perpetuates through the framework’s inadequacy rather than despite it, since each round of inadequacy justifies the next attempt (March, 1991).

The displacement is not the product of stupidity; it is anxiety management (Heidegger, 1927/1962). Acknowledging that organisational phenomena exceed formalisation entirely creates existential difficulty for management structures premised on measurement and control. The institutional response prefers the belief that the right framework has not been found yet to the recognition that no framework can exist (Alvesson and Willmott, 2002).

The Part V mathematical argument bears on the displacement directly. The cascade demonstrates that even the most expressive available formalisms, sheaf cohomology, derived categories, homotopy type theory, articulate rather than capture the chiasmic features of organisational phenomena. If those formalisms cannot capture the phenomena (and Part V's argument is that they cannot), then elementary formalisms (arithmetic for counting tickets, basic statistics for velocity calculations) fall short a fortiori. The demand for "better metrics" amounts to the demand for a more elaborate formalism within a register the horizon-mismatch theorem of §24 has already established as inadequate to the substrate. The demand is not unreasonable in its register; it is incoherent in its category.

Part III: The Phenomenological Turn to Chiasmic Organisation

§9. Merleau-Ponty's Chiasmic Ontology

§9.1 The Chiasm

Merleau-Ponty's late philosophy, in *The Visible and the Invisible* (1968), supplies the ontological framework for chiasmic organisation. The chiasm, from the Greek chi (X), names the crossing or intertwining structure underlying perception and action. Merleau-Ponty's central formulation (1968, p. 139): "the chiasm, reversibility, is the idea that every perception is doubled with a counter-perception. One no longer knows who speaks and who listens." The reversibility is a simultaneous ambiguity rather than an alternation between distinct states, both aspects present at once. This is not a property some organisations have; it is what organisations are once the observer-observed separation is recognised as untenable.

§9.2 The Touching-Touched Illustration

The classical illustration is tactile self-perception: when the right hand touches the left, Merleau-Ponty observes that "my hand while it is felt from within is also accessible from without. If it takes its place among the things it touches, [it] is in a sense one of them" (1968, p. 133). The touching hand is also touched. Subject and object are aspects of a single phenomenon. The reversibility occurs without coincidence: the touching hand and the touched hand are not identical at any moment, yet neither precedes the other as a pre-existing entity. The *écart* (the gap) between touching and touched is what makes both possible. As Merleau-Ponty puts it elsewhere, "there is a circle of the touched and the touching, the touched takes hold of the touching" (1968, p. 148).

§9.3 Flesh as Element

Merleau-Ponty introduces *flesh* (*chair*) as the element of reversibility itself. His own definition (1968, p. 139): "the flesh is not matter, is not mind, is not substance. To designate it, we should need the old term 'element,' in the sense it was used to speak of water, air, earth, and fire, that is, in the sense of a general thing, midway between the spatio-temporal individual and the idea." Flesh is what subject and object share as their common tissue, the formative medium of both. Merleau-Ponty's later formulation (1968, p. 147) describes flesh as "the formative medium of the object and the subject." The element is not added to the objects it connects; it is the medium through which they emerge as objects at all.

§9.4 Organisational Implications

The chiasmic ontology has direct organisational implications: at every organisational boundary, inside becomes outside and outside becomes inside. The customer service representative embodies both company and customer in a single role. The salesperson creates needs while discovering them. As Merleau-Ponty notes, "the boundary is not a line but a zone of reciprocal insertion" (1968, p. 123). The recognition recovers what Beer's VSM could not formalise: organisational boundaries that constitute the entities they distinguish, rather than separating antecedent units.

Organisations do not observe environments; they are visible within the same field they observe. As Merleau-Ponty puts it, “he who sees cannot possess the visible unless he is possessed by it, unless he is of it” (1968, p. 134). The reversibility of seeing and being-seen makes the observer-stance the VSM presupposes unavailable in principle.

What allows organisational visibility is itself invisible, a position Merleau-Ponty’s formulation captures: “the invisible is not another visible. It is the invisible of this world, that which inhabits this world, sustains it, and renders it visible” (1968, p. 152). The constitutive invisibility is the chiasmic-organisation parallel of what the present paper develops elsewhere as the constitutive outside: the conditions that must remain non-thematic for the visible aspects of organisational life to operate at all.

§10. Barad’s Agential Realism

§10.1 Intra-Action and Interaction

Karen Barad (2007) extends the chiasmic insight through agential realism. Where interaction presupposes pre-existing agents that subsequently encounter each other, Barad’s neologism *intra-action* describes the mutual constitution of agencies through their entanglement. Barad’s foundational formulation (2007, p. 33): “the neologism ‘intra-action’ signifies the mutual constitution of entangled agencies. That is, in contrast to the usual ‘interaction,’ which assumes that there are separate individual agencies that precede their interaction, the notion of intra-action recognizes that distinct agencies do not precede, but rather emerge through, their intra-action.”

For organisations, the implications are direct: workers and managers do not pre-exist their organisational relations. The cut that separates them, the agential cut in Barad’s terminology, creates the entities it separates rather than recognising antecedent ones. Customer and provider, employee and employer, internal and external: each pair emerges through the apparatus that distinguishes them, with the apparatus itself part of what the distinction produces.

§10.2 Cutting Together-Apart

Barad’s formulation captures the dual operation: “cutting together-apart (one move) is how differentiating-entangling is enacted” (2007, p. 176). The move differentiates (creates the cut between agencies) and entangles (binds the agencies in mutual specification) in a single operation. The organisational examples are immediate: organisational charts cut roles together-apart, the chart simultaneously separating the roles (employee A is not in role X) and binds them in a single specification (the entire role system is defined together). Strategic plans cut markets together-apart: defining the target market separates it from non-target markets, and the operation that defines it is also the operation that produces the entire market typology in which the target sits.

§10.3 Apparatus and Phenomenon

Barad’s central claim about measurement (2007, p. 142): “apparatuses are not passive observing instruments. On the contrary, they are productive of (and part of) phenomena.” Organisational measurement apparatus produces what it measures, becomes part of the organisational phenomenon, and shapes the future operation of the organisation that deploys it. KPIs do not measure performance; they construct what counts as performance. Dashboards do not display organisational state; they constitute the organisational state they display. The recognition extends Maturana’s “everything said is said by an observer” into the domain of measurement: every measurement is made by an apparatus that is part of what it measures. As Barad puts it, “apparatuses are boundary-making practices” (2007, p. 206).

§11. Conway's Law and Its Reversal: Empirical Validation

§11.1 The Chiasmic Nature of System Architecture

Melvin Conway's (1968) original observation, that "organizations which design systems are constrained to produce designs which are copies of the communication structures of these organizations," reveals a chiasmic dynamic in software practice. The empirical literature (MacCormack et al., 2012; Colfer and Baldwin, 2016) demonstrates that the reverse direction is equally operative: systems shape the organisations that maintain them. The two-way relation, what may be called Conway's Chiasm, runs in both directions simultaneously: organisational structure constitutes system architecture, and system architecture constitutes organisational structure, with neither term temporally prior to the other.

§11.2 Evidence from Software Development

Three empirical studies anchor the case. Mockus, Fielding, and Herbsleb (2002) studied Linux kernel development and the Apache project, finding that code architecture shapes developer communities (architectural modularity drives community formation around components) and that developer communities reshape code architecture (community structure drives architectural evolution). The two directions co-evolve, with neither preceding the other in any temporal sequence the study could identify.

Howison and Crowston (2014) studied open source projects more broadly, finding the same pattern of continuous mutual specification: project architectural decisions and contributor relationships shape each other through ongoing recursive operation. Cataldo, Wagstrom, Herbsleb, and Carley (2006) studied commercial software development at Bell Labs and confirmed that technical dependencies create organisational dependencies and that organisational dependencies create technical dependencies, with the intertwining inseparable in the empirical record.

The strongest single empirical confirmation comes from MacCormack, Baldwin, and Rusnak (2012), discussed at length in §33.3 of this paper. Their test of the mirroring hypothesis uses matched pairs of software products with identical functional requirements developed by organisationally tightly-coupled commercial firms versus organisationally loosely-coupled open-source communities. In every matched pair examined, the product developed by the loosely-coupled organisation exhibited significantly higher architectural modularity than the product from the tightly-coupled organisation. The effect size reached a factor of six in propagation cost between component pairs. The mirroring hypothesis was supported without exception across the matched pairs.

§11.3 Why Planned Boundaries Do Not Hold

The empirical literature on organisational boundary work converges on the same chiasmic conclusion. Mergers collapse at rates of 50-70% according to Cartwright and Schoenberg's review (2006), with the collapses typically attributed to "cultural incompatibility." The chiasmic reading interprets cultural incompatibility as incompatibility of intertwined chiasmic structures: what cannot merge is what is internally intertwined differently in each pre-merger organisation, and the post-merger organisation cannot impose a coherent intertwining on the combined units through any planning operation external to the units' own ongoing self-constitution.

Studies of digital transformation (Vial, 2019) document the same pattern across a different domain. Boundaries between technology and organisation prove resistant to the planned reconfiguration that transformation programmes assume. The boundaries are zones of transformation rather than lines of separation, and the transformation programmes themselves become part of what they were intended to operate on. Studies of agile transformations (Dikert, Paasivaara, and Lassenius, 2016) document the same pattern within a more bounded scope: top-down agile transformations collapse at rates above 80% in large enterprises, with the collapses attributable to the assumption that an agile organisation can be designed and then implemented, rather than the chiasmic recognition that the agile organisation emerges through the practices that constitute it.

§12. Gibson, Varela, and Enactive Cognition

§12.1 Affordances as Relational

James Gibson's (1979, p. 127) concept of affordances anchors the chiasmic reading of organisational capabilities: "the affordances of the environment are what it offers the animal, what it provides or furnishes, either for good or ill." Affordances exist only in the organism-environment relation. They are real but properties of neither party alone, emerging through interaction history rather than residing in either organism or environment in isolation. The organisational analogue is direct: markets afford opportunities only relative to organisational capabilities, capabilities afford action only relative to environmental conditions, and neither markets nor capabilities exist independently of the chiasmic relation through which both come into focus.

§12.2 Enaction and Sense-Making

Varela, Thompson, and Rosch (1991, p. 173) develop the enactive account: "cognition is not the representation of a pregiven world by a pregiven mind but is rather the enactment of a world and a mind on the basis of a history of the variety of actions that a being in the world performs." Organisations enact opportunities through their actions, co-create environments through their participation, and perform capabilities rather than possessing them as static properties. As Varela and colleagues argue, "the organism and environment are mutually specified" (1991, p. 205).

Karl Weick's organisational sensemaking research aligns with the enactive account. His central formulation about retrospective sensemaking has already appeared in §2.3: "people often don't know what they think until they see what they say" (Weick, 1995, p. 30). Organisations create environments through their categorical operations, enact constraints through their procedural decisions, and produce opportunities through their attentional patterns. Weick's closing formulation captures the position: "enacted environments are outputs of organizing, not inputs to it" (1995, p. 34).

§12.3 The Chiasmic Organisation

The four concepts: reversibility (Merleau-Ponty), intra-action (Barad), enaction (Varela, Thompson, and Rosch), and flesh (Merleau-Ponty), together specify the chiasmic-organisation ontology. The argument is ontological rather than metaphorical: organisations are patterns within fields they continuously co-create, exhibiting reversibility of observer and observed, intra-action through which agencies emerge, enactment through which environments come into being, and the flesh-medium through which the recursive relations are sustained.

§13. The Mathematical Requirement

Set theory cannot represent chiasmic phenomena. The formalism requires that elements be either in a set or not, that boundaries be sharp (characteristic functions $\chi: X \rightarrow \{0, 1\}$), and that sets be distinct or identical without interpenetration. Organisational reality systematically presents elements simultaneously inside and outside their containing categories, boundaries operating as zones of transformation rather than separation, interpenetration without coincidence between organisational subsystems, and reversibility without identity between observer and observed. The mismatch is direct, and no refinement of set-theoretic formalism resolves it: set theory's foundational commitments are the commitments the chiasmic ontology has displaced.

Systems theory in the von Bertalanffy (1968) and Checkland (1981) tradition assumes definable boundaries, input-output relations, feedback loops connecting separate entities, and observers who model systems objectively. Each assumption has been undermined in the preceding sections. The chiasmic ontology requires mathematics capable of articulating boundaries as zones rather than lines, organisations without inputs or outputs because of operational closure, simultaneous feedforward and feedback, and the inseparability of observer and observed.

Luhmann (1995, p. 37) approaches the requirement with his formulation that "the system is the difference between system and environment." The formulation still treats the difference as pre-existing the operation that

produces it, where the chiasmic ontology and Barad's intra-action analysis treat the difference as itself enacted in operation. The mathematical apparatus required to articulate the chiasmic features must accommodate this enactment.

Four formal apparatuses point in the right direction. Category theory (Mac Lane, 1971) articulates relations between relations, with morphisms primary rather than objects. Topos theory (Johnstone, 1977) provides internal logic and sheaf-theoretic local-to-global structure. Homotopy type theory (Univalent Foundations Program, 2013) handles paths and equivalences as constitutive rather than derived structure. Derived categories (Verdier, 1996) accommodate transformations with irreducible remainder. The four apparatuses operate predominantly in the structural-analogue register that §25.3 below makes explicit. They articulate the chiasmic features rather than constructively embedding chiasmic organisations into specific mathematical objects, because the chiasmic argument shows constructive embedding from any observer position would itself violate the chiasm it sought to formalise. The proof-register foundation supplied by the horizon-mismatch theorem (§24) establishes the floor of the cascade. The four apparatuses extend the structural argument upward.

Part IV: Mundane Complexity in Everyday Organisational Life

§14. Methodological Note

The narrative that follows is a phenomenological reduction rather than an anecdote. It examines an ordinary organisational moment through the theoretical apparatus developed in Parts II and III. The chiasmic, autopoietic, and recursive phenomena it identifies are the substance of daily organisational life rather than rare exceptions requiring special circumstances.

Every element of the scenario, including Sarah's embodied knowledge, the bug report's resistance to categorisation, the manager's demand for metrics, Tom's emergent debugging strategy, Mark's marginalised support work, instantiates principles derived in Parts II and III. Practice and theory specify each other through the narrative; the chiasmic relation between them is itself an instance of the chiasm the paper has been articulating.

The scene is ordinary because the extraordinary is ordinary. A Tuesday morning at a software company contains the multiplicities the theoretical apparatus identifies.

§15. Tuesday Morning, 9:47 AM: A Bug Report

Sarah, a backend developer at a mid-sized software company, receives a bug report from customer support. The customer's API calls are failing intermittently. The situation seems simple enough on its face: debug the code, find the issue, deploy a fix. The narrative that follows traces what actually happens through the theoretical lenses Parts II and III have developed.

The bug report arrives through Jira. Sarah sees the notification because she is designated as the API team lead, a role that exists because the organisation drew a distinction (Spencer-Brown) between API team and not-API team. She also works on the payment system, which makes her simultaneously inside and outside the API team boundary. When she opens the ticket, she is the API team. When she switches tabs to her payment code, she is not. The boundary exists only through her enactment of it.

§15.1 The Information That Is Not

The bug report contains what Bateson would recognise as news of difference: the API worked before, now it does not. But which differences do not make a difference here? The customer's frustration level is not captured in the ticket. The support agent's previous interaction with this customer is not linked. The fact that three other customers experienced similar issues but did not report them is invisible non-difference. Sarah receives a transformed, coded version of events that preceded it rather than information in any unmediated sense. The map (the bug report) is not the territory (the actual system fault).

Meanwhile the support agent, Mark, has already told the customer “our engineering team will fix this quickly,” creating what will become a double bind. Sarah must fix it quickly to satisfy the customer promise and also fix it thoroughly to prevent recurrence. She cannot do both. She also cannot comment on the contradiction because doing so would mean throwing Mark under the bus. The four conditions of Bateson’s double bind (1972, p. 208) are all present: primary injunction (fix quickly), secondary injunction (fix thoroughly), inability to escape the field, and prohibition on commenting on the contradiction.

§15.2 The Autopoietic Response

Sarah’s response is not determined by the bug report itself; it is determined by her organisational structure (Maturana and Varela). The bug report is a perturbation rather than an instruction. Her internal structure, shaped by years of similar tickets, previous fixes, and learned patterns, determines her response. She immediately thinks “probably a rate limiting issue,” not because the ticket says so, but because her structural coupling with the system over time has produced this pattern recognition.

She sends a Slack message to Tom on the infrastructure team: “are we seeing elevated API traffic?” Tom is in a meeting but sees the notification. He is now simultaneously in the meeting (bodily present) and in the Slack conversation (digitally present). His manager is talking about fourth-quarter planning while he checks Grafana dashboards. The organisation exists in multiple overlapping presents that cannot be synchronised into a single temporal frame.

§15.3 The Chiasmic Moment

Tom discovers that API traffic is up three-fold because a large customer has started using a new integration. The chiasm becomes visible at exactly this point: the internal infrastructure problem is caused by external customer behaviour, but the customer’s behaviour was enabled by the company’s own business development team who promised unlimited API calls. Inside becomes outside; the company created the customer behaviour that now appears as external pressure.

Sarah realises she needs to implement rate limiting, but this will affect the customer causing the problem. She must communicate with sales, who sold the unlimited plan. Sales reports to a different vice president who has different OKRs. Sarah’s world (engineering efficiency) and their world (customer satisfaction) operate with different logics, different truth values. What is “true” in engineering (we need rate limits) is “false” in sales (we promised unlimited). The two organisational subsystems operate with incompatible codes, in exactly the form Luhmann’s social-systems theory predicts.

§15.4 The Observer Creating Reality

The vice president of engineering, David, asks for a status update. In preparing her response, Sarah creates the reality she is reporting. She writes: “we have identified the root cause as insufficient rate limiting.” But was that really the root cause? Or was it the sales promise? Or the infrastructure capacity? Or the business model? The act of observation, writing the status update, creates “the root cause” from a field of possible causes. Von Foerster’s recognition holds with full force at the keyboard: Sarah constructs the reality she will subsequently report as discovered.

David reads the update and thinks “we need better monitoring.” He has now created a different reality, one in which the problem is observability rather than capacity or promises or business models. Each observer creates a different environment through their categorical operations on the same underlying field.

§15.5 Recursive Blind Spots

Sarah cannot see her own pattern matching as construction; she takes it as discovery. David cannot see his framing as construction; he takes “better monitoring” as identifying what is real. Sales cannot see how their promises shape what becomes a problem; the promises arrive in their consciousness as customer needs. Mark in support cannot see how his interaction shapes what gets reported; the reports arrive as just what the customer said.

Each participant operates with blind spots that are constitutive rather than incidental. The blind spots are not failures of attention or competence; they are the consequence of the participant's situated position in the field, which makes some aspects of the field visible by making others invisible. The pattern is Merleau-Ponty's chiasm operating at organisational scale: every visibility is doubled with a counter-invisibility that makes it possible.

§15.6 Learning Levels Colliding

The automated test suite operates at Learning Zero, mechanical responses to predetermined inputs without learning. Sarah operates at Learning One, applying pattern recognition from past tickets to the immediate fix. Tom operates at Learning Two, questioning capacity planning at the framework level: should the infrastructure be designed for these traffic patterns at all? The framework-level question (should the company sell unlimited API calls?) belongs to Learning Three. Nobody asks the Learning Three question. The organisation has the cognitive capacity for it; what it lacks is the institutional standing for the question to gain traction. This is exactly the §3.4 structure: VSM-style organisations operate at Learning Two with occasional moments aspiring to Learning Three, but the question that would constitute Learning Three is unavailable from within the framework that would have to be questioned.

§15.7 Eigenforms Emerging

"The way we handle scaling issues" emerges as a stable form across the operations: implement rate limiting, communicate with affected customers, log the incident in the runbook, postpone the deeper question about business model. The pattern stabilises through recursive operation. Each time the team executes the pattern, the pattern reinforces. After enough iterations, "the way we handle scaling issues" appears as an object in organisational reality, something with properties and characteristics that can be referred to and reasoned about. It is, in von Foerster's terms, an eigenform: the stable product of recursive operation rather than an antecedent feature of the organisation.

§15.8 Re-Entry and Paradox

The system attempting to observe itself creates Spencer-Brown re-entry. The monitoring system that watches API performance is itself making API calls, contributing to the load it measures. The Grafana dashboard showing system health affects system health by consuming resources. The Slack conversation about communication problems is itself a communication problem (too many channels, too many interrupts). Sarah's fix, implementing rate limiting, will be deployed through the CI/CD pipeline that itself makes API calls that could be rate limited by the fix she is implementing. The system enters itself, creating the kind of paradox Spencer-Brown showed requires imaginary values for formal resolution.

§15.9 Structural Coupling Across Time

While all this happens, the organisation and its environment are structurally coupled across multiple timescales. At milliseconds, the API and the customer request cycle. At minutes, Sarah and the debugging session. At hours, the ticket and the customer satisfaction metric. At days, the sprint and the feature delivery. At quarters, the OKR and the market position. At years, the technical debt and the capability erosion.

Each timescale has its own coupling dynamics, and they are not synchronised. A decision made at the quarterly scale (promise unlimited API calls) creates perturbations at the millisecond scale (API overload) that create responses at the hourly scale (debugging) that accumulate at the yearly scale (technical debt). The horizon mismatch developed formally in §24 appears here in concrete form: the regulator (the quarterly OKR-setter) operates at a temporal horizon of quarters; the substrate (the API system under load, the engineering team, the customer base, the technical debt) operates at temporal horizons spanning milliseconds to years. The regulator's measurement frame cannot resolve the substrate's full variety. The theorem's inequality $V_M \leq V_R \leq N(T_R) < V_\Sigma$ is exactly what the team experiences as their fix at one timescale generates problems at others.

§16. The Ordinary Impossibility

By 11:30 AM, Sarah has deployed a fix: rate limiting that excludes the large customer. The fix is not really a fix. It is a temporary stabilisation that creates technical debt (the architecture should handle this load), preserves the double bind (Mark will continue to promise quick fixes), defers the framework-level question (should the business model continue to sell unlimited capacity?), deepens structural coupling (the mobile team will now need to handle the cascade effects from rate limiting), adds to recursive blind spots (the monitoring will now treat rate-limited traffic as success), and operates at Learning One (the deeper question never gets asked).

The Jira ticket is marked resolved. The dashboard shows green. The KPI for mean time to resolution improves. Every theoretical lens applied in this section shows that the situation has become more complex, more entangled, more problematic. The metric improves while the substrate degrades.

§16.1 What KPIs Cannot See

The KPIs show the resolution time of one hour and forty-three minutes, customer impact contained, two systems affected. The actual occurrence involves Sarah's autopoietic pattern matching, three boundary reversals across the engineering-sales-customer field, two double binds requiring her bodily absorption, observer-dependent construction of "the root cause," five recursive blind spots across the participants, the emergence of "how we handle scaling issues" as an eigenform, three re-entry paradoxes in the monitoring and CI/CD systems, structural coupling across six timescales, learning-level mismatches between automated tests and the questions that would constitute Learning Three, and the accumulation of technical debt in derived categories that have no metric. The first list is the surface visible to measurement. The second list is the substrate the measurement cannot resolve.

§16.2 The Mundane Reality

This is not a crisis. This is Tuesday morning. Sarah has three more tickets to review before lunch. Tom has two more meetings. Mark in support has seventeen more customer messages. David has a board deck to prepare in which he will show that mean time to resolution is down 15% this quarter.

Every ticket, every meeting, every customer interaction involves all these phenomena simultaneously. They are not special cases or dramatic events. They are the mundane, everyday reality of organisational life. The complexity is generated by the ordinary operation of organisations as chiasmic, autopoietic, recursively coupled, paradox-generating systems, rather than by unusual circumstances or exceptional events.

§17. The Human Absorption of Impossibility

The resolution time of one hour and forty-three minutes looks good on the dashboard. The ticket is closed. The KPI is green. The question for this section is how this actually happens, and what it costs.

Sarah skipped her morning break to debug the issue. She has been skipping breaks for three weeks because ticket velocity is being tracked. Her neck hurts from hunching over her laptop. She is drinking her fourth coffee of the day. When she implemented the rate-limiting fix, she knew it was a workaround. She knew it would create problems later. The ticket needed to be closed; the metric needed to be green.

The pattern is structurally similar to fighters cutting weight for competition. The double bind is direct: you must be strong enough to fight (requiring muscle mass and hydration) and light enough to qualify (requiring dehydration and muscle loss). Fighters resolve the impossibility by destroying their bodies in the weight-cut period, sitting in saunas until they pass out, wearing plastic suits while running, taking diuretics that damage their kidneys. Some die. Others suffer permanent organ damage. They make weight. The metric is achieved at the cost of the substrate that produces it.

Sarah is making weight in the same structural sense. She must be fast enough to close tickets (requiring quick fixes) and thorough enough to prevent future issues (requiring careful design). She resolves the impossibility by destroying her well-being: skipping meals, working weekends, implementing solutions she knows are incomplete.

The ticket closes. The KPI turns green. Her body absorbs the contradiction.

Tom checked the Grafana dashboards while his manager was presenting fourth-quarter plans. He absorbed the double bind into his body: physically present in the meeting, mentally debugging infrastructure. Tonight he will dream about API endpoints. His partner will ask why he seems distant at dinner. He will not be able to explain that he is simultaneously at home and at work, that the boundary dissolved at 9:47 AM when Sarah's Slack message arrived.

Mark in support has managed customer expectations through careful emotional labour, crafting messages that sound helpful while promising nothing. He has absorbed customer frustration into his body, metabolising their anger into careful phrases. His empathy is being extracted and commodified, measured as customer satisfaction while destroying his capacity to feel. The linguistic labour, the careful crafting of non-promises that sound like promises, exhausts him in ways that no metric captures.

The KPIs work because humans absorb the impossibility. The absorption operates through three modes simultaneously: somatic absorption (Sarah's neck, Tom's sleep, Mark's empathy), social absorption (the boundary dissolution between work and home, the strained relationships at dinner), and biographical absorption (the cumulative erosion of capacities over months and years that no quarterly measurement registers).

§18. The Three Sites of Extraction

When KPIs create impossible demands, three sites absorb the impossibility: the worker, the substrate, and the externality. The three-site framework articulates what is happening across the Tuesday-morning scenario in its underlying structure.

Worker-site absorption. Sarah works sixty-hour weeks to keep the tickets flowing. She makes it work by coding during meetings (partial attention to everything), working evenings (time taken from her life), taking shortcuts she knows will create future problems, building mental models too complex to document, and maintaining relationships that hold the system together despite the system itself. The KPIs look good. Sarah is looking for a new job. She will leave in four months and take irreplaceable knowledge with her. Her replacement will not understand why certain things are done in certain ways. The system will degrade slowly and invisibly, until it collapses without warning.

Substrate absorption. The team makes the KPIs work by marking tickets resolved when they are really worked around, implementing fixes that are technical debt, creating special cases that make the codebase unmaintainable, skipping documentation because velocity is measured and documentation is not, and building fragile solutions that work today and break tomorrow. The KPIs improve quarter over quarter. The codebase becomes increasingly brittle. Every change takes longer. Every fix creates new bugs. This quarter's metrics look good. Next quarter is someone else's problem.

Externality absorption. The organisation makes the KPIs work by pushing complexity elsewhere. Rate limiting pushes load to the mobile team. Quick fixes push problems to future quarters. Local optimisations push failures to interfaces. Departmental successes push costs to other departments. Current solutions push debt to future teams. The dashboards stay green. The debt is invisible. It accumulates in derived categories that have no metric. One day, adding a simple feature becomes impossible and no one can explain why. The organisation has painted itself into a corner one green KPI at a time.

The three sites operate simultaneously rather than in alternation. Worker burnout is co-occurring with quality destruction is co-occurring with externality accumulation, and the metric improvement that drives all three is the same metric improvement.

§19. Tuesday's Hidden Cost

Sarah has handled four tickets today. The metrics show four resolutions, average resolution time of one hour and twelve minutes, customer satisfaction score of 4.3 out of 5. The actual cost includes seventeen micro-decisions that will haunt the codebase in months to come, three double binds absorbed and manifest as shoulder tension, forty-three context switches without achieving deep focus on any of them, a promise to Tom to review his infrastructure proposal (which she will not have time to do), and accumulating mental models that exist only in her head.

Tom has attended three meetings while monitoring seventeen Slack channels. He has agreed to commitments in meetings while debugging issues in Slack. His brain maintains parallel tracks that cannot be synchronised. He will wake at 3 AM thinking about API rate limits.

Mark has managed customer expectations through careful emotional labour. He has absorbed customer frustration into his body, metabolising their anger into careful phrases. His empathy is being extracted and commodified, measured as customer satisfaction while destroying his capacity to feel.

The KPIs show tickets resolved, customer satisfaction trending up, system uptime at 99.9%. The same KPIs do not show three people approaching burnout, technical debt approaching critical mass, knowledge residing in exhausted brains and undocumented mental models, and systems maintained by human sacrifice.

§20. The Logical Impossibility

That the KPIs work does not make them logical, sustainable, or even short-term viable. The KPIs operate by extracting unsustainable human effort, hiding accumulating complexity, pushing problems to invisible futures, destroying quality while measuring quantity, and creating the pressure that ensures the degradation continues.

The arrangement is not logical. It is not sustainable. It is barely short-term viable. The dashboard is green so it must be working. Tomorrow Sarah will quit, or the system will crash, or the technical debt will make further change impossible. Or all three. The KPIs will look good right up to the moment they do not. Then everyone will wonder why no one saw it coming.

What the KPIs cannot see is human exhaustion accumulating, technical debt compounding, knowledge evaporating, quality degrading, and complexity accumulating. The metrics work by not measuring what matters. They succeed by defining success as what can be measured while ignoring everything else. They function by forcing humans to absorb impossibility into their bodies, their relationships, their lives.

Sarah's resolution time of one hour and forty-three minutes is both true and a lie. True: the ticket closed. Lie: everything else about what happened, what it cost, what it means, and what it is building toward.

§21. The Tracking Impossibility

To track what actually happened on Tuesday morning, one would need to map autopoietic responses (impossible: internal states are inaccessible to external measurement), trace boundary reversals (impossible: boundaries are enacted in operation rather than fixed for measurement), identify all double binds (impossible: most are unspeakable within the discourse that generates them), track observer constructions (impossible: requires observing the observation), map blind spots (impossible: one cannot see what one cannot see), detect eigenforms (impossible: they appear as natural reality rather than as recursive products), resolve re-entry paradoxes (impossible: requires Spencer-Brown's imaginary values for formal treatment), synchronise time scales (impossible: different scales have different logics), measure learning levels (impossible: Learning Three is the level at which the measurement framework would have to be questioned), and calculate accumulating remainder (impossible: it exists in derived categories with no metric structure).

The ten impossibilities are the empirical manifestation of the horizon-mismatch theorem from §24 applied to the substrate of Tuesday morning. The substrate's variety exceeds the measurement frame's resolution. Every

element of the morning is invisible in principle to the apparatus that is supposed to manage the morning.

§22. The Mundane Conclusion

By lunch, Sarah is eating a sandwich while code-reviewing a colleague's pull request. The API issue is resolved without being solved. The organisation has adapted without learning. The complexity has increased while appearing managed. The KPIs look good while measuring nothing real.

This is not a malfunction; it is normal operation. Organisations work through this complexity rather than despite it. The impossibility of tracking does not mean organisations cease to function. It means they function through processes that exceed any possible measurement system.

The question is not how to measure organisational reality but how to navigate skilfully through unmeasurable complexity. The answer cannot be metrics, dashboards, or KPIs. It must be something that engages the simultaneous operation of all these phenomena rather than fighting against them, which is the philosophy of cultivation §27 below sets out in detail.

Tomorrow Sarah will discover that the rate-limiting fix causes cascade failures in the mobile app. The mobile team's retry logic, combined with the rate limiting she deployed today, will create an unexpected resonance pattern. The cascade will require coordination across three teams, each operating with different assumptions, different priorities, different worldviews. The CEO will ask why a simple bug fix is affecting mobile performance. No one will be able to explain that there are no simple bug fixes in structurally coupled, autopoietic, chiasmic systems with recursive blindness and emergent eigenforms. They will create another dashboard to track cross-team dependencies. It will not help. The organisation will survive, adapt, continue, through the mundane complexity of organisational life that exceeds any possible formalisation. The organisation will persist because it is unmeasurable, maintaining itself through processes that no mathematics, from arithmetic to ∞ -topoi, can capture.

This is Tuesday. This is Wednesday. This is every day.

Part V: The Mathematical Apparatus

§23. The Constituent Cannot Total the Process

§23.1 What the Argument Is Not

The argument of this part is easy to mistake for a claim it is not making, and the misreading disarms it, so the boundary is drawn first. The argument does not hold that organisations cannot count. A firm can count units shipped, headcount, cash on hand, revenue booked in a quarter, defects per thousand, days from order to delivery. These quantities are real. They are not approximations of something truer that the firm has not reached. When a company records forty million in sales for a quarter, the forty million happened. No part of this paper denies it, and any reading that takes the paper to deny it has mistaken the target.

The economists' complexity argument is also not the argument here, though it sits nearby and is worth separating out. Hayek (1945) held that the market's information is distributed across millions of agents, that no central mind can assemble it, and that the price system economises on the knowledge any single participant needs. The Austrian case is epistemic. It says the calculator is too small for the object: the object is vast, the knowledge is dispersed, and the central planner falls short for want of reach. The complexity economists who followed (Arthur, 1994; the broader Santa Fe programme) sharpened the point with non-linearity and path dependence, but kept its shape. The system is too complex to compute from outside.

The present argument runs underneath that one. The trouble is not that the calculator is too small for the object. The trouble is that there is no outside from which the calculation could be performed, because the

would-be calculator is a constituent moment of the process it proposes to total. This is the lesson the second-order cyberneticists drew and the lesson Beer's own architecture could not absorb without dissolving. The firm does not observe the market and then act in it. The firm is one of the processes through which the market continuously constitutes itself. A retailer setting a price is not reading a market price and responding to it; the price it sets is one of the constituting acts out of which the market price, that statistical artefact, is afterward assembled by an observer who is also inside. The same holds within the firm. Management does not stand outside the organisation observing its state and adjusting inputs. Management is one of the organisation's own processes, and its observations are operations of the organisation upon itself, in exactly the recursive sense §5 and §6 developed.

The wave cannot measure the ocean's displacement including its own contribution to the displacement, because the measurement would have to complete before the wave's own motion completed, and the wave's motion is part of what is being measured. The organisation is in this position with respect to itself and its market. The incalculability is not a size mismatch between a small calculator and a large object. It is the structural impossibility of a constituting process totalling, from within, a whole that includes its own constituting activity as an unfinished term.

The transition from participant to constituent is not a matter of degree, and the distinction can be drawn without appeal to size. An agent participates in a market when it responds to conditions the collective has established; it constitutes the market when its own acts generate the conditions to which others respond. Four marks separate the two. The first is control of interface: the agent determines how others encounter the economic space, what is visible to them, what alternatives are surfaced, so that others meet the market as the agent presents it rather than directly. The second is control of alternatives: the agent's position ensures that competing options are absent, absorbed, or rendered unviable, so that exit is formally available and practically empty because there is nowhere to exit to. The third is insulation from exit: network lock-in, data accumulation, and the absence of interoperability ensure that the agent's offering can degrade without proportionate loss of participation, suspending the discipline that departure would otherwise impose. The fourth is the capacity to set conditions rather than respond to them: the agent does not take prices, it makes them; it does not move within rules, it writes the rules others move within. When these four obtain together, the agent has crossed from participant to constituent.

The crossing shows itself in the sign of the feedback. Below it, perturbations to an agent's position are self-correcting: raise prices and customers leave, let quality slip and a competitor takes the demand, and the agent is returned toward the collective condition. This negative feedback is the discipline the cybernetic tradition and the market tradition alike rely on. Above the crossing, the same perturbations become self-reinforcing: the agent with the largest base draws more participants, the agent processing the most exchange builds the most complete model of the space, the agent mediating the most flow captures the rents that fund further mediation. The feedback that disciplined the participant now entrenches the constituent. The attractor has changed from distributed condition to monopolistic constitution, and the change is a change of kind rather than of magnitude, because the very mechanism that held the participant in place is what now consolidates the constituent's hold. A regulator inside such a system is regulating conditions it is itself generating, and the deeper the consolidation, the less its measurements report on anything external to its own constituting activity. This is the precise sense in which the constituent cannot total the process: at the moment its measurements would matter most, they have become the most fully circular.

§23.2 The Two Things a Metric Cannot Do

Two distinct incapacities follow, and they must be kept apart because they hold for different reasons and a defender of metrics will try to collapse them.

The first incapacity concerns people. A workforce is not a stock of interchangeable units whose output sums. It is a tissue of relationships, accumulated judgement, tacit coordination, and mutual adjustment built up over years, the substance Part IV showed Sarah, Tom, and Mark absorbing into their bodies on an ordinary Tuesday. The relationships are not a quantity with measurement error. They are not a quantity. A person's standing in a

team, the trust that lets one engineer accept another's verdict without re-deriving it, the half-sentence in a corridor that prevents a fortnight of wasted work: these are constitutive of the team's capacity and they have no scalar. One can count the engineer's commits. One cannot count what the engineer is to the team, because what the engineer is to the team is a relation in a continuous process, not a magnitude in a state vector. This is the sense in which people in a complex system are incalculable, and the claim is held without qualification.

The second incapacity concerns trajectory. Even granting every real number the firm books, the metric cannot register what the firm is building or dismantling in the act of booking it. The quarter's profit is a real number. Whether that profit is the yield of a renewable process or the harvest of a substrate now spent is not a further number the firm forgot to compute. It is a question about the firm's position on a trajectory, and the trajectory is constituted by processes operating at horizons the quarter cannot see. A firm can post record efficiency in the same quarter it consumes the last of an engineering substrate twenty years in the building, and the efficiency figure and the consumption are not in tension within the measurement frame, because the frame has no term for the consumption. The number climbs. The substrate is gone. The instrument reads success.

The two incapacities are independent. The first would hold even if trajectories were legible, because relations are not magnitudes. The second would hold even if people were interchangeable units, because the trajectory question is about horizons rather than about persons. A defender of metrics who concedes the first sometimes tries to recover the second by promising better instruments, longer dashboards, leading indicators. The horizon-mismatch theorem of §24 closes that recovery in the proof register: the second incapacity is not a gap in current instruments but a structural feature of any regulator operating at a horizon shorter than its substrate.

§23.3 The Projection: Real Part, and a Component That Is Not a Number

The relationship between what the metric reports and what the firm is can be stated precisely, and the precision matters because the imprecise version invites the dismissal that there is more to measure.

Let the firm's condition at a moment be written as an ordered pair

$$z = (\text{Re}, \text{Im})$$

where the first component collects the firm's scalar magnitudes, revenue, output, headcount, the quantities that do admit a number, and the second component collects the firm's constitutive process: its relations, its position on trajectories unfolding at horizons longer than the measurement window, its ongoing participation in the constitution of its own market. The measurement frame is a projection

$$\pi(z) = \text{Re}.$$

The notation borrows the form of the complex number deliberately, and the borrowing must be qualified at once so that no more is claimed than holds. The complex number's imaginary part is itself a real magnitude, a member of the same field, recoverable from the modulus and argument. The claim here is sharper and stranger than the complex case. The second component is not a scalar of the same kind as the first, waiting on a finer instrument that would read it off the same axis. It is a process, relational and constitutive, and the reason the projection π discards it is not that π has too few channels but that π ranges over a single scalar axis and the second component does not lie on that axis at all. This does not make the second component structureless or unnameable. It can carry rich structure, with structure that has its own invariants, as §27 sets out when it names that structure cohomologically. What it cannot do is sit on the one-dimensional axis π projects onto, or be recovered there by refining π . A projection onto the real axis discards the imaginary part because the imaginary part lies off the axis; in the complex case one rotates the axis and recovers it. Here no rotation brings the constitutive process onto the scalar axis, because the constitutive process is not displaced from the axis by an angle. It is of another order than the scalar. The analogy to the complex number holds in form, the metric reports one component of a two-component condition, and breaks in substance, the discarded component is not a scalar recoverable by re-aligning the scalar frame but a structured process that no choice of scalar frame renders as a scalar.

This is the precise sense in which the metrics measure something imaginary. They do not measure nothing; Re is real and the firm's forty million is in it. They report Re and present it as z . The error is not in the reading of Re , which is accurate, but in the identification of $\pi(z)$ with z , the treatment of the projection as the whole. An organisation steering by $\pi(z)$ steers by a real shadow of itself and takes the shadow for the body. The shadow can grow while the body fails, and nothing in the shadow registers the failure, because the failure lives in the component the projection does not range over.

The cohomological articulation of §27 names the structure of the discarded component for the cases where it admits topological description: the higher cohomology where local data fails to glue into global, the obstruction that has no representation in H^0 . The naming is structural-analogue, in the sense §25.3 makes explicit, and it does the work of characterising the shape of Im without claiming Im is computable. What secures the existence of Im , prior to any characterisation of its shape, is the horizon-mismatch theorem, to which the argument now turns.

§24. The Horizon-Mismatch Theorem

§24.1 The Claim and Why It Carries the Paper

The theorem establishes, in the proof register and without any appeal to topos theory, derived categories, or homotopy type theory, that a regulator operating at a temporal horizon shorter than its substrate cannot capture the substrate's variety within its measurement frame. Its corollary, holding wherever the firm goes on persisting, locates the uncaptured variety in the absorbing terms rather than leaving it to vanish. The theorem is the floor on which the rest of Part V stands. If it holds, the existence of the discarded component Im of §23.3 is secured by proof rather than by phenomenological assertion, and the cohomological apparatus of §§27-28 is freed to characterise the component's shape without bearing the burden of establishing that there is anything to characterise. If the floor gives way, the apparatus above it is decoration. For this reason the theorem is stated and defended at length, and the defence is built to survive a hostile reader who grants nothing that is not forced.

§24.2 Premises

The theorem rests on three premises, each of which is either a structural identity or an elementary fact about finite measurement, and none of which is an empirical conjecture about organisations.

The first premise is Ashby's law of requisite variety (1956, p. 207), in the form Beer adopted for organisational regulation (1985, p. 89): a regulator can absorb disturbance variety only up to the regulatory variety it commands. Variety here is the count of distinguishable states. The law is a structural identity rather than an empirical hypothesis. It says that if a disturbance can present more distinct states than the regulator can distinguish and respond to differentially, then some distinct disturbance states are mapped by the regulator to the same response, and control over the difference between those states is lost by definition. Ashby's law is not in dispute among cyberneticians; it is the founding constraint of the field, and Beer built the VSM on it. The present argument takes it as given precisely because the tradition the argument addresses already takes it as given. A defender of metrics who rejects Ashby's law has abandoned the cybernetic ground on which managerial control theory itself stands.

The second premise is the finite resolution of any measurement frame within a finite window. A measurement frame is whatever apparatus the regulator uses to distinguish states of the substrate: its instruments, reports, dashboards, sampled metrics. Any such frame has finitely many channels, each resolving finitely many discrete levels, sampled at a finite rate over a finite observation window. This is not a contingent limitation of present technology. It is a property of measurement as such: an instrument that made infinitely many infinitely fine distinctions in finite time would transmit unbounded information in bounded time, which Shannon's source-coding theorem (1948) forbids for any channel of finite capacity. The premise is therefore secured by information theory, not assumed.

The third premise is the horizon relation. The regulator operates at a characteristic horizon T_R : the window over which it observes before it must act, the reporting period, the planning cycle, the sprint, the quarter. The substrate, the workforce, the codebase, the customer relationships, the market position, the accumulated engineering, operates at a characteristic horizon T_Σ over which its own variety unfolds, and for the substrates that matter $T_\Sigma \gg T_R$. A codebase accumulates structure over years; a quarter cannot see it. A workforce's tacit coordination builds over the careers of its members; a sprint cannot see it. The horizon relation $T_R \ll T_\Sigma$ is not assumed as universal; it is the condition under which the theorem bites, and §24.5 identifies exactly the organisational substrates that satisfy it.

§24.3 Statement and Proof

Theorem (Horizon-Mismatch). Let Σ be a substrate whose variety V_Σ unfolds over characteristic horizon T_Σ . Let R be a regulator coupled to Σ , operating at horizon T_R , with a measurement frame resolving at most $N(T_R)$ distinguishable states within any window of duration T_R . Let V_M denote the variety the frame actually captures and V_R the regulatory variety R commands. Then

$$V_M \leq V_R \leq N(T_R),$$

and whenever $V_\Sigma > N(T_R)$, which obtains generically when $T_R \ll T_\Sigma$, the frame strictly understates the substrate:

$$V_M \leq V_R \leq N(T_R) < V_\Sigma.$$

The difference $V_\Sigma - V_M$ is the variety the frame leaves unregistered. The corollary below establishes where it goes once the further condition of persistence is added.

Proof. The first inequality, $V_M \leq V_R$, holds because the regulator cannot act differentially on distinctions its frame does not capture; the variety available for regulation is bounded above by the variety measured. The second, $V_R \leq N(T_R)$, is the resolution bound of premise two: the frame distinguishes at most $N(T_R)$ states in its window, so the regulatory variety predicated on that frame cannot exceed $N(T_R)$. To bound $N(T_R)$ explicitly, let the frame have k channels each resolving r levels, sampled with period τ . A single measurement act distinguishes at most r^k states. Over a window T_R the frame takes at most T_R/τ such acts, giving $N(T_R) \leq r^k \cdot (T_R/\tau)$ when states are read independently and $N(T_R) \leq r^{k(T_R/\tau)}$ when the temporal sequence itself carries state. The sampling period τ is the frame's temporal resolution, the finest interval at which it can register a change in substrate-state; a later paper in this series carries this resolution as a separate quantity, the model-cycle, alongside the action-cycle T_R and the substrate's replenishment-cycle, where the same theorem is restated for the political economy of extraction. Either bound is finite, and by Shannon's theorem neither can exceed $2^{\rho T_R}$ where ρ is the channel capacity of the regulator's information pathway. Thus $N(T_R)$ is finite and fixed by the regulator's internal constitution.

Now V_Σ is bounded by no regulator-internal quantity. It is set by the substrate's own configuration space unfolding over T_Σ . When $T_\Sigma \gg T_R$, the substrate passes through state-distinctions that mature over durations longer than the regulator's entire window, and these escape the frame because they are not complete within any single window the frame observes. A codebase's architectural decay, a team's erosion of trust, a customer base's slow disaffection: each is a variety that exists only over T_Σ and has no representative within T_R . Hence $V_\Sigma > N(T_R)$, and the strict inequality follows.

A regulator might try to escape the bound by integrating across many consecutive windows: sixteen quarters of reports, read together, to recover the decay that no single quarter shows. The escape relocates the regulator rather than defeating the bound. A regulator that integrates to a horizon of sixteen quarters has become a regulator operating at that longer horizon, and the same inequality applies to it at T_Σ , where it can no longer act on the quarter-scale disturbance it was built to handle. The dynamical regime changes with the horizon: the fast regulator corrects fast disturbance and cannot see slow decay; the slow regulator sees slow decay and cannot correct fast disturbance. No single regulator occupies both horizons at once, because the horizon is what fixes which disturbances it can resolve and which it can act upon. The integration manoeuvre is taken up again in §24.6, where the contradiction in demanding both horizons of one instrument is drawn out.

The theorem as stated establishes the measurement gap: the frame understates the substrate, and the understated variety $V_\Sigma - V_M$ goes unregistered. Where the unregistered variety goes is a further question the theorem does not settle on its own, and the corollary below settles it under one added condition.

Corollary (Absorption under Persistence). Suppose the coupled system persists: the firm continues to hold its essential variables in range, continues to ship, continues to meet the metric the regulator reports on. Ashby's law, applied now to the whole coupled system rather than to the regulator alone, then forces a consequence. The disturbance variety V_Σ has to be met by matching regulatory variety somewhere, or the essential variables would leave their range and persistence would not obtain. The regulator supplies at most $N(T_R)$ of the required variety. The remainder is supplied by the other elements coupled to the substrate, because there is nowhere else in the coupled system for it to come from: the people operating the substrate, the substrate's own material reserve, or systems downstream to which disturbance is passed. The displacement follows as the contrapositive of Ashby's law: if the firm persists and the regulator is short of requisite variety, some other element made up the difference. Read across many such firms and many such periods, the displacement hardens into a constraint, the variety a formal system cannot hold is conserved against the demand for coherence and reappears wherever the coupled system can still absorb it. The constraint, and the full accounting of where the absorbed variety lands and what its absorption costs the bodies and substrates that take it on, is the work of a later paper in this series. Here the corollary establishes only that the variety must land somewhere, and the cannibalisation case of §24.5 reads which element takes it on and at what price. ■

§24.4 What Absorption under Persistence Means

The corollary is the part of the argument that converts an epistemic result into a material one, and it is where the force gathers, so it is drawn out.

The unregistered variety has to be met somewhere once the firm goes on persisting. This is not rhetoric; it is Ashby's law read as a contrapositive over the whole coupled system rather than over the regulator alone. If the substrate generates more distinguishable disturbance than the regulator can meet, and the firm holds together rather than coming apart, then by the same law that grounds the regulator's own operation the difference is met by the other terms coupled to the substrate. There are three such terms, and the metabolic cost their absorption carries, traced through the bodies and substrates that bear it, is followed in detail by a later paper in this series; the present argument establishes the three sites and leaves the accounting of their cost to that work.

The first absorbing term is the people. When the measurement frame cannot represent the variety of what holds the work together, that variety is met by the workers as unmeasured labour: the evenings, the held-together relationships, the judgement exercised in the gap the metric cannot see. Sarah making weight is the corollary read on a body. The metric reads her resolution time; the variety her resolution time cannot represent is met by her skipped breaks and her decision to leave in four months.

The second absorbing term is the substrate itself, as accumulating degradation. When the variety is not met by people it is met by the thing, as wear that does not appear in the window in which it is generated. The codebase takes on the debt. The architecture accretes the special cases. The degradation is real, and it accrues from the moment it is generated, but it matures over T_Σ and so stays out of view across every T_R until the horizon at which it surfaces all at once, as the sudden incapacity Part IV described, the day a simple feature becomes impossible and no one can say why.

The third absorbing term is the externality, the variety discharged to coupled systems outside the regulator's boundary: the other team that inherits the rate-limited load, the future quarter that inherits the deferred problem, the commons that inherits the discharged cost. The regulator's boundary is drawn so that these lie outside it, which is precisely why its frame does not register them and its metric can improve as it discharges them.

The corollary has a consequence that the cannibalisation case of §24.5 turns on. Because the unmet variety is met by the absorbing terms rather than vanishing, an improvement in the measured component can be produced *by* moving variety onto those terms. Efficiency, read as a rising ratio of measured output to measured input, can be manufactured by taking variety off the measured input and placing it on the unmeasured absorbers. The metric

does not simply miss the cost. The metric improves by incurring the cost, because the cost falls in the component the metric does not range over.

§24.5 Cannibalisation Read as Efficiency

The theorem now meets the case you raised, and the case is not an illustration bolted on; it is what the corollary says, read at the horizon of a firm consuming its own substrate.

A firm may carry an engineering substrate built over twenty years: a codebase, yes, but more to the point a workforce whose tacit coordination, accumulated system knowledge, and mutual trust constitute the firm's actual capacity to build. The substrate's variety unfolds over T_Σ of years to decades. The regulator, the quarterly operating review, the efficiency programme, the headcount ratio, operates at T_R of a quarter. By the theorem, $V_\Sigma > N(T_R)$: the substrate carries variety the quarter cannot resolve.

The firm undertakes an efficiency drive, reducing the slack, the apparently unproductive time in which the substrate was in fact maintained: the senior engineers' availability to answer questions, the documentation time, the redundancy of knowledge across people, the headroom that absorbed shocks. Measured productivity rises: the same output, fewer measured inputs, the ratio improves quarter on quarter. The frame reports efficiency, and the report is accurate *as a reading of* Re , since output over headcount has indeed risen.

What the frame cannot report is that the slack it removed was the mechanism by which the substrate maintained itself, and that its removal has begun to consume twenty years of accumulated coordination as current output. The variety formerly met by the slack is, by the corollary, now met by the two remaining terms: the people, who do the same work with the maintenance margin gone, which the firm also reads as efficiency because their increased load appears only as steady output against reduced input; and the substrate, which begins to degrade as the maintenance that the slack performed ceases, the degradation maturing silently over T_Σ and invisible across every quarterly T_R in which it accrues.

The firm is on a slope toward incapacity and the instrument reads the slope as ascent. Careful measurement does not avert this; the structure of the theorem produces it. The efficiency gain and the substrate consumption are the same event viewed in the two components of §23.3: in Re , output per head rising; in the discarded component, the trust and knowledge and maintenance-margin that constituted the firm's capacity being burned as fuel. The two are not in tension within the frame because the frame projects onto Re , where only the gain appears. The cannibalisation is legible only at T_Σ , by which horizon the substrate that would have been needed to recover has already been spent, and the workers who carried the accumulated knowledge have already, like Sarah, begun to leave, taking the unmeasured component out of the door with them.

A firm in this position does not discover its error by measuring better. Every measurement it can take is a reading of Re within T_R , and within that frame the error does not exist. The error is the identification of $\pi(z)$ with z , sustained over the horizon across which the discarded component matures into incapacity. By the time the incapacity surfaces into the measured component, as missed deliveries, as the simple feature that has become impossible, as the exodus of the people who knew why things were done as they were, the substrate that would have been required to respond is the substrate that was consumed to produce the efficiency that the response would now have to reverse. The slope is not observed until it is the floor.

§24.6 Why Better Instruments Cannot Close the Gap

A defender concedes the cannibalisation case and proposes the obvious remedy: measure the substrate. Add leading indicators, track technical debt, survey engagement, instrument the trust. The remedy does not work, and the theorem says why, so the breakdown is structural rather than a matter of insufficient diligence.

Every proposed instrument is a measurement frame, and every measurement frame is subject to the theorem. An instrument that tracks technical debt operates at some horizon $T_{R'}$ and resolves some $N(T_{R'})$ states; if the debt it must track matures over $T_\Sigma \gg T_{R'}$, the new instrument stands in the same relation to the substrate as the old one and understates it by the same structure. One can make $T_{R'}$ longer, but a regulator whose horizon approaches T_Σ

has, by the third premise, ceased to be a regulator operating faster than its substrate; it now acts on year-scale information and cannot regulate quarter-scale disturbance, which is the opposite breakdown. The horizons cannot both be short enough to regulate and long enough to see the substrate, because they are the same horizon. This is the precise content of the claim that the incalculability is structural and not technological: there is no instrument because the demand on the instrument is internally contradictory.

And beneath the horizon argument lies the deeper one of §23.1. Even an instrument operating at T_{Σ} would still be an instrument deployed by a constituent of the process, and the component it is asked to capture, the relational and constitutive process of §23.2, does not lie on the scalar axis any frame projects onto. The trust between two engineers is not a low-frequency signal awaiting a slow enough sensor. It is not a scalar signal at all. The horizon-mismatch theorem shows the gap cannot be closed from within the measured component even for the variety that is in principle scalar; the projection argument of §23.3 shows that the residual does not lie on the scalar axis in the first place. The two results bound the firm from both sides. What is scalar cannot be captured at the horizon that matters; what matters at that horizon is not scalar.

§25. Strategic Deployment of Mathematical Sophistication

§25.1 The Vulnerability the Apparatus Addresses

The chiasmic ontology developed in Part III, grounded phenomenologically and demonstrated empirically through the Tuesday-morning narrative of Part IV, faces a vulnerability of position. It appears insufficiently rigorous compared to the metric-driven, algorithmic approaches that dominate contemporary management discourse. When executives demand KPIs, dashboards, and metrics, phenomenological arguments about reversibility and intertwining may seem insubstantial in the institutional register where decisions are made.

The vulnerability is more than rhetorical. Power (1997, p. 97) documents in *The Audit Society* that organisations increasingly require quantification and measurement as a condition of legitimacy. Meyer and Rowan (1977, p. 340) identify the structural pressure as institutional isomorphism: organisations incorporate the practices and procedures defined by prevailing rationalised concepts of organisational work. An argument that abandons the institutional register surrenders the field; an argument that engages the register must do so with apparatus the register recognises.

The chiasmic model requires protection from reductive demands. The protection comes not through better arguments alone but through mathematical demonstration of the structural limits of formalisation itself. The paper deploys advanced mathematical apparatus for diagnostic rather than representational purposes: to articulate the structural features any adequate formalisation would need to capture, and to establish in the proof register that the simplest formalisations underlying KPIs and dashboards cannot meet the conditions Ashby's law imposes.

§25.2 Mathematics as Diagnostic Apparatus

The strategy uses mathematics diagnostically rather than representationally. Representational mathematics attempts to model and capture the phenomenon. Diagnostic mathematics reveals where and why such capture breaks down. A physician uses diagnostic tools to understand a patient's limits without becoming the patient, and the present argument uses mathematical formalism to mark its own boundaries without claiming to formalise the chiasmic phenomena it discusses.

The approach has direct precedent in mathematical logic. Gödel (1931) used arithmetic to prove arithmetic's incompleteness. Turing (1936) used computation to prove computation's limits. Arrow (1951) used preference logic to prove preference aggregation's impossibility. The present argument follows the same tradition: using formal apparatus to establish the structural limits of formal apparatus when applied to organisational phenomena.

The strategy constitutes an immanent critique. The argument uses the opposition's preferred language (mathematical rigour) to demonstrate the inherent limits of that language when applied to organisational reality.

The argument is not an external critique from a softer discipline that management science can dismiss as lacking rigour. The argument is mathematics demonstrating, with mathematical precision, that formalisation pursued with rigour reveals its own structural limits. The master's tools cannot dismantle the master's house, in Audre Lorde's (1979) formulation, but they can prove the house was never built on the foundation its architects believed it rested on.

MacKenzie (2006, p. 89) captures the feature the strategy exploits: "mathematical models gain authority not from accuracy but from sophistication." The argument leverages the authority strategically, using mathematical rigour to protect phenomenological insight from the dismissive moves the institutional register makes against work in softer disciplines.

§25.3 Two Registers

The mathematical apparatus this paper deploys operates in two distinct registers. The first register supplies proof. The second register supplies structural analogue. Both registers do work, but the work each does is different, and a reviewer or consultant encountering the apparatus deserves to know which is which.

In the proof register, mathematical claims are demonstrated by deduction from premises that hold for any system meeting the relevant conditions. The horizon-mismatch theorem of §24 above is the paradigm case. Ashby's law and finite-cardinality arguments yield the variety-absorption inequality without requiring any commitment to what specifically the substrate is. The Novikov-Boone result on undecidability of the word problem in finitely presented groups is another proof-register result, applied in §24 below.

In the structural-analogue register, mathematical apparatus is deployed as articulation. Sheaf cohomology articulates the phenomenon of local solutions that do not extend to a global one. Derived categories articulate the phenomenon of transformation producing irreducible remainder. Higher inductive types articulate the phenomenon of identity emerging through paths rather than through pre-given containers. In each case the mathematics names what the phenomenology has identified. The naming is rigorous in the formal sense, with the mathematical object's properties matching the structural features the phenomenon exhibits. It is not rigorous in the sense that the mathematical object has been constructed from the phenomenon through a justified embedding.

The distinction matters for three reasons. First, the cascade argument depends on the proof register at its foundation. If the floor is hand-waved, the cascade collapses. The horizon-mismatch theorem secures the floor; from the floor the cascade extends upward through structural-analogue articulations. Second, the distinction matters because reviewers and consultants with mathematical training will probe the apparatus. A paper that claims more than it can defend invites dismissal of even the parts it can defend. Better to mark the structural analogue explicitly and let it do the work it can than to dress it as proof and have the dressing torn off. Third, the distinction matters because the paper's commitment is to mathematics as reasoning rather than as modelling. Reasoning includes both proof and structural articulation. Modelling is precisely what the paper rejects. The structural-analogue register does not propose that organisations *are* topoi or *are* chain complexes; it proposes that organisations exhibit features that topos-theoretic and homological language can name with precision.

§26. The Word Problem and Algorithmic Undecidability

The second proof-register result applied in the cascade is the algorithmic-undecidability result of Novikov (1955) and Boone (1958).

Theorem (Novikov-Boone). There exist finitely presented groups whose word problems are algorithmically undecidable. That is, there exist groups G with finite presentation $\langle S | R \rangle$ such that no algorithm decides, for arbitrary words w_1, w_2 in S , whether $w_1 = w_2$ in G .

The result is standard combinatorial group theory from the mid-twentieth century. Its application to organisational categories proceeds in two steps. The first step is a proof-register result. Consider an organisational category $\mathcal{C}$ whose objects are organisational states and whose morphisms are the available transitions between

states. The nerve $N(C)$ is the simplicial set whose n -simplices are sequences of n composable morphisms. The geometric realisation $|N(C)|$ is a topological space, and its fundamental group $\pi_1(|N(C)|)$ admits a presentation derived from the morphism structure. When the morphism structure of C is rich enough to encode the generators and relations of a Novikov-Boone group, $\pi_1(|N(C)|)$ contains a subgroup with unsolvable word problem, and algorithmic equivalence-checking between sequences of morphisms in C is undecidable. This step follows directly from the embedding theorems of combinatorial group theory.

The second step is a structural-analogue claim. Organisations of even moderate complexity meet the conditions for such encoding. The justification is not a constructive proof of the embedding for any specific organisation. The justification is structural: organisational categories combine the cardinality (large state space), the composition structure (sequential operations), and the relational density (non-commuting transformations) that the standard embedding theorems require. The structural-analogue claim is that the conditions are met by organisations as a class. Constructive proof for any particular organisation would require fixing the category C explicitly, which the chiasmic argument of Part III shows is itself not possible without observer-imposition.

The consequence is that strategic-planning procedures requiring algorithmic equivalence-checking, including procedures that ask whether two sequences of organisational operations lead to identical states, are undecidable in the proof-register sense for any sufficiently rich organisational category, and undecidable in the structural-analogue sense for organisations as encountered.

The cascade does not require this result to operate. The horizon-mismatch theorem of §24 secures the foundation. The Novikov-Boone result articulates an additional dimension of breakdown: even if the horizon mismatch were somehow overcome, algorithmic procedures over the resulting structure would still break down by undecidability. The two modes of breakdown are independent.

§27. Cohomological Obstruction Theory: A Structural Articulation

The apparatus that follows operates in the structural-analogue register. Sheaf cohomology is deployed to articulate, with the formal precision of standard algebraic geometry (Grothendieck, 1957; Hartshorne, 1977), the phenomenon that the chiasmic argument has identified: organisational intelligence resides in the obstruction to uniform extension of local solutions across the organisation's relational field. The intuitive face of this apparatus was given in 3.2, where the impossibility of flattening a curved surface onto a plane without distortion stood as the geometric form of the map-territory relation. The patching obstruction below is that impossibility stated exactly: where the flat map cannot represent the globe without relocating distortion, the local sections cannot be glued into a global section without an obstruction that the higher cohomology measures.

Principle (Intelligence as Obstruction). Organisational intelligence appears where uniform solutions do not extend. In the language of sheaf cohomology, the higher cohomology groups $H^n(F)$ of a sheaf F over a site measure exactly such obstruction to extension. The structural analogue is direct: where H^0 captures globally-extensible sections (uniformly applicable solutions), H^n for $n \geq 1$ captures obstructions to global extension. The principle states that organisational intelligence corresponds to the higher-cohomology content rather than the H^0 content.

The articulation is structural-analogue rather than constructive embedding. The correspondence between organisational structure and the sheaf-theoretic object is not a constructive embedding because the chiasmic argument shows constructing the embedding would require committing to a specific sheaf on a specific site, a commitment that cannot be made from any observer position. What the principle captures is rigorous in the structural sense: the formal features of sheaf cohomology (sequence exactness, the local-to-global spectral sequence, the connecting homomorphisms relating H^n to H^{n-1}) exhibit the structural features the phenomenology identifies (irreducible local-global gap, propagation of obstruction through the relational field, dependence of higher obstruction on lower).

The implication for organisational practice is direct. Measurement instruments designed to capture H^0 content (uniformly applicable metrics, standardised KPIs, dashboards with fixed dimensional resolution) capture

precisely the content that is not the intelligence. They project onto the trivial cohomology and discard the higher cohomology where the intelligence resides. When management consultants propose breaking down silos or imposing departmental KPIs, they propose operations that destroy higher-cohomology content while preserving H^0 content. The intelligence is then no longer present in any measurable register.

The Leray spectral sequence provides further structural articulation. The standard result $E_2^{p,q} = H^p(X, \mathcal{H}^q(F)) \Rightarrow H^{p+q}(I^\bullet)$ articulates the entanglement of cohomological levels: organisational intelligence cannot be computed layer by layer because the differentials d_r entangle the levels essentially. McCleary (2001) provides the standard reference. The structural-analogue claim is that organisational reform programmes which proceed level by level (first restructure leadership, then departments, then teams) do not register the entanglement that the spectral-sequence structure articulates.

§28. Derived Categories and the Necessity of Remainder

The apparatus of derived categories (Verdier, 1996; Gelfand and Manin, 2003) is deployed here in the structural-analogue register to articulate the phenomenon of transformation producing irreducible remainder.

Principle (Necessary Remainder). Every organisational transformation produces a non-trivial cone: a residual structure that cannot be transferred into either the prior state or the subsequent state without loss. In the language of derived categories, the distinguished triangle

$$A^\bullet \rightarrow B^\bullet \rightarrow C^\bullet \rightarrow A^\bullet[1]$$

captures the structure: $C^\bullet$ is the cone of the morphism $A^\bullet \rightarrow B^\bullet$, and it is non-trivial whenever the morphism is not an isomorphism. The principle states that organisational transformations correspond to non-isomorphism morphisms, and the cone $C^\bullet$ corresponds to the irreducible remainder.

The articulation is structural-analogue. The cone construction is well-defined, the non-triviality condition has precise meaning, and the propagation of remainder through composed transformations follows the formal rules of the derived category. The analogue is not constructive in the sense of producing $A^\bullet$, $B^\bullet$, $C^\bullet$ as complexes for any specific organisation. The chiasmic argument of Part III shows why such construction is not available from any observer position.

The phenomenological content the principle articulates is direct. Transformation programmes that promise “seamless transitions,” “frictionless adoption,” or “lossless migration” promise the impossible. The derived-category articulation does not prove this directly; it names with formal precision the structural feature that the phenomenology identifies. The Fourier-Mukai apparatus (Huybrechts, 2006) extends the same articulation to knowledge transfer between organisations: every transfer of practice involves transformation, and transformation produces non-trivial cone, and the cone is precisely the part of the practice that does not transfer.

§29. Higher Inductive Types and Chiasmic Boundaries

The apparatus of homotopy type theory (Univalent Foundations Program, 2013) is deployed here in the structural-analogue register to articulate the chiasmic identity that Merleau-Ponty’s late ontology identifies: the touching-touched, inside-outside, observer-observed structure of §9.

The articulation proceeds by treating organisational boundaries as higher inductive types. A higher inductive type is a type defined not only by constructors that build elements (as in ordinary inductive types) but also by constructors that build paths between elements. The constructor

$$\text{chiasm} : \prod (u : \text{Unit})(e : \text{Environment}), \text{inside}(u) = \text{outside}(e)$$

posits, as a primitive of the type, that inside and outside are connected by a path. The path is not derived from prior identity of the two; it is constitutive of the type itself.

The structural analogue captures the chiasmic feature with precision. Where ordinary set-theoretic and category-theoretic formalisms treat boundaries as separations between pre-existing entities, the HoTT formalism

treats them as paths constitutive of the entities they connect. The Univalence Axiom (Voevodsky, 2006) extends the articulation: identity and equivalence are themselves equivalent, so organisations that operate equivalently are operationally identical. Claims about “culture” or “values” that have no operational manifestation correspond to types with no path-constructors and are therefore empty in the formal sense.

The articulation is structural-analogue rather than constructive embedding. HoTT has precise rules for forming higher inductive types, the path-constructor formalism is well-developed, and the Univalence Axiom has its standard consequences. The articulation does not construct the type `OrgBoundary` for any specific organisation, because the chiasmic argument shows that observer-imposed construction would itself violate the chiasm it aims to formalise. The articulation captures the structural feature without claiming to instantiate it.

§30. The Cascade Footing

With the proof-register foundation established by the horizon-mismatch theorem of §24 and the Novikov-Boone result of §26, and with the structural-analogue articulations of §§27-29 explicitly marked as such, the cascade argument proceeds on solid ground. The cascade does not claim that ∞ -topoi have been proved insufficient for organisations in the strict sense of constructive impossibility. The cascade claims four things, distinguished by register.

In the proof register: for any regulator operating at horizon shorter than its substrate, the substrate’s variety exceeds the regulator’s measurement frame (the horizon-mismatch theorem of §24); and, under persistence, the unregistered variety is met by the absorbing terms (the corollary of §24). In combinatorial group theory: the word problem is undecidable for finitely presented groups of sufficient richness (Novikov-Boone, §26), a result proved about groups, not about organisations.

In the structural-analogue register: the claim that organisational state-categories encode such groups, and so inherit the undecidability, is an application argued structurally rather than by constructive embedding (§26); and organisational intelligence exhibits the formal features that sheaf cohomology, derived categories, and homotopy type theory articulate, namely local-global obstruction (§27), transformation remainder (§28), and path-constitutive identity (§29).

The consequence: mathematical formalisations operating below the level of §§27-29’s articulations cannot capture the features the articulations name. Elementary formalisms (arithmetic in KPIs, basic real analysis in OKRs, linear algebra in balanced scorecards) operate well below this level. The cascade therefore concludes that the inadequacy of elementary formalisms is structural and not a contingent shortfall of present technique, for the reasons the proof-register results establish and the structural-analogue articulations clarify.

The argument is one theorem about variety-absorption with its persistence corollary, one proof-register theorem about algorithmic undecidability with a structural-analogue application to organisations, and three structural-analogue articulations of features that any mathematical formalisation would need to capture. The cascade rests on what is proved at its base and on what is articulated at its higher levels, with the join between them marked at each step. The consultant or reviewer who attempts to attack the apparatus must engage either Ashby’s law (structurally the regulatory premise of cybernetics itself), the Novikov-Boone embedding theorems (standard mid-twentieth-century combinatorial group theory), or the structural-analogue articulations (explicitly marked and not overclaimed). The attack surface is narrow.

§31. Creating Protected Space

The mathematical apparatus creates protected space for phenomenological investigation by making reductive demands provably misdirected. Any critique that engages the apparatus must engage the mathematics (requiring years of study to assess), ignore the apparatus (appearing anti-intellectual within the institutional register the critique would have to operate in), or dismiss the apparatus (conceding the field). The trilemma is exhaustive: there is no fourth move available to a critique that wishes to remain within the rationalist register that motivated the demand for metrics in the first place.

Inside the protected space, phenomenological investigation can proceed without defending itself against reductive interruption. The chiasmic intertwining of organisation and environment can be examined for what it is. The reversibility and flesh structures of Part III can be developed. The temporal depth across multiple horizons that Part IV demonstrated empirically can be characterised without being reduced to scalar metrics. The organisational becoming that the genealogy from Beer through second-order cybernetics to chiasmic ontology has traced can be the substance of further investigation without each step being challenged on its measurement compatibility.

As Ashby (1956, p. 207) formulated the regulatory principle that the horizon-mismatch theorem of §24 sharpens: “only variety can destroy variety.” The mathematical variety of the apparatus this paper deploys destroys the reductive simplicity of the institutional demand for metrics. The argument is not that organisations should not be reduced to metrics. The argument is that the horizon-mismatch theorem and the Novikov-Boone result in the proof register, supplemented by the structural-analogue articulations of §§27-29, together establish the lower bound: any regulator operating at horizon shorter than its substrate cannot capture the substrate’s variety in any measurement frame, and algorithmic equivalence-checking over sufficiently rich organisational categories is undecidable. The structural-analogue apparatus articulates what any mathematical formalisation would need to capture and demonstrates that the most expressive available formalisms operate in the structural-analogue register only. The elementary formalisms underlying KPIs and dashboards fall short by both registers.

Part VI: Strategic Adjustments and Empirical Grounding

§32. Methodological Declaration

This paper employs a methodological strategy that requires explicit acknowledgement. The paper uses mathematical impossibility results to establish rigorous boundaries of formalisation in organisational theory rather than to model organisations. The approach follows a distinguished tradition in mathematics itself, where impossibility theorems have proven as valuable as constructive results. Lakatos (1976, p. 9) captures the position in *Proofs and Refutations*: “mathematics does not grow through a monotonous increase of the number of indubitably established theorems but through the incessant improvement of guesses by speculation and criticism.” The argument of this paper represents one such criticism, using mathematical rigour to demonstrate the limits of mathematical application to organisational phenomena.

The method has four steps. First, deploy the advanced mathematics available in the contemporary field. Second, articulate the structural features any formalisation would need to capture (marking the apparatus as structural analogue rather than constructive embedding, per §25.3). Third, establish in the proof register that the elementary formalisms underlying institutional metrics cannot meet the conditions Ashby’s law imposes (per §24). Fourth, create protected space for non-reductive organisational understanding (per §31).

The approach is rigorous in its application of mathematics to its own limits, which is what Gödel (1931) did for arithmetic and what the present argument attempts for organisational theory. A serious critique might argue that organisations do not necessarily instantiate the mathematical properties the apparatus invokes, and that the paper has not proven organisations *are* topoi. The response is threefold. First, the paper does not claim organisations are topoi; it claims that if mathematics could capture organisations, topos theory would be among

the most promising frameworks given its capacity to handle multiple truth values (organisational ambiguity), internal logic (organisational reasoning), and sheaf conditions (local-global properties), and the structural-analogue framing of §§27-29 makes this explicit. Second, if organisations do not instantiate advanced mathematical properties, they certainly do not instantiate the elementary properties assumed by KPIs (scalar addition), percentages (ratio measurement), averages (statistical aggregation), and rankings (total ordering). Third, the burden of proof lies with those claiming mathematical capture is possible. The paper has shown the most advanced attempts fall short at the level of constructive embedding while remaining usable as articulation. Defenders of elementary metrics must explain why arithmetic succeeds where topos theory operates only at the structural-analogue register.

§33. Empirical Grounding: Three Documented Cases

The mathematical argument of Part V claims that the horizon-mismatch theorem of §24 and the structural-analogue articulations of §§27-29 hold across organisational contexts. Three publicly documented cases instantiate the argument: a digital-transformation lawsuit between a Fortune 500 company and a global consultancy, a healthcare scandal that prompted the largest NHS public inquiry of the post-war period, and the only published empirical test of the architecture-organisation mirroring relationship in the software industry. The cases are deliberately chosen across contexts (commercial technology, public healthcare, open-source software) and methodologies (court documents, statutory inquiry, peer-reviewed quantitative study) to test whether the argument holds invariantly across them.

§33.1 The Hertz-Accenture Digital Transformation (2016-2019)

In January 2015, Hertz Corporation hired a new Chief Information Officer and committed to a comprehensive digital transformation of its enterprise IT environment. At that time, Hertz operated approximately 1,800 IT systems, six database vendors, and thirty separate rental-processing systems, according to documentation produced in subsequent litigation and a MuleSoft conference presentation of 2018 (Dolfing, 2022). In March 2016, Hertz engaged Accenture under a consulting services agreement that had been in place since 2004. Accenture recommended a five-layer digital platform architecture: a front-end layer using the Angular JavaScript framework, a content-management layer using Adobe Experience Manager, a microservices layer, an enterprise-services bus, and an integration layer connecting to legacy back-end systems. The recommendation and its consequences are set out in *Hertz Corporation v. Accenture LLP*, the complaint filed in the United States District Court for the Southern District of New York on 19 April 2019, with the amended complaint filed 20 June 2019 (Case No. 1:19-cv-03508).

The project was governed throughout by the standard agile-at-scale measurement apparatus: story-point velocity, sprint completion rates, integration milestones, user-acceptance test pass rates. By the standard metrics, the project was tracking acceptably through 2017. In September 2017, Hertz signed the statement of work for the first phase, with go-live scheduled for December 2017. The go-live was postponed one month later. By late 2017, Hertz had removed Accenture from the project and engaged IBM to complete the work. Hertz never deployed the system Accenture had been paid approximately USD 32 million to build.

The Hertz complaint alleges, in Count 2, that Accenture engaged in deceptive and unfair trade practices, specifically that Accenture elected, without Hertz's authorisation, to design the system in ways that locked Hertz into ongoing dependency on Accenture for maintenance; staffed the engagement with personnel lacking the experience Accenture had represented; and delivered code that did not satisfy basic requirements for security, accessibility, or performance. The complaint reproduces internal Accenture communications, obtained in discovery, in which Accenture personnel acknowledged the project was coming apart while public reporting through the agile-velocity framework remained on track. Approximately 2,000,000 documents were reviewed in discovery; approximately 500,000 were placed in evidence (Dolfing, 2022).

The case instantiates three of the paper's claims with documentary precision. First, the horizon-mismatch theorem of §24: Hertz's measurement frame operated at sprint-level horizon (two-week cycles, story-point velocity) while the substrate the project was reorganising operated at the horizon of Hertz's 1,800 legacy systems, six database vendors, and thirty rental-processing systems, accumulated across two decades. The regulator's variety was bounded by what the agile measurement frame could resolve; the substrate variety exceeded that frame by orders of magnitude that the metrics could not represent. Second, the N-th Order State of §8.1: the four features are documented in the Hertz pleadings. Internal communications show recursive acknowledgement that the framework's reporting did not capture project reality. Third, the three-mechanism reform foreclosure of §8.3: every reform proposal that arose within the project was processed by absorption (rephrased into agile-compatible velocity targets), attenuation (engineers' concerns dismissed as not following the methodology), or decomposition (hybrid proposals split, with the agile-compatible component implemented and the structural critique discarded).

The case ended through external shock rather than internal reform. The shock was provided by Hertz's new Chief Marketing Officer, hired in late 2017, who appears from the public record to have brought sufficient organisational standing to override the agile-velocity reporting framework. Within months of her arrival, Accenture had been removed. The reform required someone with institutional authority operating outside the project's measurement frame, exactly the pattern the impossibility result of §8.3 predicts.

The Hertz-Accenture case demonstrates that the three-mechanism foreclosure is not peculiar to public-sector organisations or planned economies. It operates with full force in commercial technology projects backed by Fortune 500 capital and the largest consultancy firms in the world.

§33.2 Mid Staffordshire NHS Foundation Trust (2005-2009)

Between January 2005 and March 2009, the Mid Staffordshire NHS Foundation Trust operated under intensified performance-management arrangements designed to secure Foundation Trust status, which the Trust obtained in February 2008. The arrangements required the Trust to meet specific operational targets: maximum accident-and-emergency waiting times of four hours, financial-balance requirements, and a range of performance indicators tracked through the Healthcare Commission's annual health-check ratings.

The Healthcare Commission investigation, published March 2009, found that the Trust's mortality rate had been substantially elevated throughout the period under review. Press reports during the public inquiry cited figures of between 400 and 1,200 excess deaths. The Francis Report itself, the Mid Staffordshire NHS Foundation Trust Public Inquiry chaired by Robert Francis QC, three volumes totalling 1,782 pages published 6 February 2013 (Francis, 2013), concluded that it would be unsafe to attribute a specific number of avoidable deaths from the available statistical data alone, focusing instead on documenting the conditions of care and the institutional mechanisms that permitted those conditions to persist.

The Francis Report's findings instantiate the three-mechanism reform foreclosure of §8.3 with documentary specificity that no constructed case could match. The Report reviewed evidence from more than 250 witnesses and approximately one million pages of documentary material. It found that nursing staff, doctors, patients, families, and external regulators had all attempted at various points to raise concerns about the standard of care at Stafford Hospital. The Report documents how those concerns were processed.

The first mechanism, absorption, appears throughout. Concerns articulated in the Trust's own performance-management vocabulary were processed as routine quality-improvement matters. Specific examples in the Report include incident reports filed by nursing staff that were classified as low-severity and absorbed into the Trust's monthly clinical-governance dashboard, where they appeared alongside routine items and lost institutional visibility (Francis, 2013, Volume 1, Chapter 8).

The second mechanism, attenuation, dominates the Report's findings. Nurses and patients who articulated concerns in non-target-compatible language (clinical judgement, dignity, basic patient care) lacked the institutional standing required for uptake. The Report documents that the Trust's senior management consistently prioritised attainment of the four-hour A&E target, financial balance, and Foundation Trust assessment criteria over the clinical concerns raised by ward-level staff. Concerns that did not translate into the

target framework were dismissed as anecdotal, isolated, or attributable to individual nursing-staff inadequacy rather than systemic conditions (Francis, 2013, Volume 2, Chapter 18).

The third mechanism, decomposition, operates on the families' campaigning. Julie Bailey's Cure the NHS campaign, founded after the death of her mother at Stafford Hospital in 2007, articulated concerns in a hybrid register: specific clinical harms (target-compatible language) and a broader claim that the institutional culture was producing those harms (non-target-compatible language). The Trust's response, documented in the Report, addressed the specific clinical harms through individual incident investigations while declining to engage with the systemic claim. Local managers handled the specific deaths; the systemic critique was not taken up. The hybrid was decomposed exactly as the theorem of §8.3 predicts.

The Mid Staffordshire case also instantiates the homeostatic triad of §8.2. Credentialism operated to attenuate ward nurses' concerns relative to management's metric-based reports, because management's reports met the Healthcare Commission's data requirements while nursing concerns did not. Performance amplification followed the predicted pattern: as the gap between target attainment and clinical reality widened, the Trust intensified its target-focused performance management rather than questioning the targets themselves, and the Foundation Trust assessment was pursued through the intensification. Source isolation marginalised staff who persistently raised concerns through transfers, performance-improvement procedures, and the cultural classification "not a team player." The Report names specific individuals subjected to these mechanisms (Francis, 2013, Volume 3, Chapter 22).

This state ended through external shock that arrived in two stages: first, the Healthcare Commission's independent investigation of 2008-2009, conducted without the Trust's authorisation; second, the public inquiries chaired by Robert Francis. Both shocks required institutional authority external to the Trust's measurement-and-incentive frame. The first Francis Report (2010) was conducted under the Inquiries Act 2005 by an authority the Trust could not absorb, attenuate, or decompose. The reform that followed was substantial: the Trust was dissolved in 2014, with services transferred to neighbouring trusts. The pattern matches the prediction of §8.3: internal reform was foreclosed by the three mechanisms, and change required perturbation of sufficient magnitude to exceed the regulatory frame's absorption capacity.

The Mid Staffordshire case demonstrates that the state this paper diagnoses operates in public-sector contexts with statutory accountability requirements, professional regulation, and decades of established quality-improvement practice. The professionalism of clinical staff and the existence of statutory regulators proved insufficient to defeat the three-mechanism foreclosure once the target-attainment regime had stabilised.

§33.3 Conway's Chiasm in Software Architecture

Section 11 of this paper introduced Conway's Chiasm: the proposition that organisational structure and product architecture co-constitute one another, rather than one being the cause and the other the effect. The proposition is testable. MacCormack, Baldwin, and Rusnak (2012) provide the test.

The MacCormack-Baldwin-Rusnak study appeared in *Research Policy* volume 41, issue 8, pages 1309-1324 (October 2012). The paper takes advantage of a natural experiment in the software industry: identical functional requirements being implemented by organisational forms at opposite ends of the coupling spectrum. At one end, commercial software firms with tightly-coupled organisational structure (shared employment, hierarchical authority, co-located development). At the other end, open-source communities with loosely-coupled organisational structure (distributed contributors, no shared employer, asynchronous communication). The mirroring hypothesis predicts that the products developed by loosely-coupled organisations will exhibit more modular architecture than the products developed by tightly-coupled organisations.

MacCormack and colleagues construct matched pairs of software products that fulfil the same function but were developed under the two organisational forms: matched pairs in word processing, operating-system kernels, database management, and other categories. They measure product modularity through Design Structure Matrix analysis, which quantifies the level of coupling between a product's components and the propagation potential of design changes across components, with the underlying method developed in Baldwin and Clark (2000).

The empirical finding (MacCormack, Baldwin, and Rusnak, 2012, p. 1318): in every matched pair examined, the product developed by the loosely-coupled organisation exhibited significantly higher modularity than the product developed by the tightly-coupled organisation. The magnitude of the difference reached a factor of six in propagation cost between component pairs. The mirroring hypothesis was supported in every pair, without exception.

The finding instantiates Conway's Chiasm rather than Conway's Law alone. Conway's Law (1968) is a one-way claim: organisation structure determines product architecture. The MacCormack-Baldwin-Rusnak finding is consistent with Conway's Law and cannot establish its directionality from cross-sectional data. What the finding does establish is mirroring: organisational coupling and product coupling are statistically inseparable. Subsequent work by Colfer and Baldwin (2016) reviews the broader literature and finds that the mirroring relationship persists across context, methodology, and time period. Where mirroring breaks down, it does so through organisations consciously decoupling their structure from the product they are building, in order to achieve the modularity the structure-imposed coupling would otherwise prevent (Sosa, Eppinger, and Rowles, 2004). The conscious decoupling itself confirms the underlying mirroring relationship; it is the mirroring's mechanism being overridden by deliberate organisational design.

The case supports Conway's Chiasm in the form stated in §11: organisation and architecture co-constitute one another through ongoing mutual specification, with neither term temporally prior. The universal support for mirroring in the matched-pair study and the documented difficulty of decoupling architecture from organisation in practice are inconsistent with treating organisation and architecture as separable entities that can be measured, optimised, or transformed independently. The proposition that an organisation can transform its architecture while holding its organisational structure constant, or restructure while leaving the architecture in place, is exactly the proposition the mirroring literature shows does not hold empirically across more than fifty years of software-industry observation.

The §31.3 case differs in evidentiary register from §§31.1 and 31.2. The Hertz and Mid Staffordshire cases are documentary: court filings and statutory inquiry reports. The MacCormack-Baldwin-Rusnak case is quantitative empirical work in a peer-reviewed journal. The convergence of the three cases across three widely different evidentiary registers is itself an empirical argument: the mechanisms the paper identifies operate across contexts that share no domain features beyond being organisations.

§33.4 Summary

Three cases, three registers, one underlying pattern. The Hertz-Accenture litigation demonstrates the three-mechanism reform foreclosure operating in commercial digital transformation under conditions of maximum measurement intensity. The Mid Staffordshire NHS public inquiry demonstrates the same three mechanisms operating in public healthcare under conditions of statutory regulation and professional accountability. The MacCormack-Baldwin-Rusnak study demonstrates the chiasmic co-constitution of organisation and product architecture as a quantitative empirical regularity across the software industry. The convergence across the three cases is the empirical argument. The mechanisms identified in Parts II through V and formalised in §§8 and 23-27 are not artefacts of one institutional context; they appear, with the structural features the paper specifies, across contexts that share no domain features beyond being organisations of human action coordinated by measurement and incentive.

§34. Constructive Elements: What Chiasmic Understanding Enables

The impossibility of calculation does not entail impossibility of intelligent action. It calls for a shift from a philosophy of control to one of cultivation and navigation. The following pairs map reductive approaches to chiasmic alternatives.

Where the reductive approach measures boundaries, the chiasmic alternative recognises transformation zones: the shift moves from lines to gradients, with the practical consequence that the organisation can work with

permeability rather than enforce walls. Where the reductive approach quantifies intelligence, the chiasmic alternative senses obstruction patterns: the shift moves from metrics to topology, preserving the conditions that enable adaptation. Where the reductive approach calculates change, the chiasmic alternative anticipates remainder: the shift moves from planning to expecting debt, preparing for what cannot be eliminated. Where the reductive approach optimises locally, the chiasmic alternative respects entanglement: the shift moves from silos to fields, recognising that changes reverberate across the entire structure. Where the reductive approach controls variance, the chiasmic alternative maintains requisite variety: the shift moves from reduction to preservation, with variety understood as the condition of adaptation rather than its enemy. Where the reductive approach defines objectives, the chiasmic alternative enacts purposes: the shift moves from targets to emergence, with objectives discovered through pursuit rather than fixed in advance.

The pairs translate into concrete leadership practices. Instead of enforcing strict departmental silos with defined handoff procedures, the organisation can create formal boundary-spanning roles whose responsibility is the quality of interaction between teams. The success of such roles is assessed not by ticket closure rates but by the richness of knowledge transfer and emergence of shared understanding across supposed boundaries. A Customer Experience Translator sitting between support and engineering, responsible for ensuring each team understands the other's constraints and possibilities, is the kind of role the chiasmic ontology supports.

Instead of sprint retrospectives focused on velocity improvement, the organisation can hold sessions where teams collectively map recurring friction points, discovering that delays cluster around specific handoffs and revealing architectural rather than performance issues. The resulting patterns, not individual complaints, become organisational intelligence made visible through obstruction-mapping practice. Management's role shifts from measuring output to cultivating the capacity to sense and navigate these obstruction topologies.

For every major initiative, the organisation can allocate a formal remainder budget, typically 20-30% of project capacity, pre-dedicated to addressing unforeseen consequences, cleanup, and system reintegration. The allocation is honest accounting of what change actually costs in chiasmic systems, rather than the conventional "slack" framing that treats remainder work as inefficiency. A platform migration that allocates three engineers for remainder work, whose explicit job is handling the emergent problems that were not in the plan (unexpected dependencies, legacy workarounds now surfacing, integration debt), implements the recognition the derived-category articulation of §26 supplies.

Instead of static quarterly objectives, the organisation can use directional vectors: teams assessed on demonstrated progress along a direction, not achievement of a fixed point, allowing adaptation to discoveries made en route. "Move toward deeper customer understanding of our value proposition" supports the kind of qualitative judgement the chiasmic ontology recognises, where "increase conversion rate by 15%" forces the team into the absorption-attenuation-decomposition cycle of §8.3.

The organisation can match the diversity of internal responses to the complexity of challenges faced, allowing multiple context-specific ways of working rather than imposing uniform processes. Different product teams using different methodologies (some Kanban, some Scrum, some ad-hoc) based on their specific coupling with customers and technology embodies what Ashby's law has been pointing at: regulatory variety matched to environmental variety, deployed at the practice level rather than the metric level. Leadership's role is ensuring adequate knowledge transfer between approaches, rather than enforcing uniformity that destroys the variety the practices need.

When teams consistently work around the official deployment process, the chiasmic recognition is that the workaround likely encodes critical knowledge about system behaviour that the official process ignores. The response is to study the workaround rather than punish it. The pattern recognition is the substance of organisational intelligence; the official process is a partial map that has become outdated.

The practices are not techniques to be implemented like software updates. They are stances, orientations, ways of attending to organisational reality that acknowledge its chiasmic, autopoietic nature. They cannot be reduced to checklists because they are modes of engagement with incalculability itself.

§34.1 From Control to Cultivation

The shift from reductive to chiasmic organisation is a philosophical shift rather than a methodological one. It requires leadership to relinquish the fantasy of total knowledge and embrace a different relationship with organisational reality. Leaders in chiasmic organisations cultivate sensing capacity rather than demanding reports: they create conditions where signals about organisational reality can emerge and be heard, rather than filtering everything through predetermined metrics. They work explicitly with remainder: after every major change, they ask not “did we hit our targets?” but “what emerged that we did not anticipate, and what does that teach us about the system?” They navigate by obstruction: when work becomes difficult, they ask “what is the system telling us?” rather than “how do we remove this obstacle?” because some obstructions are the system’s immune response protecting essential complexity. They protect boundary intelligence: they recognise that the most valuable organisational knowledge often exists at interfaces (between teams, between organisation and environment) and invest in roles and practices that make boundary work visible and valued. They enact purpose through experimentation: they understand that organisational purpose is not discovered through strategic planning but emerges through action, and they create fast feedback loops and preserve the capacity to change direction based on what is learned.

None of these practices are measurable in the conventional sense. Their success is visible in system behaviour: increased adaptability, reduced burnout, better retention of essential knowledge, more real innovation. These outcomes resist quantification because they are chiasmic, existing in the intertwining of organisation and environment, observer and observed, measurement and measured. The practices are rigorous in the application sense; they apply careful attention to what actually exists rather than demanding it conform to what can be calculated. Spinoza’s formulation captures the position: “no one has yet determined what the body can do” (*Ethics*, III, P2 Scholium). No one has determined what organisations can do when freed from the illusion of mathematical control.

§35. Preparing the Counterargument

The argument anticipates six characteristic objections, each addressed below.

“But we still need metrics; we cannot manage without measurement.” The objection concedes the argument while attempting to reject it. The paper does not claim organisations can exist without tracking or awareness. It demonstrates that formalisation at the level of KPIs and dashboards destroys the very intelligence it claims to measure. The horizon-mismatch theorem of §24 establishes the structural reason: any regulator operating at horizon shorter than its substrate cannot capture the substrate’s variety in any measurement frame the regulator deploys. The objection assumes measurement approximates organisational reality where the argument shows measurement transforms that reality, destroying what it claims to capture. Saying “we need metrics” after engaging this argument is structurally analogous to saying “we need phlogiston” after the oxygen theory of combustion: the need is grounded in a misapprehension of the phenomenon.

“This is too theoretical for practical application.” The objection demonstrates the blindness the paper diagnoses. Part IV’s Tuesday-morning phenomenology shows that chiasmic, autopoietic, and recursive phenomena are not rare theoretical exceptions but the ordinary substance of daily organisational life. Every bug report, every architectural decision, every attempt to categorise work for reporting contains the complexity the paper describes. The complaint that the argument is too theoretical typically means “I cannot recognise my own organisational reality in this description,” and the inability is internal to the conventional frameworks the objector has internalised. Metrics, KPIs, and objective measurement create the blindness that makes the argument seem theoretical. The objection proves the thesis: conventional management frameworks generate their own justification by destroying the capacity to recognise what they destroy.

“But companies using metrics succeed; they cannot all be wrong.” The objection confuses correlation with causation in the way second-order cybernetics warned against. Organisations succeed through the unmeasured remainder, through the shadow systems, informal networks, workarounds, and tacit knowledge that develop

around official metrics, rather than because of the metrics themselves. Part IV documents the dynamic in detail. The actual work happens in the spaces metrics cannot see, and “success” is in any case measured how, if not by the metrics whose adequacy the paper has put into question. The objection assumes the legitimacy of the measurement framework to defend the measurement framework, which is circular reasoning that the cascade argument severs. Organisations persist through their autopoietic, chiasmic capacity to recreate themselves through processes that exceed formalisation. Metrics do not enable this persistence. At best they are noisy. At worst they degrade organisational intelligence by forcing energy into performance maintenance rather than adaptive response.

“Surely some measurement is better than none.” The objection commits a category error by assuming metrics approximate organisational reality with some degree of inaccuracy. The cohomological argument of §25 articulates the structural feature: measurement at the elementary formalism level annihilates the higher-order structure (the H^n groups) where organisational intelligence resides. The structural-analogue articulation supports the substantive claim: measurement of H^0 content is not measurement with error bars; it is the systematic discarding of the higher cohomology where the intelligence lives. The analogy is direct: saying “some measurement is better than none” is structurally analogous to saying “some bloodletting is better than none” when treating bacterial infection. The intervention operates on a different ontological level than the problem, and at that level the intervention damages rather than approximates.

“Organisations are not really topoi or cohomological structures.” Precisely. The objection misunderstands the argumentative direction. Organisations are not topoi; they are more complex than topos theory can capture. The cascade argument does not claim organisations are topoi. It claims that even if they were, even if humanity’s most advanced mathematical tools could represent them, capture would still fall short in the constructive-embedding sense (per the §25.3 distinction). Therefore attempting to measure them with elementary formalisms (arithmetic in KPIs, real analysis in averages) is structurally inadequate, with the inadequacy quantified by the horizon-mismatch theorem of §24. The objector has accepted the argument while misunderstanding its direction.

“We can optimise locally even if we cannot measure everything.” The objection misunderstands the topological nature of organisational reality. In chiasmic systems, elements are not independently modifiable; they exist only through their relations. Mathematically, the feature is articulated by the fact that in derived categories and ∞ -topoi, morphisms (relationships) are not external connections between pre-existing objects; they constitute the objects themselves. The cohomological structure of §25 articulates the consequence: if organisational intelligence resides in higher cohomology groups H^n , then these groups are defined by exact sequences where each element is simultaneously the kernel of one map and the image of another, and changing any single element forces changes to propagate throughout the entire sequence. There is no local in a structure where locality is defined relationally.

The spider’s web analogy captures the feature in compressed form. The web vibrates as one because it is one: one differentiation, one continuous intertwining where the distinction between local and global is itself produced by the weaving rather than antecedent to it. The web cannot be touched at one point without the entire structure moving, not because of mechanical linkage between independent points but because the points are nothing more than nodes in the pattern of tensions that the whole web maintains. Local optimisation in such a structure is not an adjustment of one variable in a multivariate function; it is the perturbation of a node in a coherent sheaf, forcing the entire structure to reorganise around the perturbation. The optimisation succeeds locally (the velocity metric increases) through global reorganisation (work gets redefined to be measurable, architectural debt accumulates, shadow systems proliferate). The local-versus-global distinction does not apply at the ontological level the chiasmic argument operates at; it is an artefact of the measurement frame that the horizon-mismatch theorem of §24 has put into question.

§36. Conclusion

Three connected lines of argument run through the paper. The historical analysis has traced the genealogy from Beer's cybernetic insights through second-order critique to phenomenological alternatives, and has shown this to be a necessary evolution through the internal logic of the positions taken. Each phase revealed limitations in the previous, driving theoretical development through internal necessity rather than external fashion.

The philosophical development has established the chiasmic ontology as the account of what organisations are once the observer-stance has become untenable. Organisations are patterns within fields they co-create, exhibiting reversibility, autopoiesis, and emergent intelligence that exceed formal capture in the constructive-embedding sense. The chiasmic ontology recovers what the second-order critique displaced and supplies the framework within which the mathematical apparatus operates.

The strategic and mathematical contribution has created protected space for non-reductive organisational understanding through the proof-register foundation supplied by the horizon-mismatch theorem of §24 and the Novikov-Boone result of §26, supplemented by the structural-analogue articulations of §§27-29. The argument deploys the advanced apparatus to establish structural limits on what mathematical formalisation can do with organisational phenomena, while remaining honest about which register supplies which result.

§36.1 The Honest Paradox

The paper uses mathematics to establish the structural limits of mathematics when applied to organisational reality. The procedure has precedent: Gödel used arithmetic to prove arithmetic's incompleteness, Turing used computation to prove computation's limits, and Arrow used preference logic to prove preference aggregation's impossibility. The present argument follows the same tradition, using categorical mathematics and information-theoretic constraints to demonstrate that organisational phenomena exceed categorical capture in the constructive-embedding register while remaining articulable in the structural-analogue register. As Wittgenstein (1953, Part I, remark 309) put the philosophical task, "what is your aim in philosophy? To show the fly the way out of the fly-bottle." The mathematical apparatus shows the way out of the metric-fly-bottle by deploying the institutional register's preferred language to mark its own limits.

§36.2 The Space Created

The space the apparatus creates supports phenomenological investigation without reductive interruption, recognition of chiasmic intertwining as structural rather than metaphorical, respect for organisational intelligence as emergent from obstruction, acceptance of transformation remainder as necessary feature, and engagement with organisational reality on its own terms. The space is now secured, since demands for KPIs, impositions of dashboards, and the institutional pretence of optimisation face mathematical results in the proof register (the horizon-mismatch theorem and Novikov-Boone application) and structural-analogue articulations in the higher register. The protection operates at technical, philosophical, and strategic levels.

§36.3 Future Directions

The immediate next step is the development of a fully worked chiasmic praxis built upon the principles set out in §32. The development is not a call for new methodologies to be marketed to organisations; it is an invitation to practitioners to develop skilful navigation within acknowledged incalculability. A chiasmic toolkit cannot provide algorithms for control, but it can provide heuristics for cultivation: ways of attending to organisational reality that honour its autopoietic, chiasmic, and incalculable nature. Such a praxis cannot be formalised in the conventional sense because formalisation in the constructive-embedding sense is what the paper has shown is structurally inadequate to organisational phenomena. The praxis can be cultivated through communities of practice, through apprenticeship in sensing obstruction patterns, through developing the embodied skill of working with remainder rather than against it. The work that follows from this paper is inhabitation of its insights rather than implementation of its findings.

Three companion papers extend the argument in directions this one has not pursued. The first develops the thermodynamics of organisational extraction, formalising the three-site absorption (worker, substrate, externality) implicit in Part IV's Tuesday-morning narrative. The second develops the dissolution of orders, treating the breakdown of organisational orderings as a structural feature of late-modern institutional life rather than an exceptional crisis. The third develops the membrane architecture, articulating the chiasmic ontology of organisational boundaries through the formal apparatus of biological membranes. Together with the present paper, the four constitute a post-cybernetic series on organisational ontology.

Empirical studies grounded in the chiasmic model would investigate boundary zones, reversibility patterns, and emergence dynamics in actual organisations. The methodological challenge is direct: any empirical methodology that imposes the observer-stance the chiasmic argument displaces will fail to capture the phenomenon it seeks to study. The work points toward participatory and phenomenological research methodologies developed in the traditions of Schön (1983), Lave and Wenger (1991), and Stiegler (1998-2001). The philosophical implications of the argument extend beyond organisational theory to other domains where formalisation is treated as the default register of legitimate knowledge: legal reasoning, economic measurement, educational assessment, public-health metrics. The structural features the paper identifies are not peculiar to organisations; they appear in any social phenomenon where the measurement apparatus is part of what is measured.

The thesis can be stated in compressed form, with which the argument concludes. An organisation cannot calculate the field it is part of co-creating through the act of calculation. The attempt to do so generates a stable configuration in which the inadequacy is collectively known and institutionally maintained. This state absorbs the impossibility into human bodies, organisational substrates, and discharged externalities. The argument's empirical anchor is the three documented cases of §31. The argument's structural anchor is the horizon-mismatch theorem of §24. The argument's phenomenological anchor is the Tuesday-morning narrative of Part IV. The three anchors converge on a single point. The cost of insisting on calculation where chiasmic relation is the structure is paid by people whose lives, capacities, and relationships absorb what the metrics cannot register. The argument of this paper is that the cost is structural and consistent across institutional contexts, and the response is not better metrics. The response is the recognition that what the metrics cannot register is, in fact, where the work is done and the cost is paid.

References

- Abbott, A. (1988). *The System of Professions: An Essay on the Division of Expert Labor*. Chicago: University of Chicago Press.
- Alvesson, M. and Robertson, M. (2006). 'The construction, significance and effects of elite identities in consulting firms'. *Organization*, 13(2), pp. 195-224.
- Alvesson, M. and Spicer, A. (2016). *The Stupidity Paradox: The Power and Pitfalls of Functional Stupidity at Work*. London: Profile Books.
- Alvesson, M. and Willmott, H. (2002). 'Identity regulation as organizational control: producing the appropriate individual'. *Journal of Management Studies*, 39(5), pp. 619-644.
- Arrow, K. J. (1951). *Social Choice and Individual Values*. New York: Wiley.
- Arthur, W. B. (1994). *Increasing Returns and Path Dependence in the Economy*. Ann Arbor: University of Michigan Press.
- Ashby, W. R. (1956). *An Introduction to Cybernetics*. London: Chapman and Hall.
- Aumann, R. J. (1976). 'Agreeing to disagree'. *The Annals of Statistics*, 4(6), pp. 1236-1239.
- Baldwin, C. Y. and Clark, K. B. (2000). *Design Rules, Volume 1: The Power of Modularity*. Cambridge, MA: MIT Press.

- Barad, K. (2007). *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*. Durham, NC: Duke University Press.
- Bateson, G. (1972). *Steps to an Ecology of Mind*. Chicago: University of Chicago Press.
- Bateson, G. (1979). *Mind and Nature: A Necessary Unity*. New York: E. P. Dutton.
- Bateson, G., Jackson, D. D., Haley, J. and Weakland, J. (1956). 'Toward a theory of schizophrenia'. *Behavioral Science*, 1(4), pp. 251-264.
- Beckford, J. (1995). *The Viable System Model: A More Adequate Tool for Practising Management?* PhD Thesis, University of Hull.
- Beer, S. (1972). *Brain of the Firm*. London: Allen Lane.
- Beer, S. (1979). *The Heart of Enterprise*. Chichester: Wiley.
- Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester: John Wiley and Sons.
- Berliner, J. S. (1957). *Factory and Manager in the USSR*. Cambridge, MA: Harvard University Press.
- Bertalanffy, L. von (1968). *General System Theory: Foundations, Development, Applications*. New York: George Braziller.
- Boone, W. W. (1958). 'The word problem'. *Annals of Mathematics*, 70(2), pp. 207-265.
- Brown, J. S. and Duguid, P. (1991). 'Organizational learning and communities-of-practice: toward a unified view of working, learning, and innovation'. *Organization Science*, 2(1), pp. 40-57.
- Brunsson, N. (1989). *The Organization of Hypocrisy: Talk, Decisions and Actions in Organizations*. Chichester: John Wiley and Sons.
- Cartwright, S. and Schoenberg, R. (2006). 'Thirty years of mergers and acquisitions research: recent advances and future opportunities'. *British Journal of Management*, 17(S1), pp. S1-S5.
- Cataldo, M., Wagstrom, P. A., Herbsleb, J. D. and Carley, K. M. (2006). 'Identification of coordination requirements: implications for the design of collaboration and awareness tools'. *Proceedings of CSCW '06*, pp. 353-362.
- Checkland, P. (1981). *Systems Thinking, Systems Practice*. Chichester: John Wiley and Sons.
- Colfer, L. J. and Baldwin, C. Y. (2016). 'The mirroring hypothesis: theory, evidence and exceptions'. *Industrial and Corporate Change*, 25(5), pp. 709-738.
- Collins, H. and Evans, R. (2007). *Rethinking Expertise*. Chicago: University of Chicago Press.
- Conway, M. E. (1968). 'How do committees invent?' *Datamation*, 14(5), pp. 28-31.
- Dikert, K., Paasivaara, M. and Lassenius, C. (2016). 'Challenges and success factors for large-scale agile transformations: a systematic literature review'. *Journal of Systems and Software*, 119, pp. 87-108.
- Dolfing, H. (2022). *Case Study 8: How Hertz Paid Accenture USD 32 Million for a Website That Never Went Live*. Available at: <<https://www.henricodolfing.ch/en/case-study-8-how-hertz-paid-accenture-32-million-for-a-website-that-never-went-live/>> (Accessed: 25 May 2026).
- Engelen, E., Ertürk, I., Froud, J., Johal, S., Leaver, A., Moran, M., Nilsson, A. and Williams, K. (2011). *After the Great Complacency: Financial Crisis and the Politics of Reform*. Oxford: Oxford University Press.
- Espejo, R. and Harnden, R. (eds.) (1989). *The Viable System Model: Interpretations and Applications of Stafford Beer's VSM*. Chichester: John Wiley and Sons.
- Fayol, H. (1916/1949). *General and Industrial Management*. Translated by C. Storrs. London: Pitman.
- Feynman, R. P. (1974). 'Cargo cult science'. Caltech commencement address. Reprinted in *Surely You're Joking, Mr. Feynman!* New York: W. W. Norton, 1985, pp. 338-346.

- Fleming, P. and Spicer, A. (2003). 'Working at a cynical distance: implications for power, subjectivity and resistance'. *Organization*, 10(1), pp. 157-179.
- Flyvbjerg, B. (2001). *Making Social Science Matter: Why Social Inquiry Fails and How It Can Succeed Again*. Translated by S. Sampson. Cambridge: Cambridge University Press.
- Foerster, H. von (1981). *Observing Systems*. Seaside, CA: Intersystems Publications.
- Foerster, H. von (1992). 'Ethics and second-order cybernetics'. *Cybernetics and Human Knowing*, 1(1), pp. 9-19.
- Foerster, H. von (2003). *Understanding Understanding: Essays on Cybernetics and Cognition*. New York: Springer-Verlag.
- Francis, R. (2010). *Independent Inquiry into Care Provided by Mid Staffordshire NHS Foundation Trust January 2005-March 2009*, Volumes I-II. London: The Stationery Office. HC375-I and HC375-II.
- Francis, R. (2013). *Report of the Mid Staffordshire NHS Foundation Trust Public Inquiry*, Volumes 1-3. London: The Stationery Office. ISBN 978-0-10-298146-9.
- Gelfand, S. I. and Manin, Y. I. (2003). *Methods of Homological Algebra*. 2nd edn. Berlin: Springer-Verlag.
- Gauss, C. F. (1827). *Disquisitiones generales circa superficies curvas*. Göttingen: Dieterich. (Theorema Egregium.)
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Gödel, K. (1931). 'Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I'. *Monatshefte für Mathematik und Physik*, 38, pp. 173-198.
- Goldman, M. I. (1991). *What Went Wrong with Perestroika*. New York: W. W. Norton.
- Graeber, D. (2018). *Bullshit Jobs: A Theory*. New York: Simon and Schuster.
- Grothendieck, A. (1957). 'Sur quelques points d'algèbre homologique'. *Tôhoku Mathematical Journal*, 9(2), pp. 119-221.
- Hartshorne, R. (1977). *Algebraic Geometry*. New York: Springer-Verlag.
- Hayek, F. A. (1945). 'The use of knowledge in society'. *American Economic Review*, 35(4), pp. 519-530.
- Healthcare Commission (2009). *Investigation into Mid Staffordshire NHS Foundation Trust*. London: Commission for Healthcare Audit and Inspection.
- Heidegger, M. (1927/1962). *Being and Time*. Translated by J. Macquarrie and E. Robinson. New York: Harper and Row.
- Hertz Corporation v. Accenture LLP (2019). Complaint, United States District Court for the Southern District of New York, Case No. 1:19-cv-03508. Filed 19 April 2019; amended 20 June 2019.
- Howison, J. and Crowston, K. (2014). 'Collaboration through open superposition: a theory of the open source way'. *MIS Quarterly*, 38(1), pp. 29-50.
- Huybrechts, D. (2006). *Fourier-Mukai Transforms in Algebraic Geometry*. Oxford: Oxford University Press.
- Jackson, M. C. (2003). *Systems Thinking: Creative Holism for Managers*. Chichester: John Wiley and Sons.
- Johnstone, P. T. (1977). *Topos Theory*. London: Academic Press.
- Kieser, A. (2002). 'On communication barriers between management science, consultancies and business organizations'. In: T. Clark and R. Fincham (eds.), *Critical Consulting: New Perspectives on the Management Advice Industry*. Oxford: Blackwell, pp. 206-227.
- Kornai, J. (1992). *The Socialist System: The Political Economy of Communism*. Oxford: Clarendon Press.
- Korzybski, A. (1933). *Science and Sanity: An Introduction to Non-Aristotelian Systems and General Semantics*. Lakeville, CT: International Non-Aristotelian Library Publishing Company.

- Lakatos, I. (1976). *Proofs and Refutations: The Logic of Mathematical Discovery*. Cambridge: Cambridge University Press.
- Lave, J. and Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge: Cambridge University Press.
- Lewis, D. (1969). *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press.
- Lorde, A. (1979). 'The master's tools will never dismantle the master's house'. Comments at *The Personal and the Political Panel*, Second Sex Conference, New York. Reprinted in *Sister Outsider*, Trumansburg, NY: Crossing Press, 1984, pp. 110-113.
- Luhmann, N. (1995). *Social Systems*. Translated by J. Bednarz Jr. with D. Baecker. Stanford: Stanford University Press.
- Mac Lane, S. (1971). *Categories for the Working Mathematician*. New York: Springer-Verlag.
- MacCormack, A., Baldwin, C. and Rusnak, J. (2012). 'Exploring the duality between product and organizational architectures: a test of the "mirroring" hypothesis'. *Research Policy*, 41(8), pp. 1309-1324. DOI: 10.1016/j.respol.2012.04.011.
- MacKenzie, D. (2006). *An Engine, Not a Camera: How Financial Models Shape Markets*. Cambridge, MA: MIT Press.
- MacKenzie, D. (2011). 'The credit crisis as a problem in the sociology of knowledge'. *American Journal of Sociology*, 116(6), pp. 1778-1841.
- MacKenzie, D. and Spears, T. (2014). 'The formula that killed Wall Street: the Gaussian copula and modelling practices in investment banking'. *Social Studies of Science*, 44(3), pp. 393-417.
- March, J. G. (1991). 'Exploration and exploitation in organizational learning'. *Organization Science*, 2(1), pp. 71-87.
- Maturana, H. R. (1978). 'Biology of language: the epistemology of reality'. In: G. A. Miller and E. Lenneberg (eds.), *Psychology and Biology of Language and Thought*. New York: Academic Press, pp. 27-63.
- Maturana, H. R. and Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: D. Reidel.
- Maturana, H. R. and Varela, F. J. (1987). *The Tree of Knowledge: The Biological Roots of Human Understanding*. Boston: Shambhala.
- McCleary, J. (2001). *A User's Guide to Spectral Sequences*. 2nd edn. Cambridge: Cambridge University Press.
- Medina, E. (2011). *Cybernetic Revolutionaries: Technology and Politics in Allende's Chile*. Cambridge, MA: MIT Press.
- Merleau-Ponty, M. (1945/2012). *Phenomenology of Perception*. Translated by D. A. Landes. London: Routledge.
- Merleau-Ponty, M. (1968). *The Visible and the Invisible*. Edited by C. Lefort, translated by A. Lingis. Evanston, IL: Northwestern University Press.
- Meyer, J. W. and Rowan, B. (1977). 'Institutionalized organizations: formal structure as myth and ceremony'. *American Journal of Sociology*, 83(2), pp. 340-363.
- Mockus, A., Fielding, R. T. and Herbsleb, J. D. (2002). 'Two case studies of open source software development: Apache and Mozilla'. *ACM Transactions on Software Engineering and Methodology*, 11(3), pp. 309-346.
- Morrison, E. W. and Milliken, F. J. (2000). 'Organizational silence: a barrier to change and development in a pluralistic world'. *Academy of Management Review*, 25(4), pp. 706-725.
- Nove, A. (1977). *The Soviet Economic System*. London: George Allen and Unwin.

- Novikov, P. S. (1955). 'On the algorithmic unsolvability of the word problem in group theory'. *Trudy Matematicheskogo Instituta Imeni V.A. Steklova*, 44, pp. 1-143.
- Orlikowski, W. J. (2000). 'Using technology and constituting structures: a practice lens for studying technology in organizations'. *Organization Science*, 11(4), pp. 404-428.
- Pfeffer, J. and Sutton, R. I. (2006). *Hard Facts, Dangerous Half-Truths, and Total Nonsense: Profiting from Evidence-Based Management*. Boston: Harvard Business School Press.
- Power, M. (1997). *The Audit Society: Rituals of Verification*. Oxford: Oxford University Press.
- Russell, B. and Whitehead, A. N. (1910-1913). *Principia Mathematica*. 3 vols. Cambridge: Cambridge University Press.
- Ryan, K. D. and Oestreich, D. K. (1991). *Driving Fear Out of the Workplace: How to Overcome the Invisible Barriers to Quality, Productivity, and Innovation*. San Francisco: Jossey-Bass.
- Sartre, J.-P. (1943/1956). *Being and Nothingness: An Essay on Phenomenological Ontology*. Translated by H. E. Barnes. New York: Philosophical Library.
- Schön, D. A. (1983). *The Reflective Practitioner: How Professionals Think in Action*. New York: Basic Books.
- Shannon, C. E. (1948). 'A mathematical theory of communication'. *Bell System Technical Journal*, 27(3), pp. 379-423.
- Sosa, M. E., Eppinger, S. D. and Rowles, C. M. (2004). 'The misalignment of product architecture and organizational structure in complex product development'. *Management Science*, 50(12), pp. 1674-1689.
- Spencer-Brown, G. (1969). *Laws of Form*. London: George Allen and Unwin.
- Stiegler, B. (1998-2001). *Technics and Time*, Volumes 1-3. Translated by R. Beardsworth, G. Collins and S. Barker. Stanford: Stanford University Press.
- Strathern, M. (ed.) (2000). *Audit Cultures: Anthropological Studies in Accountability, Ethics and the Academy*. London: Routledge.
- Taleb, N. N. (2007). *The Black Swan: The Impact of the Highly Improbable*. New York: Random House.
- Taylor, F. W. (1911). *The Principles of Scientific Management*. New York: Harper and Brothers.
- Turing, A. M. (1936). 'On computable numbers, with an application to the Entscheidungsproblem'. *Proceedings of the London Mathematical Society*, 42(2), pp. 230-265.
- Univalent Foundations Program (2013). *Homotopy Type Theory: Univalent Foundations of Mathematics*. Princeton: Institute for Advanced Study. Available at: <<https://homotopytypetheory.org/book/>>
- Varela, F. J. (1979). *Principles of Biological Autonomy*. New York: North-Holland.
- Varela, F. J., Thompson, E. and Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.
- Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. Chicago: University of Chicago Press.
- Verdier, J.-L. (1996). *Des catégories dérivées des catégories abéliennes*. Astérisque 239. Paris: Société Mathématique de France.
- Vial, G. (2019). 'Understanding digital transformation: a review and a research agenda'. *Journal of Strategic Information Systems*, 28(2), pp. 118-144.
- Voevodsky, V. (2006). 'A very short note on the homotopy lambda-calculus'. Unpublished manuscript. Available at: <<https://www.math.ias.edu/~vladimir/>>
- Weber, M. (1922/1947). *The Theory of Social and Economic Organization*. Translated by A. M. Henderson and T. Parsons. New York: Free Press.
- Weick, K. E. (1979). *The Social Psychology of Organizing*. 2nd edn. Reading, MA: Addison-Wesley.

Weick, K. E. (1995). *Sensemaking in Organizations*. Thousand Oaks, CA: Sage.

Wiener, N. (1948). *Cybernetics: Or Control and Communication in the Animal and the Machine*. Cambridge, MA: MIT Press.

Wittgenstein, L. (1953). *Philosophical Investigations*. Translated by G. E. M. Anscombe. Oxford: Blackwell.

Zuboff, S. (1988). *In the Age of the Smart Machine: The Future of Work and Power*. New York: Basic Books.

The Archival Regulator

A General Theory of Regulation Against Metabolic Substrates

Author: Nourizadeh, Moreno ORCID: 0009-0006-0174-2585

Abstract

Four results have accumulated across the study of quantified coordination, each proved on its own ground and each apparently addressed to a different object. The first concerns representation: any valuation that assigns each unit of an organisation a single figure and reads the whole as the sum across its units represents the organisation faithfully if and only if every relational surplus among those units is exactly zero, so that the per-unit form is blind, of either sign, to the entire stratum of value that exists only in the joining of specific people (the decomposition into coalition-level dividends established by Harsanyi, 1959, 1963, following von Neumann and Morgenstern, 1944, and Shapley, 1953; the same conclusion returned independently by the vanishing of the synergistic information atom under per-member marginalisation in Williams and Beer, 2010). The second concerns conservation: the total variety required to hold an organisation coherent against its disturbances is fixed by the environment and not by the instrument, so that whatever the formal apparatus does not capture is not eliminated by the apparatus but displaced onto the human beings who bridge the gap between the representation and the work, at a cost paid in attention, cognition, and metabolism (the constraint following from Ashby, 1956, read against the impossibility of total formal capture). The third concerns time: a regulator whose action-cycle is shorter than the replenishment-cycle of the system it acts on by more than a factor of order one cannot, by construction, model that system at the system's own resolution, so that the system's slow variety is absorbed within the system itself and the regulator receives its first legible signal of depletion only after the depletion has run past the horizon at which intervention could have arrested it (a composition of Ashby, 1956, Conant and Ashby, 1970, and the rate-distortion allocation of Shannon, 1959). The fourth concerns dynamics: under the conditions the first three results specify, a reallocation driven by the per-unit instrument can produce the very improvement it reports by drawing down the relational stock it cannot see, so that a sustained record of measured success is consistent with the consumption of the capacity that produced it, and the instrument cannot distinguish the two cases from inside its own ledger.

The present paper argues that these four results are one result. What joins them is a distinction between two modes in which a coordinating capacity can exist. A substrate is metabolic when it accrues only through lived time under thermodynamic cost, through sustained shared work among specific people, through the apprenticeship transmission Polanyi identified as the sole vehicle for an art that cannot be specified in detail (Polanyi, 1958, 1966), and when its state at any moment is the sediment of a history that cannot be re-run faster than it was lived. An instrument is archival when it holds its content as retrievable record, indifferent to the duration of its own accumulation, reproducible without cost proportional to what it records, and composed of per-unit entries whose whole is their sum. The paper's central claim, which I will call the archival regulator theorem, is that any archival instrument set to regulate a metabolic substrate exhibits, jointly and by construction, the four properties the accumulated results describe: representational blindness to the substrate's relational stratum, conservative displacement of the uncaptured variety onto living absorbers, temporal belatedness with respect to the substrate's cycle, and extractive dynamics wherever the instrument drives reallocation; and that the four properties are not four defects to be repaired separately but four aspects of the single fact that the instrument's form of existence and the substrate's form of existence do not share a mode of being. The theorem is developed constructively, its instances are then read across four apparatuses that the contemporary order treats as unrelated (the delivery dashboard, the price, the prescriptive organisational framework, and the large language model), the second-order closure is established (the demand for a better instrument is itself the theorem's signature, since a finer archival instrument lands inside the same form), and the constructive direction is stated: regulation of a metabolic substrate can proceed only from within the substrate's own mode, by proximity, by disposition, and by conditions

rather than prescriptions, which is a direction the archival form can point toward and can never itself perform. The theorem's components are severally proved in their own formal homes, and their unity is stated in a single frame, the archival representation map and its kernel, developed in Part II at the level of definitions and propositions; substrates mixed in kind decompose under the same frame, the theorem applying to the metabolic component and asserting nothing about the remainder.

Keywords: requisite variety, Conant-Ashby, Harsanyi dividends, conservation of variety, horizon mismatch, cannibalisation, tacit knowledge, Polanyi, metabolic substrate, archival regulation, technique, Gestell, price system, organisational frameworks, language models.

Part I: Four Results in Search of Each Other

1.1 A Quarterly Review

An executive committee convenes on the second Tuesday of the quarter. The room receives a dashboard: delivery velocity has risen eleven per cent since the reorganisation, cost per unit of output has fallen, the dependency count between teams is down, and the utilisation figure sits at the level the planning cycle targeted. Every number on the screen is correct, and the correctness is the point of what follows. The instrument that produced them has been audited; the data pipeline is sound; nobody in the room is lying to anybody. The committee approves the next phase of the programme on the strength of what it sees, and what it sees is all that it has.

Eleven quarters later the same committee convenes over a different document. A product has shipped with a defect of a kind the organisation had not produced in twenty years; the engineers who would have caught it are enumerated in an exit-interview file; the two teams that once resolved this class of problem across a shared lunch table now resolve it across a ticketing queue whose median latency is measured in weeks; and the recovery estimate, when it finally arrives, is denominated in years. Between the first meeting and the second, no instrument in the organisation registered anything but improvement. The question the paper takes up is whether this sequence is a misfortune that better instrumentation would have prevented, or the signature of something the instrumentation could not, in principle, have carried.

The wager of everything that follows is that the second reading is correct, that it is provable, and that the proof already exists in four separated pieces. Each piece has been established on its own ground; each addresses what looks like a different object; and the labour of this paper is to show that the four are one, to name what they are one statement of, and to draw the consequences that only the unified form makes visible.

1.2 The First Result: The Ledger Is Blind to the Joining

The first result concerns what a per-unit valuation can represent. Let an organisation's performance be written as a value functional over its coalitions, in the standard characteristic-function form of cooperative game theory (von Neumann and Morgenstern, 1944; Shapley, 1953), so that for any subset of members the functional returns the value that subset can realise through its members working in their actual coupled relation. Harsanyi established that any such functional admits a unique decomposition into coalition-level dividends, one for each non-empty coalition, such that the value of any sub-organisation is the sum of the dividends on all coalitions it contains (Harsanyi, 1959, 1963). The dividend on a singleton is the individual contribution; the dividend on any coalition of two or more members is, by construction, a relational quantity, the value that appears only when those members are joined and that belongs to no member alone.

A word on the notation this paper carries from here forward, because its discipline is part of the argument. Every symbol that follows attaches to the instrument, to its form, or to the constraints provable against that form; no symbol attaches to the metabolic substrate itself, whose mode of being is exactly what the notation will be

shown unable to carry, and any formalism that claimed to model the substrate would have conceded the paper's conclusion in its first line. Let N be the set of units, $|N|=n$, and let the value functional be a characteristic function $v: 2^N \rightarrow \mathbb{R}$ with $v(\emptyset)=0$. The Harsanyi dividends are obtained by Möbius inversion over the subset lattice and return the functional by summation:

$$d_v(S) = \sum_{T \subseteq S} (-1)^{|S|-|T|} v(T), \quad v(S) = \sum_{T \subseteq S} d_v(T).$$

A per-unit valuation, the form of every per-person and per-team ledger, assigns each unit a single figure and reads every whole as the sum of the figures across its units. The representation question then has an exact and complete answer. A per-unit valuation represents the functional faithfully if and only if every dividend on every coalition of size two and above vanishes; wherever any relational dividend is non-zero, of either sign, the per-unit form does not undervalue it and does not approximate it, since the dividend appears in no unit's figure and is therefore assigned, under the form, to nobody. The equivalence can be displayed in both directions at once. A per-unit valuation is an additive functional $\hat{v}(S) = \sum_{i \in S} a_i$ for some assignment $a: N \rightarrow \mathbb{R}$, and

$$\hat{v} = v \iff d_v(S) = 0 \text{ for all } S \text{ with } |S| \geq 2,$$

since the dividends of an additive functional vanish above order one by (1.1), and conversely the vanishing of all higher dividends reduces (1.1) to $v(S) = \sum_{i \in S} d_v(\{i\})$, which is additive. The blindness runs in both directions at once: the ledger cannot register the pairing whose joint capacity exceeds the sum of its parts, and it cannot register the coalition that carries drag, because the relational stratum as such lies outside what the form can express.

An objection from within the same theory must be met here, because it will occur to every reader trained on it, and because meeting it sharpens the result. Solution concepts of the Shapley family distribute the whole of the grand coalition's value across individual members, each member receiving, among other components, a share of every dividend on every coalition that contains them (Shapley, 1953); the relational surplus, the objection runs, is therefore not lost to a per-unit account but allocated within one. The objection conflates two operations that the theory itself keeps distinct. Allocation answers the question of who is owed what from a settled total; representation answers the question of which configurations carry which value, and the two come apart in one displayed line. The Shapley value distributes each dividend equally among the members of its coalition, $\varphi_i(v) = \sum_{S \ni i} d_v(S)/|S|$, so that for any coalition $C \subset N$,

$$\sum_{i \in C} \varphi_i(v) = \sum_{S \cap C \neq \emptyset} d_v(S) \frac{|S \cap C|}{|S|} \neq \sum_{S \subseteq C} d_v(S) = v(C)$$

whenever the weighted sum of dividends on coalitions straddling C 's boundary is non-zero: every dividend from outside C leaks into the sum in proportion to its overlap, the weighted terms cancelling only in special cases, and since representation requires equality at every coalition, the allocation furnishes no representation of the functional below the grand one. A per-unit ledger can allocate relational surplus only by destroying its relational location: two organisations with identical allocation vectors can carry disjoint dividend structures, and a reallocation that dissolves harmless redundancy in the first dissolves load-bearing coupling in the second, while the vector reads identically over both. The instrument that authorises reallocation requires representation, the knowledge of what a given dissolution destroys, and this is exactly the knowledge allocation erases in the act of distributing its value. And beneath the conflation lies a precondition the objection cannot pay: computing any allocation of the Shapley family requires the characteristic function over all coalitions as input, which is to say that the objector's ideal ledger presupposes possession of the complete relational knowledge whose absence from the instrument is the very thing under discussion. The objection concedes the premise in the act of being stated.

The same structural verdict arrives by an independent route through information theory, where the partial information decomposition resolves the joint information a group of sources carries about a target into unique, redundant, and synergistic atoms, and where the synergistic atom, the information present only in the joint configuration, is destroyed by marginalising the group to its members (Williams and Beer, 2010). The dividend and the synergy atom are not the same object, and the two decompositions share no axioms beyond arithmetic; their convergence is therefore not one proof in two notations but two independent formalisms returning one

verdict, which strengthens the verdict exactly because the routes are disjoint: whatever exists only in the joining does not survive the per-unit form.

The magnitude of what the form discards deserves a number, since the discard is habitually imagined as a rounding error at the margin of an otherwise adequate account. The coalitions of size two and above number

$$|\{S \subseteq N: |S| \geq 2\}| = 2^n - n - 1,$$

which for a single unit of fourteen people gives sixteen thousand three hundred and sixty-nine, and for an organisation of sixty exceeds ten to the eighteenth power. Every accumulated pairing, every mutual model between specific colleagues, every apprenticeship line and cross-team calibration lives on one of those coalitions; the dashboard that summarises the same organisation carries perhaps fifteen figures. The instrument does not compress this stratum; the instrument's form contains no place where the stratum could appear.

1.3 The Second Result: What the Instrument Cannot Capture, Someone Absorbs

The first result would be an accounting curiosity if the uncaptured stratum simply idled outside the books. The second result establishes that it cannot idle, because the coordination the stratum performs is load-bearing and its load is conserved. Ashby's law states the constraint from which everything follows: a regulator can hold a system stable only if its own variety, the number of distinct states it can assume and deploy, at least matches the variety of the disturbances it must absorb; in the compression that has organised the cybernetic tradition since, only variety can absorb variety (Ashby, 1956, p. 207). The law fixes the total that any arrangement must somewhere absorb, and in Shannon's entropy form it reads

$$H(Z) \geq H(D) - H(R),$$

the residual variety of outcomes Z bounded below by the variety of disturbances D less the variety the regulator R can deploy (Ashby, 1956, ch. 11). The disturbances an organisation's environment presents do not diminish because the organisation's formal apparatus represents fewer of them; the requisite variety is set by the environment, and the apparatus's share of it is set by the apparatus's form.

The consequence is a conservation constraint that any investigation of functioning organisations encounters as soon as it asks where the unrepresented coordination goes. The total variety required for coherence divides into the portion the formal system captures, the metrics, the procedures, the documented interfaces, and the portion that human beings absorb in order to bridge the distance between the representation and the work: the interpretation of what the ticket actually means, the navigation of the interface the specification does not describe, the reconciliation of targets that contradict one another at the point of execution, the shadow work Illich named as the unpaid complement of every formalised economy (Illich, 1981). The sum is fixed; the division is what management initiatives move. Written as the conservation identity, with V_{req} set by the environment, V_F the formally captured share, V_A the absorbed share, and L the count of living absorbers,

$$V_{\text{req}} = V_F + V_A, \quad a = \frac{V_{\text{req}} - V_F}{L}, \quad \delta V_A = -, \delta V_F \text{ at fixed } V_{\text{req}},$$

so that the per-capita absorption a rises under any contraction of V_F or of L against an unchanged environment, and the identity forbids the elimination the initiative reports. A real increase in effective V_F reduces V_A , by the identity itself; the pattern under diagnosis is the different and institutionally common case, in which formalisation expands the archive without increasing the variety the archive can absorb. Every such extension of the formal apparatus, enlarging the representation while leaving the underlying work and the effective V_F unchanged, transfers absorption from the instrument's ledger to the workforce's cognition, and the transfer is paid in the currencies the political economy of attention has documented, in working memory, in allostatic load, in the evenings spent holding what the system could not hold (the load taxonomy in Sweller, 1988; the stress physiology in McEwen, 1998). The dashboard that reports a leaner operation after such an initiative is reporting the visible half of a conserved quantity; the other half has moved to where the instrument does not look, which is inside the people who make the representation come true.

The conservation reading reverses the usual moral of the green dashboard. The manager's retort to any critique of metrics is that the metrics work, the targets are met, the system produces what it was designed to produce; and the retort is, within its own terms, irrefutable, because the measurements do measure something and the coherence does appear. What the conservation constraint adds is the origin of the coherence. The formal system did not produce it; the formal system cannot, by the first result, even represent the stratum in which most of it lives; the coherence was produced by the absorbers and booked by the instrument. The KPI did not manufacture the order it reports. The KPI took the credit for an order manufactured elsewhere, and the elsewhere is a set of human nervous systems whose capacity to keep manufacturing it is finite, depletable, and absent from the books.

1.4 The Third Result: The Instrument Arrives After the Horizon

The first two results are synchronic; the third introduces the dimension along which the damage compounds. Conant and Ashby proved that every good regulator of a system must be a model of that system, in the sense that maximal regulation requires the regulator's internal organisation to stand in a mapping relation to the organisation of what it regulates (Conant and Ashby, 1970). Neither their theorem nor Ashby's law speaks of time, and the temporal extension is where the third result does its work. Let the regulator act on a cycle, the interval at which it measures, evaluates, reallocates, and resets its operative account of the system; let the system possess a replenishment-cycle, the interval over which its decisive state actually rises and falls. Where the action-cycle is shorter than the replenishment-cycle by more than a factor of order one, the regulator cannot, by construction, model the system at the system's own resolution. Within any single window of the regulator's cycle, a change unfolding at the system's cycle presents a variation indistinguishable from the window's ordinary fluctuation; under the rate-distortion allocation that governs any instrument operating under a fixed representational budget, components whose within-window variance falls below the water level receive no rate at all and are reconstructed as constants (Shannon, 1959; Cover and Thomas, 2006, §10.3.3); and the slow variable therefore enters the regulator's account as a steady state exactly while it is falling. The allocation itself can be displayed in two lines. For parallel Gaussian components with within-window variances $\sigma_i^2(W)$ under total rate budget R , reverse water-filling assigns

$$R_i = \max\left(0, \frac{1}{2} \log_2 \frac{\sigma_i^2(W)}{\lambda}\right),$$

with λ the water level set by the budget (Cover and Thomas, 2006, §10.3.3); and a component whose characteristic cycle T_s is long against the window W contributes within-window variance only through its drift, at the scaling

$$\sigma_s^2(W) = \Theta((W/T_s)^2), \sigma_s^2,$$

so that for $W \ll T_s$ the variance falls below any fixed water level, $R_s = 0$, and the component is reconstructed at its mean.

The corollary carries the diagnostic weight of the whole result. The variety the regulator cannot absorb does not disappear; by the conservation constraint it is absorbed within the system itself, producing state changes the regulator does not register; and the regulator receives its first legible signal of a slow depletion only once the accumulated change has crossed a threshold that produces effects at the regulator's own resolution, which is to say only after a span comparable to the system's cycle has already run. Where the slow variable is a stock that took years to accrue and would take years to rebuild, the geometry of the two cycles guarantees that legibility arrives on the far side of reversibility. The instrument certifies the course as sound throughout the interval in which the course could still have been corrected, and delivers its warning at the moment the warning can no longer be acted on within the horizon that matters. The relation deserves to be stated in its plainest form: the belatedness is not a lag that faster reporting would close, since the constraint binds any instrument whose action-cycle stands in the stated ratio to the substrate's cycle, whatever its sampling rate; the belatedness is the ratio itself, worn as a property of the account.

The construction under which the result binds must be stated, because the obvious escape will be proposed and must be met on its own ground. The result assumes an instrument operating under a fixed representational budget; a rate allocation governed by each component's contribution, within the regulator's window, to the fidelity criterion that the action-cycle imposes; and a regulator that acts each cycle from its operative account, the reconstruction available at decision time, rather than from a passive archive consulted at leisure. The escape proposes a regulator that carries memory: longitudinal models, lagged variables, priors over the slow components. Under the stated construction the escape forks and both branches close. Either the longitudinal model is funded by expanding the budget, whereupon the water-line argument recurs at the new budget against the full variety of the organisation's disturbances, which exceeds any finite budget by the conservation constraint of the previous section; or the model is funded by reallocating rate away from the fast components, which is exactly the reallocation the action-cycle's incentive gradient forbids, since the regulator is evaluated, each cycle, on its regulation at the cycle's own scale. A record that is kept but does not enter the operative account at decision time regulates nothing; it is consulted, where it is consulted at all, in the retrospective inquiry that follows the threshold event. And the deepest closure is the one the coming synthesis will make general: the slow variable at issue is the relational stratum, which the first result showed the per-unit form cannot represent at any rate, so the lagged entry the escape proposes would be an entry about a proxy for the stratum rather than about the stratum, and the belatedness result binds hardest exactly where the blindness result binds. The two aspects are joint; Part II will establish that this jointness is the content of the theorem itself, and no accident of the present argument. The honest boundary of the result follows from the same construction: slow variables that are themselves archival in kind, headcount trends, cash positions, the ageing of a backlog, are representable per unit and per window, and longitudinal tracking of them succeeds routinely, which is why they do not produce the pattern and why their tractability is so often mistaken for evidence that the substrate's stratum could be tracked the same way.

1.5 The Fourth Result: The Improvement Can Be the Consumption

The first three results establish blindness, displacement, and belatedness; the fourth establishes that the three together license a dynamic in which the instrument's record of success is produced by the destruction of what the record cannot show. Write the output an organisation delivers in a reporting interval as a function of its visible inputs and its relational stock together: the headcount and hours the ledger measures, and the accumulated coupling, mutual model, and tacit calibration that the first result located in the higher-order dividends and that the ledger cannot carry. A reallocation driven by the per-unit instrument, a cut, a boundary redraw, a utilisation push, reduces the visible input while the demand on output stands; the demand is then met, where it is met, from the remaining term, which is the stock; and measured output per unit of visible input rises for exactly the interval during which the stock is being drawn down faster than it re-forms. The number on the dashboard and the depletion beneath it are one event under two descriptions, and the instrument possesses only the first description. Displayed: with output $F=f(L, K, A)$, $\partial f/\partial K > 0$, a cut $L \rightarrow (1-c)L$ at held output $F=F_0$ financed by drawdown $\dot{K} < 0$ gives a measured productivity ratio

$$\frac{p_{\text{after}}}{p_{\text{before}}} = \frac{F_0/((1-c)L)}{F_0/L} = \frac{1}{1-c} > 1 \quad \text{while } \dot{K} < 0,$$

the reported gain and the consumption standing on the two sides of one identity.

The result is conditional, and its conditional form is what makes it exact rather than polemical. A rise in the measured flow can have innocent causes; a process simplification, a tooling improvement, a favourable shift in demand, or the dissolution of a coalition that carried drag will each raise the figure without touching the stock, and a reallocation can therefore land well. What the result establishes is narrower and sufficient: under the stated conditions, an action-cycle short against the stock's replenishment-cycle and an instrument blind to the stock's stratum, the drawdown case and the improvement case produce identical readings at decision time, and the instrument that authorises the reallocation cannot tell, from anything within its own form, which case it is authorising. In the notation Part II will complete, with Φ the instrument's representation map and τ_{leg} the

legibility span of §1.4, the indistinguishability is the condition

$\Phi(x_{\text{drag}}(t)) = \Phi(x_{\text{drawdown}}(t))$ for all $t \in [t_0, t_0 + \tau_{\text{leg}},]$. Bastiat's oldest correction in political economy names the shape of the error: an accounting that stops at what is seen will mistake destruction for gain, because the visible activity is counted and the destroyed stock that financed it appears nowhere (Bastiat, 1850). The per-unit dashboard commits the error by construction rather than by oversight; it presents the seen as the whole because, by the first result, it has no entry in which the unseen could be recorded, and by the third, it will learn of the loss only after the horizon at which the loss could have been arrested.

1.6 The Question the Four Results Pose Together

Set side by side, the four results exhibit a symmetry too regular to be coincidence. Each takes an instrument of quantified coordination; each identifies a stratum the instrument cannot carry; each shows the uncarried stratum bearing the actual load; and each ends at the same scene, an instrument reporting order while the conditions of order deplete beneath it. The blindness result locates the stratum in the relational dividends; the conservation result shows its load displaced onto living absorbers; the horizon result shows its depletion arriving in the account after the account could act; the cannibalisation result shows the account's improvement drawing on the depletion itself. The results were proved separately, on the mathematics of games, of variety, of information rate, and of production; and their convergence poses the question a synthesis exists to answer. What is it about the instrument, and what is it about the stratum, that makes the four demonstrations come out the same way every time? Part II answers by exhibiting the two modes of being whose meeting the four results describe, and by proving that the four properties follow jointly from the meeting itself.

Part II: The One Result

2.1 Two Modes of Existence for a Coordinating Capacity

Consider first how the stratum the four results keep finding actually comes to exist. Two engineers share a desk across the lifecycle of a product. In the same hours in which they produce the code their commit logs record, they produce something the logs do not record: a shared model of how the codebase behaves under load, a mutual knowledge of where it tends to come apart and what compensates, a calibration to one another's blind spots that lets each catch what the other will miss, a felt map of the user's actual workflow that no requirements document contains. None of this was installed; all of it accrued, and it accrued under conditions that admit no acceleration. The two-hundred person-hours that a hundred engineers spend together in a week and the two-hundred person-hours that two engineers spend together in a year are identical in the ledger and disjoint in what they can deposit, because the deposit in question is laid down only along a shared duration, the way sediment is laid down only along a river's actual run. Polanyi identified the transmission condition for capacities of this kind: an art which cannot be specified in detail cannot be transmitted by prescription, since no prescription for it exists, and it passes only by example, from master to apprentice, by proximity, correction, and time (Polanyi, 1958, p. 53). The condition of transmission and the condition of accrual are the same condition, and I will name the mode of being it describes. A capacity is metabolic when it exists only as the ongoing deposit of lived duration under thermodynamic cost; when its present state is the sediment of a history that cannot be re-run faster than it was lived; when it degrades without continued exercise, as tissue degrades without continued perfusion; and when it moves between bearers only through the shared presence its accrual required in the first place.

Consider next the mode in which the instrument itself exists. A dashboard's content is a set of entries; each entry is retrievable at will, reproducible at negligible cost, indifferent to whether it was written a second or a decade ago, and meaningful in isolation from every other entry, since the form defines the whole as the sum. Nothing in the instrument's mode of being depends on duration; a copied database is the database; a report

regenerated from stored records is the report; the account can be transported, merged, and queried without any of the operations costing what the recorded events cost. I will call this second mode of existence archival. An entity is archival when its content exists as retrievable record rather than as sustained process; when reproduction is decoupled from the cost of what is reproduced; when its elements are per-unit entries composable by summation; and when time enters it only as a stamp on the entries rather than as the medium of its existence.

The two modes are not a contrast of degree, as if the archive were a coarser metabolism or the metabolism a slower archive; they are two answers to the question of how a thing persists. The metabolic persists by being continuously re-enacted, and stops existing when the enactment stops; the archival persists by being stored, and exists identically whether consulted or ignored. The engineers' mutual model dies when the pairing is dissolved, on a decay curve set by the residue of trust and memory; the record of their velocity survives the dissolution untouched, and will report the pairing's historical throughput to any query, forever, with no entry registering that the capacity behind the number has ceased to exist.

2.2 The Archival Regulator Theorem

The distinction permits the four results to be restated as one. Let a regulator be archival in the sense just given, and let the substrate it is set to regulate be metabolic. I state the claim first and derive its parts after.

The Archival Regulator Theorem. An archival regulator set over a metabolic substrate exhibits, jointly and by construction, four properties: it is blind to the substrate's relational stratum (representation); the coordination variety it does not capture is displaced onto living absorbers rather than eliminated (conservation); its account of the substrate lags the substrate's cycle by an interval that places legibility on the far side of reversibility (temporality); and wherever it drives reallocation under the stated conditions, its recorded improvement can be produced by consumption of the substrate it cannot record (dynamics). The four properties are aspects of one fact, the disjunction of the two modes of being, and no refinement of the regulator that preserves its archival form removes any of them.

The derivation of each aspect from the modes runs as follows. Blindness follows from composability: the archival form defines every whole as the sum of per-unit entries, and the metabolic stratum is constituted in the joining of units, which is exactly what a sum cannot express; the Harsanyi decomposition and the partial information decomposition are two independent confirmations of one grammatical fact, that relation is not a property of relata taken severally, and an instrument whose entries attach to relata severally has no place for it. Conservation follows from load: the substrate's coordination is load-bearing against an environment whose variety the instrument does not reduce, so by Ashby's law the variety the archival form fails to capture must be absorbed by whatever can vary at the required rate, and the only entities in the arrangement that can vary at the required rate are the living participants, who therefore receive the displaced load as work. Temporality follows from the decoupling of persistence from enactment: an account whose entries are indifferent to duration is updated at the regulator's action-cycle, while the substrate's decisive state changes at the cycle of lived accrual and decay, and the rate-distortion bound converts the ratio of the two cycles into the guarantee that the slow variable is carried at zero rate, reconstructed as a constant, and reported as steady exactly while it moves. Dynamics follows from the first three jointly: a reallocation that optimises the archival account has no term in its objective for a stratum the account cannot express, meets sustained demand from the stock wherever the visible inputs no longer suffice, and books the drawdown as productivity because the account contains the flow and not the stock; and the belatedness aspect guarantees that the booking stands uncorrected through the whole interval in which correction was available.

The joint character of the four aspects is the theorem's content, and it is what distinguishes the claim from a list of grievances against measurement. Each aspect, taken alone, invites a local repair: add a relational metric against the blindness, hire against the displacement, lengthen the reporting window against the lag, audit the reallocations against the drawdown. The theorem states why the repairs fail as a class. A relational metric is a per-unit entry about relations, and lands inside the same composable form whose grammar the first aspect indicts;

the hire adds absorbers without touching the displacement mechanism, and the new absorbers inherit the load; the lengthened window changes the sampling interval of an account whose mode of being, not whose frequency, is the obstacle; and the audit reads the same ledger the reallocation optimised. Every repair that preserves the archival form redistributes the four aspects among themselves; none removes them; and a repair that abandoned the archival form would no longer be a repair of the instrument but the adoption of a different mode of regulation, which is the subject of Part V.

2.3 What the Theorem Does Not Say

The theorem's boundary must be drawn with the same care as its content, because the claim is strong and a strong claim survives only inside its stated conditions. The theorem does not say that archival instruments are useless; over substrates that are themselves archival or near-archival, inventories, transactions, discrete stocks of fungible goods, queue lengths, the archival form is adequate and often optimal, and the operations research it supports is sound within that domain. The theorem does not say that every reorganisation destroys value; the dynamics aspect is conditional on the dissolved coalitions carrying positive dividends and on dissolution outrunning re-formation, and a reallocation can cut drag, and land well, and be recorded well. The theorem does not say that measurement corrupts what it touches by some occult influence; the mechanism is representational and thermodynamic throughout, and no intention, ideology, or bad faith appears anywhere in the derivation. What the theorem says is narrower and harder: that wherever the substrate is metabolic, the archival instrument set over it cannot tell its own success from the substrate's consumption at the moment of decision, and that this indistinguishability is a property of the meeting of the two modes, available to no participant to waive and to no refinement of the instrument to remove. The committee of the opening scene was not deceived by anyone. The committee was reading an archive about a metabolism, and the theorem describes what such a reading can and cannot contain.

2.4 One Map, One Kernel

The unity the theorem asserts can be stated in a single frame, and stating it removes the last respect in which the four aspects might be suspected of being a well-named list. Let the substrate's states form a space carrying the fine description, in which every coalition of participants has its relational configuration and every capacity has its history of accrual; let the instrument's space consist of per-unit, per-window entries; and let the archival representation map, written here as a map from the first space to the second, be the operation by which the instrument receives the substrate. The map is not chosen by anyone; it is the archival form itself, and by the definition of that form in §2.1 it factors into two operations that always apply, marginalisation to units, since every entry attaches to one unit, and windowed sampling, since every entry attaches to one reporting interval, together with, in general, a third, the aggregation of each cell to the single figure the entry carries, whose kernel is the within-cell distribution, the shape, the variance, the overloaded participant hidden inside a healthy mean; each coarsening factor contributes its own distinctions to the kernel, and every further coarsening an implementation composes onto the map enlarges the kernel rather than shrinking it; refinements that add archival entries may shrink the kernel along archival dimensions, and no such refinement removes the kernel the three operations generate, since every added entry marginalises, windows, and aggregates in its turn. Call the kernel of the map the set of distinctions it destroys, the pairs of substrate states that arrive at the instrument as one state. Displayed: with X the substrate's state space under its fine description, U the set of units and $\mathcal{W}$ the set of reporting windows, the instrument's space is $Y = \mathbb{R}^{U \times \mathcal{W}}$, and the archival representation map factors as

$$\Phi = A \circ W \circ M : X \longrightarrow Y, \quad \ker \Phi = \{, (x, x') : \Phi(x) = \Phi(x'), \},$$

with M the marginalisation to units, W the windowed sampling, and A the aggregation of each cell to its single figure; each factor coarsens the algebra the instrument can distinguish, $\Phi^{-1}(\Sigma_Y) \subseteq \Sigma_X$, and the kernel is monotone under composition, enlarged and never reduced by any further coarsening applied to the image, while

remaining, for the strata the three operations destroy, invariant under every refinement that preserves the archival form. Four propositions then hold about this one object, one per theatre.

In the theatre of representation, every relational configuration of order two and above lies in the kernel of the compositional factor, since marginalisation to units annihilates whatever exists only in the joint configuration; the Harsanyi dividends of order two and above and the synergistic atoms of the information decomposition are two decompositions supported on the same kernel, one in the value algebra and one in the information algebra, distinct functionals on distinct underlying objects whose agreement follows from their shared support: both vanish under per-member marginalisation because both live on exactly the distinctions the map destroys, which locates, and retroactively explains, the convergence of the two disjoint formalisms that §1.2 recorded. In the theatre of regulation, if the coupled arrangement remains coherent while the instrument's channel carries only the map's image, then by the law of requisite variety the regulatory load corresponding to kernel content is carried by some other channel, and the only channels in the arrangement with the requisite variety are the living participants; the conservation aspect is the kernel's load surfacing there. In the theatre of time, the windowed factor is a coarse-graining of the substrate's history, and under the construction of §1.4, fixed budget, within-window allocation, action from the operative account, components slow against the window fall below the allocation water-line and enter the kernel through the second factor; the metabolic stratum, being relational and slow at once, lies in the kernel through the compositional and the temporal factors together, which is the joint binding of §1.4 located rather than argued. In the theatre of optimisation, any objective defined on the instrument's space is constant along the kernel, and the fact is one line of calculus: for $J = g \circ \Phi$ and any direction u with $D\Phi, u = 0$,

$$\langle \nabla J, u \rangle = \langle (D\Phi)^\top \nabla g, u \rangle = \langle \nabla g, D\Phi, u \rangle = 0,$$

so an optimiser of such an objective is unconstrained exactly along the kernel's directions and moves along them wherever its visible gradient leads; where the substrate's value falls along a kernel direction, the stock, the measured improvement and the consumption are one motion, seen through the map and past it, which is (1.10) derived rather than stated.

The theorem's one fact now has an address inside a single frame. The four aspects are four properties of one map and its kernel, encountered in four theatres; the mode distinction of §2.1 is what guarantees the kernel is occupied and load-bearing wherever the substrate is metabolic, since the metabolic is exactly what accrues in the configurations and the durations that the compositional and the temporal factors annihilate; and the full formal development of the frame, the measure-theoretic setting, the regularity conditions, the composition of the factors in general, is stated here as open and bounded work rather than as accomplished, which the propositions above neither require nor pretend.

2.5 The Status of the Theorem

A claim assembled from formal components owes its reader an exact account of where the formality ends, and the account is this. Each of the four component results has a formal home in which it is proved under its own assumptions: the representation result in the dividend decomposition of cooperative game theory, with the allocation distinction of §1.2 marking its exact edge; the conservation result in the law of requisite variety, read against the impossibility of total formal capture; the temporal result in the rate-distortion allocation, under the construction §1.4 specifies, fixed budget, within-window allocation, action from the operative account each cycle; and the dynamics result as a conditional on the three, holding wherever dissolved coalitions carried positive dividends and dissolution outruns re-formation. What the synthesis adds is not a fifth formal result but a derivation of the four from a single source, the disjunction of the archival and metabolic modes, and the derivation is conceptual in the way that the identification of a mechanism is conceptual: it exhibits why the four formal results keep being proved about the same meetings of instrument and substrate, and it predicts where the pattern will and will not appear. The theorem is therefore offered as a theorem about form, whose components are severally proved and whose unity is demonstrated by construction from the mode distinction; readers who require the unity itself in a single frame are directed to §2.4, where the frame is exhibited, one map and its kernel, with each aspect as a proposition about that object, its full formal development marked as open; and the observation stands that the historical dispersal of the four results across four literatures is itself an instance of the paper's subject, since each literature's formalism is an archive of one aspect of a phenomenon that exists whole only in the meeting the paper describes. The scope conditions travel with the claim: the substrate must be metabolic in the defined sense, the instrument archival in the defined sense, the cycle ratio in the stated regime, and outside these conditions the paper asserts nothing, which is what permits it to assert so much inside them.

Part III: One Instrument, Four Costumes

3.1 The Recognition Problem

The theorem earns the title of a general theory only if the archival regulator turns out to be one instrument wearing several costumes, and the contemporary order supplies the costumes in abundance. Four apparatuses organise the coordination of the present economy, and the order's self-understanding treats them as belonging to four different subjects: the delivery dashboard belongs to management science, the price to economics, the prescriptive framework to organisational design, and the large language model to artificial intelligence. The sections that follow read each apparatus against the theorem's four aspects, and the reading returns the same verdict four times, which is the empirical form of the claim that the four subjects have been studying one object without knowing it.

3.2 The Dashboard

The delivery dashboard is the theorem's laboratory case, and Part I has already run it; what remains is to name its substrate and confirm the aspects in order. The substrate is engineering culture in the load-bearing sense: the accumulated tacit knowledge of senior practitioners, the mutual models among specific colleagues, the apprenticeship lines along which the capacity to certify, to review, and to catch reproduces itself across a generation. The substrate's replenishment-cycle is the career; the instrument's action-cycle is the quarter; the ratio sits in the regime the temporal aspect specifies. Blindness holds by the dividend result; displacement holds by the conservation constraint, and its receipts are legible in the burnout literature wherever formal capture has been extended over unchanged work; belatedness holds by the cycle ratio; and the dynamics aspect holds wherever reallocation is driven from the ledger, with the drawdown booked as productivity. The documented record of metric-driven institutions supplies the pattern at scale: the aircraft manufacturer whose executive account improved across two decades while the certification culture beneath it thinned toward the threshold that two crashes finally rendered legible (House Committee on Transportation and Infrastructure, 2020; Joint Authorities Technical Review, 2019; Robison, 2021); the bank whose cross-sell figure climbed toward its target for fifteen years by consuming the customer trust the figure was supposed to proxy (U.S. Securities and Exchange Commission, 2020; Tayan, 2019). In each case the instrument reported order until the substrate's depletion crossed into the instrument's resolution, and in each case the crossing arrived denominated in a currency, lives, licences, settlements, that the account had never carried.

3.3 The Price

The price system is the dashboard written at the scale of an economy, and the tradition that built the strongest case for it supplied, in the same gesture, the conditions the theorem requires. Hayek's argument for the price mechanism rests on knowledge that exists only dispersed among particular participants, the knowledge of the particular circumstances of time and place, never given to any single mind and never capable of being so given (Hayek, 1945); and the argument's force is exactly that the price economises on what cannot be centralised. Read against the present distinction, Hayek's dispersed knowledge is the metabolic substrate of the economy, accrued in particular lives at particular sites, and the price is the archival entry that summarises its local verdicts. The theorem then specifies what the summary can and cannot do. The price carries the relata's severally-expressible valuations and cannot carry the relational stratum, the supplier ecosystem whose redundancy is a property of the network and of no firm within it, the engineering commons whose capacity belongs to a regional web of practice rather than to any priced transaction; the coordination the price does not capture is absorbed by participants, households, and hinterlands whose absorption appears in no market; the price's action-cycle, compressed by the contemporary financial architecture to the quarter and below, stands to the replenishment-cycles of the productive substrates, the decade of a supplier network, the generation of an engineering culture, the century of an ecological system, in ratios that place the account's legibility far beyond the substrates' reversibility; and the reallocation the price drives, capital moving toward the favourable figure, can produce the favourable figure by consuming the substrate, the buyback financed from the research pipeline, the margin financed from the maintenance backlog, the quarter financed from the commons. The conditions under which the price's defenders claimed it as an instrument of coordination were conditions in which its entries still stood coupled to the substrates they summarised; the theorem describes the regime the order enters when the coupling thins, and the empirical record of the financialised period, the hollowed supplier bases, the retired unreplaced engineers, the infrastructure run past its renewal cycle, is that regime's signature written across sectors (Lazonick and O'Sullivan, 2000; Lazonick, Tulum, and Sakinc, 2020; Mazzucato, 2018).

3.4 The Framework

The prescriptive organisational framework, of which the team-pattern literature is the dominant contemporary instance (Skelton and Pais, 2019), presents itself as the corrective to metric-driven management, and the theorem places it instead among the costumes. A framework is a book of configurations: team types, interaction modes, boundary heuristics, a grammar that can be drawn on a chart and applied to any organisation. Its entries are per-unit at the team grain rather than the individual grain, and the change of grain leaves the form intact, since a per-team ledger is a per-unit ledger whose units are teams, and the dividends between teams, the cross-boundary coupling, the shared product knowledge, the practitioners whose work stitches the units together, vanish from it exactly as the dividends between individuals vanished from the finer account. The framework's mode of transmission seals its position under the theorem. What makes a team a team rather than a set of allocated names is metabolic, accrued through shared duration, transmissible by the apprenticeship form and by nothing else; a framework is archival by genre, explicit content in a portable book, and Polanyi's condition states the consequence without needing any addition: the rules of an art can be useful, and they do not determine the practice of the art (Polanyi, 1958, p. 50). The framework can describe the configurations within which metabolic capacity might accrue; it cannot install the capacity, and where its adoption licenses reallocation against per-team entries, it operates as the dashboard operates, with the four aspects intact and the additional irony that the vocabulary of the reallocation is borrowed from a literature written to protect what the reallocation consumes.

3.5 The Model

The large language model completes the set, and its inclusion is the synthesis's furthest reach, since the model is not usually described as a regulator at all. The description fits as soon as the deployment is examined rather than the artefact. The model is installed over exactly the substrate the previous costumes were installed over, the living practice of an organisation's language, judgement, and coordination, and it is installed to do what regulators do, to absorb variety on the organisation's behalf: to answer, summarise, draft, decide, and route in place of the practice that formerly did so. Its mode of being is archival in the strict sense of Part II. Its content exists as retrievable record, the compressed sediment of a corpus whose accumulation cost it nothing to reproduce; its persistence is storage, identical whether exercised or idle; its account of any practice is composed from entries about the practice's past products, severally taken. The four aspects then follow in the same order as before. The relational stratum of a living practice, the mutual attunement of particular speakers, the situated judgement that Polanyi's condition reserves to proximity, does not survive into a record of the practice's outputs, and the model is blind to it in the constitutive rather than the corrigible sense; the coordination the model does not carry is displaced onto the humans who supervise, correct, contextualise, and repair its products, an absorption the deployment's ledger books as productivity gain; the model's account of the practice is fixed at the cycle of its training against a practice that moves at the cycle of living use, a horizon mismatch worn as a property of the artefact; and where the deployment reallocates, where the practice's bearers are reduced against the model's measured throughput, the recorded improvement can be financed by the consumption of the very practice whose sediment the model replays, a stock drawn down that no entry in the deployment's account can register. The model is the archival regulator in its purest costume, an archive so complete that the order has begun to mistake it for the metabolism it records; and the theorem's verdict on it is the same verdict four times confirmed, delivered now on the instrument the order proposes as the successor to all the others.

3.6 The Cascade

Read jointly, the four costumes expose a further structure that no single instance shows. The displaced load and the drawn-down stock do not distribute themselves evenly; they flow, and the theorem constructs two properties of the flow while leaving a third to the record. The conservation aspect constructs the flow's persistence: when any absorber saturates or departs, the load relocates rather than lapses, so that the flow has no terminus short of a substrate able to regenerate faster than it is drawn. The temporal aspect constructs the flow's legibility: among the substrates the flow reaches, depletion crosses the instrument's resolution in the order of their replenishment-cycles, the fastest-cycle substrate first and the slowest last. What the theorem does not construct is the sequence of stations the flow visits; the sequence belongs to the documented record rather than to the derivation, and the record exhibits it with a regularity that invites, without yet possessing, a construction of its own. The first station in the record is the worker, whose cognition and physiology receive the bridging labour and whose departure, when absorption exceeds regeneration, removes a carrier of the stock and increases the absorption load borne by those who remain, as the identity (1.6) requires when L falls. The second station is the organisation's own substrate, the codebase that accumulates the workaround, the procedure that accumulates the exception, the institution that accumulates the informal debt no account carries. The third station is the product, whose margins, redundancies, and integrities are the next reserves the sustained demand draws upon. The fourth station is the customer, whose trust is mined where the product's reserves are spent. The last station lies past the organisation's boundary altogether, in the supplier networks, the public infrastructures, and the ecological systems whose replenishment-cycles are the longest of all. The record's sequence therefore runs, station by station, toward substrates of ever longer cycle, and at this point the derivation resumes on the record's ground: wherever the flow does move outward in this way, the temporal aspect converts the movement into the otherwise puzzling calm of the aggregate account, since each relocation carries the depletion into a stratum whose legibility arrives later than the last, so that the account stays green at every step for reasons the previous station's exhaustion has just supplied, and registers the cascade only at whichever station finally presents a threshold event inside the instrument's frame. Read under the constructed ordering, the documented record condenses to a single figure rather than to a further result: the two crashes, the enforcement action, the supply shock, and the burnout epidemic are one phenomenon surfacing at four stations of the record's sequence, and the instrument that presided over each records each as a surprise.

Part IV: The Closure From Inside

4.1 The Demand for a Better Instrument Is the Signature

Institutions that encounter the theorem's consequences respond, with a regularity that itself demands explanation, by demanding a better instrument. The metrics are not working; the organisation needs better metrics. The model does not capture the culture; the organisation needs a richer model. The dashboard missed the depletion; the dashboard needs leading indicators. The response sounds like learning, and the theorem identifies it as the pattern's continuation by other means, because every candidate improvement that preserves the archival form lands inside the form's grammar. A leading indicator is a per-unit entry with an earlier timestamp; a culture metric is a per-unit entry about relations, which the composability aspect proved cannot express relation; a richer model is a larger archive, whose enlargement changes what the record contains and cannot change what a record is. The displacement from an ontological limit to a methodological deficit is the institution's way of keeping the problem inside the register it knows how to fund, and the consultancy cycle that supplies each successive framework depends on the displacement remaining unexamined, since a diagnosis at the level of the form would end the succession rather than extend it (Power, 1997; Muller, 2018, on the cross-sector regularity of the pattern). The demand for a better archive, made in response to damage the archival form produced, is therefore not evidence that the theorem has been absorbed; it is the theorem's fourth costume change, and its most durable.

4.2 The Gestell Defends Itself

4.2.1 The Phenomenology of Dismissal

The resistance will come. It always comes. And it will feel, to those who offer it, like the most reasonable thing in the world.

The first sensation is a kind of tightening. Not in the critic, but in the room. When someone trained in quantitative methods encounters the claim that organisations cannot be mathematically modelled, something contracts. You can feel it before the words come: a slight withdrawal, a repositioning, the subtle gathering of force that precedes objection.

Then the words: "But surely we can measure *something*." "This sounds like you're saying data doesn't matter." "How would we know if this is working?" "What's the evidence for these claims?"

The questions feel neutral. They feel like due diligence, like the proper scepticism of the scientifically trained mind. The questioner experiences themselves as simply asking for rigour. They are not attacking; they are requesting. The request seems minimal: just show us the evidence, just give us something measurable, just demonstrate that this works.

But attend to the bodily quality of the exchange. The questioner leans forward slightly, or back; the posture is one of demand, however politely framed. There is an implicit threat: satisfy this request or be dismissed. The request is not an invitation to dialogue; it is an ultimatum dressed as curiosity. The felt sense is of a door closing, of permission being withheld, of the conversation being channelled back toward familiar ground.

This is what resistance feels like from the inside: not as hostility but as reasonableness. The critic experiences themselves as the adult in the room, the one asking sensible questions while others float off into abstraction. The sensation is of ground, of solidity, of having one's feet planted in reality while others drift.

And this felt sense of groundedness is precisely the Gestell at work. The ground that feels so solid is the computational grammar that makes certain questions askable and others nonsensical. The "reasonable" demand for measurement is reasonable only within a framework that has already decided what counts as reason. The feeling of solidity is the feeling of operating within unquestioned assumptions, assumptions that are invisible precisely because they constitute the ground on which the questioner stands.

4.2.2 The Philosophical Diagnosis

The resistance to post-cybernetic organisation theory is not incidental but structural. It emerges necessarily from the framework the theory dismantles. Understanding this structure transforms the experience of resistance from obstacle to confirmation.

Heidegger's analysis of *Gestell* identifies not merely a set of beliefs but a mode of revealing, a way in which beings show up for understanding (Heidegger, 1954/1977, pp. 19-23). Within the technological *Gestell*, everything appears as resource awaiting deployment, as standing-reserve (*Bestand*) to be optimised. This revealing is not chosen; it is the horizon within which choices become possible. One does not decide to see the forest as timber resource; the forest *appears* as timber resource before any decision occurs.

The computational *Gestell* that governs organisational thought operates identically. Organisations appear as information-processing systems; workers appear as variety-attenuating mechanisms; outcomes appear as measurable states to be optimised. This appearing is not a theory that could be argued against; it is the condition under which organisational arguments become intelligible. The demand for measurement is not a conclusion derived from premises; it is a presupposition that makes certain conclusions possible while excluding others.

When post-cybernetic theory claims that organisations cannot be mathematically modelled, the claim registers within the computational *Gestell* as either false (better models will succeed where current models fail) or meaningless (what could "organisation" mean if not something capturable in formal terms?). The *Gestell* cannot hear the claim as it is intended because the *Gestell* constitutes the hearing. The words pass through a filter that

transforms them into something manageable: a technical challenge, a temporary limitation, an expression of frustration with current tools.

This is why the resistance feels so reasonable to those who offer it. They are not refusing to understand; they are understanding in the only way their framework permits. The demand for evidence is not bad faith; it is the *Gestell* defending itself through agents who experience themselves as simply being sensible. The framework's most effective defence is that it does not appear as a framework at all but as reality itself, as the way things obviously are, as what any reasonable person would think.

Wittgenstein's analysis of hinge propositions illuminates the structure (Wittgenstein, 1969, pp. 15-18). Certain propositions function not as claims within a language game but as the hinges on which the game turns. They are neither true nor false within the game; they are the conditions that make truth and falsity within the game possible. "Organisations can be measured" functions as a hinge proposition for management discourse. To question it is not to make a move within the game; it is to refuse to play, which registers within the game as incompetence or bad faith.

The resistance to post-cybernetic theory is thus a defence of hinge propositions by agents who cannot perceive them as hinge propositions. The defence is passionate precisely because what is at stake is not a particular claim but the framework that makes claiming possible. To accept that organisations cannot be modelled is to lose the ground on which one has been standing, and no one surrenders ground without a fight.

4.2.3 The Self-Refuting Structure of the Demand

The demand for measurement, applied to a framework demonstrating measurement's limits, contains a self-refuting structure that can be made formally precise.

Let D denote the demand: "Provide measurable evidence that measurement cannot capture organisational reality."

The demand presupposes that valid claims require measurable evidence. Call this presupposition P :

P : Claim C is valid $\Rightarrow \exists$ measurable evidence E supporting C

The claim under consideration, C^* , states:

C^* : Organisational reality exceeds what measurement can capture

If C^* is true, then by definition there exist features of organisational reality that measurement cannot capture. Any evidence E that is measurable captures only what measurement can capture. Therefore:

E is measurable $\Rightarrow E$ does not capture what C^* asserts

The demand D thus requires evidence that, by the nature of what it demands, cannot support the claim in question. The demand is structurally unfulfillable not because the claim is false but because the demand's own presuppositions exclude the possibility of fulfilment.

This is not a paradox in the claim but a paradox in the demand. The self-refutation lies in D , not in C^* . The demand assumes P while demanding evidence for a claim that, if true, demonstrates P 's limitations.

The structure is analogous to demanding a proof, within a formal system, that the formal system is incomplete. Gödel demonstrated that sufficiently powerful formal systems cannot prove their own consistency (Gödel, 1931). The demand "prove within the system that the system has limits" is structurally unfulfillable, but this unfulfillability demonstrates precisely what it cannot prove: the system has limits that the system cannot capture.

Similarly: the demand "provide measurable evidence that measurement has limits" is structurally unfulfillable, but this unfulfillability demonstrates precisely what it cannot demonstrate: measurement has limits that measurement cannot capture.

The formal structure can be expressed as a fixed-point theorem. Let M be the operation of measurement and V be organisational variety. The Conservation Law states:

$$V_{\text{Total}} = V_{\text{Captured}}(M) + V_{\text{Remainder}}(M)$$

where $V_{\text{Remainder}}(M) > 0$ for any measurement operation M . The demand for measurable evidence requires:

$$\exists M : V_{\text{Remainder}}(M) \in \text{Range}(M)$$

But by definition, $V_{\text{Remainder}}(M)$ is precisely what M does not capture:

$$V_{\text{Remainder}}(M) \notin \text{Range}(M)$$

The demand requires what its own terms exclude. The theory demonstrates itself through the structure of demands made against it.

12.4 Transmuting Resistance into Evidence

The self-refuting structure of the demand transforms resistance from obstacle into confirmation. Each demand for measurement, each insistence on evidence, each request for quantification *enacts* the very phenomenon the theory describes: the *Gestell* that cannot perceive its own limits, the computational grammar that cannot question its own categories, the framework that defends itself through agents who experience the defence as reason itself.

This transmutation is not rhetorical judo; it is phenomenological description. The resistance is evidence because the resistance exemplifies what the theory claims: that the computational *Gestell* operates invisibly, that it constitutes rather than merely influences thought, that it defends itself by making alternative frameworks appear unreasonable.

Consider the specific objections and what they reveal:

"This is unfalsifiable." The objection presupposes Popperian demarcation: valid theories must be falsifiable through empirical test (Popper, 1959, pp. 78-92). But the post-cybernetic theory does not claim empirical falsifiability; it claims structural necessity. The cohomological obstruction is not an empirical hypothesis that might be falsified by better data; it is a mathematical theorem about the structure of patching operations. Demanding falsifiability of a structural claim applies criteria from the wrong domain, like demanding that a proof of the Pythagorean theorem be tested by measuring triangles.

"Where's the data?" The objection presupposes that valid claims rest on data, where "data" means quantified observations organised for statistical analysis. But the claim that organisations exceed quantification cannot be supported by quantified observations without self-contradiction. The appropriate evidence is phenomenological: the lived experience of organisational reality that exceeds any metric, the tacit knowledge that cannot be formalised, the contextual judgment that no algorithm captures. This evidence is available to anyone who attends to it; it is not hidden but overlooked because the *Gestell* trains attention elsewhere.

"How would you know if this works?" The objection presupposes that practices must be evaluated by measurable outcomes, that "working" means achieving quantified targets. But the Membrane Architecture does not promise measurable outcomes; it promises liveable organisations. The question "is this working?" must be answered not by metrics but by those who dwell within the organisation: do they flourish? Is the work sustainable? Does the organisation regenerate its conditions of possibility? These questions have answers, but the answers are judgments, not measurements.

"This is just philosophy." The objection presupposes that philosophy is mere speculation, that real knowledge is scientific (meaning: quantified, empirical, falsifiable). But this presupposition is itself philosophical: it is a claim about the nature of knowledge that cannot be established by the methods it privileges. The demand that philosophy justify itself scientifically applies criteria that philosophy must evaluate before accepting. The objection is not a refutation but an assertion of the *Gestell*'s authority.

Each objection, when examined, reveals not a flaw in post-cybernetic theory but a presupposition of the framework the theory dismantles. The objections are symptoms. They are the *Gestell* speaking through those who have not yet seen that they are spoken through.

12.5 The Formal Structure of Framework Defence

The phenomenon of framework defence can be given formal expression through the concept of immune response in belief systems.

Let F denote a framework consisting of:

- A set of hinge propositions $H = \{h_1, h_2, \dots, h_n\}$ that are not questioned within F
- A set of inference rules R that operate on propositions within F
- A set of acceptability criteria A that determine which claims are admissible

A challenge C to framework F is a claim that, if accepted, would require revision of some $h_i \in H$.

The immune response $I_F(C)$ is the set of moves available within F for neutralising C :

$$I_F(C) = \{m : m \text{ is derivable in } F \text{ and } m \text{ blocks acceptance of } C\}$$

For the computational *Gestell* as framework F_G , the hinge propositions include:

h_1 : Valid knowledge is measurable knowledge h_2 : Rigour requires quantification h_3 : Organisation

The challenge C^* (organisations cannot be mathematically modelled) threatens all four hinges. The immune response $I_{F_G}(C^*)$ includes:

m_1 : Demand measurable evidence (invokes h_1) m_2 : Dismiss as unrigorous (invokes h_2) m_3 : Reinforce

The immune response is complete if every challenge can be neutralised:

$$\forall C : I_F(C) \neq \emptyset$$

A framework with complete immune response cannot be challenged from within. External challenge requires operating outside the framework, which the framework's acceptability criteria A will classify as inadmissible (unscientific, unrigorous, mere philosophy).

The post-cybernetic strategy is not to defeat the immune response within F_G (which is impossible by the completeness of the response) but to exhibit the immune response as phenomenon, to make visible what F_G makes invisible: that the "reasonable" objections are not reason evaluating claims but framework defending itself.

This exhibition is itself outside F_G and will therefore be classified by F_G as inadmissible. The classification confirms rather than refutes the theory. The framework that cannot see its own limits will classify the demonstration of those limits as illegitimate. This is predicted; it is not a problem but the phenomenon itself, made visible.

12.6 Living With Resistance

The transmutation of resistance into evidence does not eliminate resistance; it transforms the experience of encountering it. The post-cybernetic practitioner who faces dismissal need not feel defeated or defensive. The dismissal is the *Gestell* revealing itself, becoming visible precisely in its attempt to remain invisible.

This is not triumphalism. The *Gestell* is not an enemy to be defeated but a condition to be acknowledged. Those who dismiss post-cybernetic theory are not villains; they are inhabitants of a framework that shapes what they can perceive. The appropriate response is not argument (which operates within shared framework) but invitation: can you see that you are seeing from within assumptions you have not chosen? Can you feel the ground you stand on as constructed rather than given?

Some will not accept the invitation. The *Gestell* is comfortable; it provides orientation; it makes the world manageable. Leaving it requires tolerating disorientation, uncertainty, the loss of familiar tools. Not everyone is prepared for this, and the preparation cannot be forced. The Membrane Architecture offers practices for those ready to practice; it cannot compel readiness.

The resistance, then, is to be neither fought nor fled but witnessed. Witnessed as phenomenon, as the visible form of the invisible framework, as evidence that confirms by denying, as the *Gestell* defending the only world it can perceive.

4.3 The Constituent Cannot Total the Process

A second closure operates beneath the first and forecloses the imagined position from which even a perfected instrument would read. The regulator the order pictures stands outside the organisation, receiving its state and returning prescriptions; and the position does not exist, for a reason the second-order cybernetic tradition established as a matter of structure rather than of practical shortfall (von Foerster, 2003; Maturana and Varela, 1980). An instrument that measures and acts upon an organisation is coupled to it; the coupling makes the instrument a constituent; the constituent's readings and prescriptions are events inside the process being read, altering relations among the measured, including relations the instrument does not represent; and a complete current model held by a constituent would have to include the effects of its own act of modelling, which are not settled until after the act and which further change the system. There is always a remainder, the instrument's own constitutive effect, unrepresented at the moment of action, and the remainder lands, by the conservation aspect, where every unrepresented load lands. The archival regulator is therefore doubly enclosed: its form cannot carry the substrate's stratum, and its position cannot supply the exterior from which a form without that limit could operate. The two enclosures explain why the theorem admits no engineering escape, and they hand the argument to its final question, which is what regulation of a metabolic substrate could look like if it proceeded from inside the substrate's own mode.

Part V: Regulation in the Metabolic Mode

5.1 What the Theorem Leaves Open

The theorem closes a class of instruments, and in closing it, draws the outline of what it does not close. The four aspects follow from the archival form; a mode of regulation that did not hold its content as per-unit record, that persisted by enactment rather than storage, that moved at the substrate's own cycle because it lived inside the substrate's own duration, would inherit none of the derivations. Such a mode exists, has always existed, and has been carrying the load throughout, in the shadow of the instruments that booked its work. It is the mode Polanyi described as the transmission of an art: proximity, example, correction, shared duration, the senior practitioner whose presence in the review is itself the regulatory act, the pairing whose mutual model is the sensor no dashboard replaces, the working arrangement that holds the meaning of the work across its layers because the arrangement's members hold it together in time. Regulation in this mode is metabolic regulation, capacity regulating capacity within one mode of being, and its variety matches the substrate's variety for the oldest cybernetic reason: it is made of the same living material as what it regulates, and accrues at the same cost.

5.2 Conditions Rather Than Prescriptions

The constructive direction cannot take the form its own analysis forbids, and the discipline of Part V is to respect the boundary the theorem draws through the paper itself. A list of prescribed configurations for metabolic regulation would be an archive of the metabolic, a book about the way offered in place of the way, and the genre analysis of Part III applies to it without exemption. What the analysis permits is the specification of conditions: the circumstances under which metabolic capacity can accrue, which an organisation can protect without pretending to install what only duration deposits. The conditions are legible in every case the record supplied. Shared duration among specific people, protected against the reallocation cycle long enough for the slow variable to accrue; proximity between the generations of a practice, so that the apprenticeship line along which the capacity reproduces is a fact of the floor plan and the calendar rather than a value in a slide; the seniority of practice over instrument at the points where the work is judged, so that the account informs the judgement of those who can see the substrate rather than replacing it; and restraint in the one act the theorem proves the instrument cannot price, the dissolution of standing coalitions, undertaken, where it must be undertaken, in the knowledge that the account authorising it is blind to what it dissolves and late with respect to what it destroys. None of these conditions is a metric, none can be verified from the ledger, and each is visible at once to anyone standing inside the practice, which is the theorem's point restated as a criterion: the health of a metabolic substrate is legible in the metabolic mode and only there.

5.3 The Finding

The paper set out from four results that had accumulated separately: a per-unit valuation is blind to every relational dividend; the variety a formal system does not capture is displaced onto living absorbers; a regulator whose cycle is short against its substrate's cycle learns of depletion after reversibility; and a reallocation driven from the blind account can finance its reported improvement from the stock it cannot see. The synthesis has shown the four to be aspects of one theorem about the meeting of two modes of being, the archival and the metabolic; has read the dashboard, the price, the framework, and the language model as four costumes of the one instrument the theorem describes; has traced the cascade by distinguishing what the theorem constructs, the persistence of the displaced load and the cycle-ordered legibility of its depletion, from the station sequence the documented record supplies; has established the double closure that forecloses repair from inside the archival form; and has stated the conditions under which the other mode of regulation, the one that was doing the work all along, can be protected where it lives. The committee of the opening scene can now be answered in the theorem's terms. The sequence it presided over was not a misfortune of instrumentation, and no refinement of the dashboard would have carried the warning, because the warning belonged to a stratum the account's form excludes and arrived on a cycle the account's position cannot occupy. What the committee lacked was not a better archive. What the committee lacked was standing inside the metabolism it was reading about, and that standing is not a document; it is a way of remaining with the work, transmitted by the only means such things have ever been transmitted, and spent, wherever the instrument is allowed to decide alone, faster than any account will ever say.

References

Ashby, W. R. (1956). *An Introduction to Cybernetics*. London: Chapman & Hall.

Bastiat, F. (1850). *Ce qu'on voit et ce qu'on ne voit pas, ou l'économie politique en une leçon*. Paris: Guillaumin. English translation: "That Which Is Seen, and That Which Is Not Seen," in *Selected Essays on Political Economy*, ed. G. B. de Huszar, trans. S. Cain. Irvington-on-Hudson, NY: Foundation for Economic Education, 1995.

Conant, R. C. and Ashby, W. R. (1970). "Every Good Regulator of a System Must Be a Model of That System." *International Journal of Systems Science* 1(2): 89-97.

- Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. 2nd ed. Hoboken, NJ: Wiley-Interscience.
- Harsanyi, J. C. (1959). "A Bargaining Model for Cooperative n-Person Games." In A. W. Tucker and R. D. Luce, eds., *Contributions to the Theory of Games IV*. Annals of Mathematics Studies, No. 40. Princeton: Princeton University Press, pp. 325-355.
- Harsanyi, J. C. (1963). "A Simplified Bargaining Model for the n-Person Cooperative Game." *International Economic Review* 4(2): 194-220.
- Hayek, F. A. (1945). "The Use of Knowledge in Society." *American Economic Review* 35(4): 519-530.
- House Committee on Transportation and Infrastructure (2020). *The Design, Development & Certification of the Boeing 737 MAX. Final Committee Report*. Washington, DC: U.S. House of Representatives. September 16, 2020.
- Illich, I. (1981). *Shadow Work*. Boston: Marion Boyars.
- Joint Authorities Technical Review (2019). *Boeing 737 MAX Flight Control System: Observations, Findings, and Recommendations*. Washington, DC: Federal Aviation Administration.
- Lazonick, W. and O'Sullivan, M. (2000). "Maximizing Shareholder Value: A New Ideology for Corporate Governance." *Economy and Society* 29(1): 13-35.
- Lazonick, W., Tulum, Ö., and Sakinc, M. E. (2020). "Why Stock Buybacks Are Dangerous for the Economy." *Institute for New Economic Thinking Working Paper Series* No. 145.
- Maturana, H. R. and Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: D. Reidel.
- Mazzucato, M. (2018). *The Value of Everything: Making and Taking in the Global Economy*. London: Allen Lane.
- McEwen, B. S. (1998). "Protective and Damaging Effects of Stress Mediators." *New England Journal of Medicine* 338(3): 171-179.
- Muller, J. Z. (2018). *The Tyranny of Metrics*. Princeton: Princeton University Press.
- Polanyi, M. (1958). *Personal Knowledge: Towards a Post-Critical Philosophy*. London: Routledge & Kegan Paul.
- Polanyi, M. (1966). *The Tacit Dimension*. Garden City, NY: Doubleday.
- Power, M. (1997). *The Audit Society: Rituals of Verification*. Oxford: Oxford University Press.
- Robison, P. (2021). *Flying Blind: The 737 MAX Tragedy and the Fall of Boeing*. New York: Doubleday.
- Shannon, C. E. (1959). "Coding Theorems for a Discrete Source with a Fidelity Criterion." *IRE National Convention Record* 7(4): 142-163.
- Shapley, L. S. (1953). "A Value for n-Person Games." In H. W. Kuhn and A. W. Tucker, eds., *Contributions to the Theory of Games II*. Annals of Mathematics Studies, No. 28. Princeton: Princeton University Press, pp. 307-317.
- Skelton, M. and Pais, M. (2019). *Team Topologies: Organizing Business and Technology Teams for Fast Flow*. Portland, OR: IT Revolution Press.
- Sweller, J. (1988). "Cognitive Load During Problem Solving: Effects on Learning." *Cognitive Science* 12(2): 257-285.
- Tayan, B. (2019). "The Wells Fargo Cross-Selling Scandal." Stanford Closer Look No. 62. Stanford, CA: Stanford Graduate School of Business, Corporate Governance Research Initiative.
- U.S. Securities and Exchange Commission (2020). *Wells Fargo & Company/MN, Form 8-K (Current Report), February 21, 2020*, including Deferred Prosecution Agreement and Civil Settlement Agreement.

Available at SEC EDGAR database.

von Foerster, H. (2003). *Understanding Understanding: Essays on Cybernetics and Cognition*. New York: Springer-Verlag.

von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.

Williams, P. L. and Beer, R. D. (2010). "Nonnegative Decomposition of Multivariate Information." arXiv:1004.2515.

Three Walls

A Formal Limit on the Computational Paradigm, Its World Models, and Their Deployment

The purpose of a system is what it does.

— *Stafford Beer*

Abstract

The systems the current AI industry sells are sold under three promises: that they understand the world, that they are safe enough to be deployed at scale, and that they learn how the world works from watching it. This paper argues that each promise meets a limit the industry's own resources cannot remove, that the three limits are independent of each other, and that the systems now in deployment meet all three limits at once. The first limit says that a system that has only seen the world through a fixed body of data cannot vouch for the cases that lie outside that data, and adding more data or more compute does not close the gap. The second limit says that a system whose probability of a serious failure on each action does not fall in response to its own serious failures will, deployed long enough, produce them without end, no matter how small the per-action probability is. The third limit says that an agent acting on a system that recovers slowly, an institution, a supply chain, an ecology, cannot tell a healthy state from a depleting one when the agent only ever sees the short run, and ends up spending what it cannot see. The three are established from sources already canonical in the relevant fields and the limits are then turned back on the field's own descriptions. The world-model programme is the central case, because two of the three limits strike it independently. For each limit the paper states the condition under which the limit would be false.

Keywords: grounding, world models, computational paradigm, requisite variety, Borel-Cantelli, foreclosure, AI safety, reference, identifiability

Part I. The Claim and the Frame

1. Three walls, one structure

This section says what the paper is going to do. The argument is built around three claims about the systems the current AI industry has put into the field. Each claim is a wall, in the sense that it sets a limit beyond which the systems cannot pass on the resources their builders have given them. The three walls are different from each other: each is about a different aspect of how the systems work, and no repair to one of them reaches the others. The section names the three walls in order, says what each one is going to claim, and fixes the small set of terms the rest of the paper uses to talk about them. The proofs come later, in Parts II and III. The point of this section is to put the structure in view.

A language model produces a sentence; a world model trained on video produces a prediction of the next frame; a proof procedure decides a theorem; a control system holds a process within a tolerance band. Each produces outputs over a range of cases for which it is held responsible, and each operates from a base that covers only a portion of that range. The paper is the investigation of the gap between the two, and the argument is laid out as a numbered structure to which the rest of the paper returns at every step. Three terms are fixed before any claim is made.

A system has a *domain*, the full range of cases for which it is held to answer. It has a *substrate*, the mechanism that produces its outputs: a trained network, a procedure of derivation, a body of measurements, a regulator's control law. It has a *cut*, the portion of the domain to which the substrate has operational access: the training distribution, the axioms and rules, the regime of measurement, the disturbances the regulator's loop can sample within its cycle. The three are not theoretical conveniences; they are the parts the field's own descriptions already

presuppose whenever it speaks of training data, model behaviour, and deployment context. The argument operates entirely on these terms.

To these three the paper adds a fourth: the *claim*, Λ , the assertion made on behalf of the system that its substrate operating on its cut suffices to determine correct accounts for elements of its domain, including elements outside the cut. The claim is not an act of the substrate but a commitment of those who build and deploy the system. A model trained on curated video is presented as a model of physical dynamics; a score on a fixed evaluation is presented as competence on open-ended work; a regulator tuned on a fast signal is presented as a governor of the slow process it sits inside. In each case a result about the cut is sold as a result about the domain, and the paper's investigation is the investigation of how that distance is to be closed and what it would take.

The argument proceeds along three axes; each is a different property of the same arrangement, and each is the variable that the corresponding wall governs. The first axis is the relation between cut and domain at a single instant: what the substrate can reach. The second axis is the temporal behaviour of the substrate's failure mode: whether catastrophe on a mode deforms that mode's future probability. The third axis is the relation between the agent's evaluation horizon and the cycle on which its target system varies: whether the agent's feedback span can resolve the distinction between a sustainable state of the target and a depleting one. The three axes are distinct, and the paper's claim of independence among the three walls is the formal consequence of that distinctness; Part IV proves the independence, and the no-exit result follows from it.

The first wall is the wall of foreclosure. In plain terms: a system that has only seen the world through a fixed body of data cannot certify what is correct on cases that lie outside that data, because two different ways the world might extend past the data are equally consistent with what the system has been trained on, and the system has no resource for choosing between them. Adding more data shifts the gap but does not close it. The formal version of the claim turns on what counts as "outside the data" and what counts as the system "certifying" a correct account, and the proof is the underdetermination argument familiar from Quine (1960), Putnam (1980), and Kripke (1982): two completions of the world that agree on the data produce identical system behaviour, and no operation of the system can distinguish their differing accounts of a case outside the data. The wall stands at a single instant; it says nothing about deployment or repetition; it says only that the part the system can reach does not contain the part the claim is made about. Part IV will show, drawing on Lawvere (1969) and Yanofsky (2003), that the three lenses through which the first wall is proved are instances of a single categorical structure.

The second wall is the wall of compounding hazard. In plain terms: if the probability that a system produces a serious failure on a given action stays the same no matter how many times that same failure has already occurred, then deploying the system for long enough will produce that failure without end, however small the probability is. A system whose probability collapses sharply after each realised failure, so that the failure suppresses itself, produces only a few. The "ninety-nine percent safe" figure the industry reports is a per-action probability; the working lifetime of a deployed system is many millions of actions; the per-action figure and the deployed total behave oppositely under repetition. The formal version of the claim is the pair of Borel-Cantelli lemmas, with the divergent regime closed by the conditional extension due to Lévy (1937) that drops the independence assumption the classical second lemma requires (Williams, 1991, §12.15). The wall is temporal; the variable it governs is not the size of the hazard but its law of motion.

The third wall is the wall of horizon-mismatch. In plain terms: an agent that acts on something that takes a long time to recover, an institution, a supply chain, an ecology, will deplete what it acts on if it only ever sees the short run. The healthy state and the depleting state of the slow thing look the same inside the window the agent can evaluate; the difference between them only shows up across the cycle on which the slow thing recovers, and the agent's window does not reach that cycle. The agent acts on what it sees, sees only the short cycle, and spends the slow thing without noticing until it has been spent. The formal version of the claim joins finite-horizon observability (Chen, 1999) to the Conant-Ashby theorem (Conant and Ashby, 1970): states the agent's feedback cannot distinguish inside its window do not enter the agent's model, and a regulator whose model lacks the distinction cannot, by the Conant-Ashby result, hold the system in the sustainable state (Ashby, 1956; Conant and Ashby, 1970). The variable here is not the rate dynamics of the second wall; it is the relation between the

agent's window and the slow thing's recovery cycle. A system can have a perfectly cost-collapsing hazard and still sit on the wrong side of this wall.

The independence of the three is the structural claim that carries the rest of the paper. Each wall governs a different variable; each variable can be set independently of the others; current systems fail on all three at once. The consequence for repair is sharp: alignment and patching address, at most, the second variable, and leave the first and third where they were; enlargement of the training distribution addresses, at most, the first variable, and leaves the second and third; extension of the evaluation horizon addresses, at most, the third, and leaves the first and second. The toolkit the paradigm has assembled reaches one variable at a time. The result is the no-exit result, proved formally in Part IV: a system that is foreclosing, cost-stationary, and horizon-mismatched, which is what the deployed systems are, requires three independent repairs, and no single repair the field has proposed reaches more than one of the three.

The paper proceeds along seventeen sections. Part II establishes the first wall and its three lenses: the semantic underdetermination of reference (Frege, 1892; Putnam, 1980; Kripke, 1982; Sellars, 1956; Searle, 1980; Carnap, 1928; Quine, 1960), the statistical underdetermination of the law over a training distribution, and the formal lens through which the field's own positive results assume the grounding they were taken to produce (Ashby, 1956; Conant and Ashby, 1970). Part III establishes the two temporal walls and the distinction between them (Williams, 1991; Chen, 1999). Part IV proves the independence of the three walls and the identity behind the first wall's formal instances (Lawvere, 1969; Yanofsky, 2003), and marks the boundaries the argument does not cross. Part V supplies the genealogy that accounts for the field's blindness to the limits from inside (Shannon, 1948; Newell and Simon, 1976; Fodor, 1975; Dreyfus, 1972), the convergence of the failure pattern across architectures that supports the paradigm-level cause, and the deployment consequence that follows from the three walls together. Part VI states the escape conditions and the falsifiers for each wall, marks what the paper does not claim, and closes.

2. What would have to be true for the paradigm to deliver

This section sets up what the three walls are going to falsify. The paradigm's three promises are not assertions out of nothing; each rests on a condition about how the systems work. If those conditions held, the promises would follow. The argument of the paper is that the conditions do not hold. To make the argument testable, the section states each condition cleanly first, in the form that would have to be true for the promise to be delivered. The rest of the paper falsifies each condition in turn.

The argument is sharpened by stating, before the walls are raised, what the paradigm's three descriptions would require in order to hold. The requirements are the negations of the three theorems, and the paper's task is the falsification of these antecedents one by one. The conditions are laid out as a numbered set; the rest of the paper falsifies each in turn.

The first condition: the correct account of every consequential element of the domain is fixed by the substrate's behaviour on the cut. There would have to be no consequential element of the domain on which two cut-indistinguishable extensions of the world disagree. If reference, the rule, the law governing an unobserved configuration were all determined by the statistics of the observed sample, a system that captured those statistics would have captured the relation; the understanding claim requires that grounding be a property internal to the data. The classical results on the underdetermination of reference (Frege, 1892; Quine, 1960), of rule by record (Kripke, 1982), of theory by formal constraint (Putnam, 1980), of meaning by trace (Sellars, 1956), and of semantic relation by symbol manipulation (Searle, 1980) all bear on this antecedent and are taken up in Part II.

The second condition: the per-action probability of catastrophe is either zero, which the architecture of probabilistic generation forbids by construction, or its dynamics are such that the mode's hazard probabilities sum in deployment. The system would have to learn from each realised catastrophe in the indexed and immediate way that drives the rate on the mode toward zero. The safety claim at scale requires that the failure distribution be deformed by its own catastrophic outputs in the right way. The Borel-Cantelli dichotomy (Williams, 1991,

§12.15), with the Lévy (1937) conditional extension for the dependent case, is the formal apparatus that distinguishes the two regimes and is taken up in Part III.

The third condition: where the system is an agent coupled to a slow-replenishing system, the agent's evaluation horizon reaches the cycle on which the slow variable becomes legible, so that the sustainable and depleting states are distinguishable inside the horizon. The world-learning-and-governance claim requires that the evaluation horizon be matched to the replenishment cycle of the systems on which the agent is to act. Finite-horizon observability (Chen, 1999) together with the Conant-Ashby theorem (Conant and Ashby, 1970) and the Law of Requisite Variety (Ashby, 1956) supply the formal apparatus that fixes when the antecedent fails and is taken up in Part III.

The three conditions are the antecedents the paper falsifies. The conjunction fails in three independent places: grounding is shown not to be internal to the data; the failure distributions of the deployed systems are shown not to be deformed by realised catastrophes on the consequential modes; and the agentic systems are shown to evaluate the consequences of their actions on horizons that fall short of the cycles on which the systems they act upon replenish. What remains, after the three antecedents are denied, is a system that cannot reach what it claims, cannot survive its own deployment, and cannot resolve the slow systems it is held to keep sustainable, sold as a system that does all three.

Part II. The First Wall: Foreclosure

3. The foreclosure theorem

This section states the first wall as a theorem and proves it. The plain content of the theorem is that a system trained on a fixed body of data cannot, by training alone, certify what is correct on cases outside that data when two equally consistent extensions of the world disagree on those cases. The proof is short: if two extensions of the world agree on the data, the system trained on the data cannot tell them apart, and whatever account it gives of the case outside the data is the account it would have given under either extension. The work of the section is to state the theorem cleanly, prove it in three lines, and identify the conditions under which the gap matters in practice. The three lenses of §§4-6 then show that the conditions are met by the systems the paper treats.

The first wall is stated in a single proposition and proved by underdetermination. The terms are those fixed in §1: a domain D , a substrate S , a cut $C \subseteq D$ to which S has operational access, and a claim Λ asserting that the action of S on C suffices to determine correct accounts for elements of D . To these the proof requires one further notion. An element $d \in D$ is *consequential under Λ* when the claim explicitly extends to d , that is, when Λ is taken to license action or assertion at d ; and a consequential element is *off-cut* when $d \notin C$. The foreclosure theorem then takes the following form.

Theorem 1 (Foreclosure). *Let (D, S, C, Λ) be a system in the sense of §1, with $C \subseteq D$ and Λ unrestricted to C . If there is a consequential off-cut element d on which there exist two extensions of the world to D , both consistent with the substrate's behaviour on C and both compatible with the conditions Λ imposes, that yield distinct correct accounts of d , then no operation of S on C alone certifies the correct account of d under Λ .*

The proof proceeds in three lines. By hypothesis the two extensions are indistinguishable to the substrate, since S sees only C and the extensions agree on C ; hence the substrate's behaviour is identical under either; hence any account S produces for d is produced equally under either extension; hence S alone cannot select between the differing correct accounts; the consequential off-cut element d therefore stands without certification by S . The structure of the argument is the underdetermination structure familiar from Quine (1960), Putnam (1980), and Kripke (1982): a fixed sample is consistent with a multiplicity of extensions to the unsampled region, and the sample's behaviour does not discriminate among them.

The theorem is unsurprising in form and inescapable in scope; the cut does not contain its own off-cut completion. What it states formally is the structural condition under which the gap matters: there must exist an off-cut element on which two cut-indistinguishable extensions disagree in their correct account, and that element must be consequential under the claim Λ . Both conditions are met by the systems the paper treats. The off-cut elements are the unprecedented configurations that constitute the cases of consequence; the disagreement is the disagreement among lawful continuations of the data the system has not seen; the claim Λ is the unrestricted competence the paradigm advertises. The paper now establishes these conditions through three lenses: the semantic lens of reference, the statistical lens of law, and the formal lens of the conditions of identifiability. The lenses converge on a single structural finding, and §10 will demonstrate, drawing on Lawvere (1969) and Yanofsky (2003), that the convergence is identity rather than analogy.

4. The semantic lens: reference is not in the cut

4.1 The first lens introduced

This section runs the foreclosure theorem through the first of three lenses, the one philosophy of language has been working out for over a century. The point of the lens is to show that the gap the foreclosure theorem identifies between training data and the world is the same gap analytic philosophy has been calling the underdetermination of reference. The argument is constructed entirely from the field's own canonical sources, on its own terms; the lens does not import an outside standard. The reader who knows Frege, Putnam, Kripke, Sellars, Quine, and the rest will recognise the moves; the reader who does not will be walked through them. The case-by-case construction shows that the same structure appears at five canonical points in the philosophy of language and one engineering case, which is what the lens needs to establish.

The first lens passes the foreclosure structure through the semantic theory of reference; the resources used are the canonical resources of analytic philosophy of language, and the line of argument is the line by which that tradition itself learned to formulate the underdetermination thesis. The discipline of immanent critique requires that the field's own classical sources do the work; the canon is the warrant. The lens proceeds case by case through Frege (1892), Putnam (1980), Kripke (1982), Tarski (1936), Gödel (1931), Sellars (1956), Searle (1980), Carnap (1928), and Quine (1960): reference, the rule, the inferential role, the formal limits of definability, the indeterminacy of translation. The lens closes in §4.10 with a contrast taken from the engineering field's own working categories.

4.2 Frege: reference is not given with structure

Frege's distinction between sense and reference (Frege, 1892, "Über Sinn und Bedeutung") is canonical for a reason the present argument requires. The morning star and the evening star are presented under distinct senses while sharing a single reference, the planet Venus; the identity of reference is not given to the language user by the structural composition of either sense; the identification is an achievement of astronomy, not of grammar; and the achievement was made not by closer inspection of the senses but by extending the perceptual cut to include the planet itself across its orbital phases. The structural lesson the paper takes from Frege (1892) is exact: that the reference of a term is not read off the structure of the expression in which the term occurs; reference is fixed by the world the expression is taken to be about, and the world is not in the expression.

4.3 Putnam: permutation does not perturb the sentences

Putnam's model-theoretic argument (Putnam, 1980, "Models and Reality", *Journal of Symbolic Logic* 45(3)) sharpens the point. A theory's formal constraints, even when taken to be the full set of constraints the theory affords, do not fix the reference of its terms; the theory admits non-standard interpretations on which the truth-values of all sentences are preserved while the terms are reassigned to wholly different objects in a permuted universe (Putnam, 1980). The constraints internal to the theory are not constraints on the reference of the terms; the link from term to object is supplied from outside the theory, by the world in which the theory is applied. The argument is the underdetermination structure of Theorem 1 stated for the formal apparatus of theories: a fixed body of sentences is consistent with multiple reference-assignments, and no operation on the sentences alone selects among them. The two extensions of the world that Theorem 1 requires are, in Putnam's formulation, the standard and permuted models of the theory.

4.4 Kripke: the rule is not in the record

Kripke's reading of Wittgenstein on rule-following (Kripke, 1982, *Wittgenstein on Rules and Private Language*) establishes the temporal-extension form of the same structure. A finite record of applications of a rule is consistent with infinitely many rules that agree on the record and diverge beyond it; the famous case (Kripke, 1982) is the rule *quus*, which agrees with the addition function on every pair the speaker has so far computed and returns five on every pair larger than the largest the speaker has yet considered. The record does not contain the rule it is a record of; the extension of the rule beyond the record is undetermined by the record. The form is again the foreclosure form: a finite cut, a domain that exceeds the cut, and a claim of competence on the domain that the cut underdetermines.

4.5 Tarski and Gödel: the structural limit from inside

The formal closure of the lens is supplied by results internal to logic itself. Tarski (1936, "Der Wahrheitsbegriff in den formalisierten Sprachen") establishes that the truth predicate for a sufficiently rich formal language is not definable within that language; the language cannot internalise its own semantics; the relation of the sentences to the world they describe is not a relation the sentences themselves can capture. Gödel (1931, "Über formal unentscheidbare Sätze...") establishes that a consistent, sufficiently strong formal system contains true sentences it cannot prove; the apparatus does not exhaust the truths in its own domain. These are not external critiques; they are theorems of the discipline that the field's own formalisms inherit. The lesson the paper takes from them is the same lesson the previous cases have delivered: the cut, however formal and however complete in its own terms, does not contain what the claim made on its behalf requires it to contain. Part IV will return to Tarski (1936) and Gödel (1931) as instances of a single categorical structure (Lawvere, 1969; Yanofsky, 2003); here they establish only that the semantic reach a grounded system would need is not securable from within a formal substrate. The argument has now encountered, for a third time inside the semantic lens, the same form of insufficiency. The recurrence is the lens's confirmation that the structure of Theorem 1 is the structure of reference itself, not an analogy.

4.6 Sellars: meaning lives in the space of reasons, not in the trace

Sellars's argument against the Myth of the Given (Sellars, 1956, "Empiricism and the Philosophy of Mind") supplies the further specification the lens requires. Meaning is not a property of an isolated item, sensory or symbolic; it lies in the normative inferential relations the item bears to other items, and these relations are constituted in a practice that determines what counts as a reason for what (Sellars, 1956). A linguistic substrate that produces tokens whose statistical distributions match those of a competent speaker captures the surface trace of that practice, the regularity of inferential transitions as they appear in text; it does not, by capturing the trace, participate in the practice itself, which is constituted by what counts as a reason for what, not by the statistical pattern of what follows what.

The sharpening this argument requires is the structural form of the Sellarsian point against the consciousness-dependent reading. The lens does not turn on whether the system has phenomenal experience; that question, however settled, is not what the argument needs. The argument needs only that the normative inferential practice in which meaning is constituted is not internal to the trace the practice leaves in the substrate's training data; the practice is the world relation the cut does not contain, and the substrate operates on the cut. The trace is the cut; the practice is the off-cut remainder; and the failure of the trace to constitute the practice is the form of Theorem 1 once more.

4.7 Searle: the formal apparatus does not constitute the relation

Searle's argument (Searle, 1980, "Minds, Brains, and Programs") is read here in its structural form, set apart from the consciousness-intuition reading the original argument made famous. The structural form is the form the foreclosure theorem requires: a system that operates by purely formal manipulation of symbols, with no causal connection to the objects the symbols are taken to denote, has nothing in its operation that constitutes the semantic relation between symbol and object (Searle, 1980). The semantic relation is constituted by the relation between the formal apparatus and the world it is applied to, and the system, by hypothesis, lacks that relation. The argument is the foreclosure argument: the apparatus is the cut, the world is the domain, and the relation from cut to domain is not internal to the cut. Whether the system has experience, what experience would be, what role experience plays, all of these are downstream of the structural point, and none is required to make it. The paper's use of Searle (1980) is therefore deliberately re-grounded in logic rather than in the intuition-pump form for which the original argument is best known, and the structural reading is defensible against the original on the terms the original itself supplies.

4.8 Carnap and Quine: the world is required, the cut is not enough

The *Aufbau* (Carnap, 1928, *Der logische Aufbau der Welt*) attempts the constructive form of the opposing position: the world is to be constructed from sense data and logical operations, that is, from a cut consisting of immediate sensory tokens and a set of operations defined over them. The project's failure to deliver the world is acknowledged by Carnap (1928) himself and amplified by Quine's later analysis (Quine, 1960, *Word and Object*). The gavagai example (Quine, 1960) fixes the lesson: pointing at a rabbit while uttering the term is consistent with reference to the whole animal, to the temporal slice, to the undetached parts, to the colour, to the act of attention; the pointing does not select among these, and no extension of pointing closes the gap; the reference is not determined by the perceptual cut, however refined.

The two cases close the lens at its strongest point. Carnap (1928) failure is from inside the most ambitious form of the cut-suffices position; Quine's (1960) analysis is the diagnostic that explains why the failure is structural rather than incidental. The substrate that operates on a cut consisting of tokens, however richly structured, does not by that operation acquire the world-relation the claim made on its behalf requires; the underdetermination of reference by sample is the underdetermination of the world by the cut, and the foreclosure theorem is its formal statement.

The lens has now run through the canon's central cases: Frege (1892) on sense and reference; Putnam (1980) on permutation; Kripke (1982) on rule-following; Tarski (1936) and Gödel (1931) on the formal limits from inside; Sellars (1956) on the normative space; Searle (1980) on the formal apparatus and its world-relation; Carnap (1928) and Quine (1960) on construction and indeterminacy. Each is a case of the same structure; each makes the structural point through different resources; each is a theorem or argument internal to the canon of philosophy of language. The recurrence is what the lens needs; the canon has been encountering the foreclosure structure for a century, and the canon's own resources are what have been forcing the encounter.

4.9 The figure of camera and carving

To anchor the lens before the next one is taken up, the paper introduces a structural figure to which §5 will return. A camera, however precisely engineered and however dense the resulting image, records the surface of a scene; it does not, on the strength of the image, recover the bones of what it photographs. A carving, by contrast, is the product of an act that engages the material on which it is performed; the carver works against the wood, the wood resists; the resulting form is the joint product of a hand and a substance whose internal grain has been encountered and worked.

The figure is not a metaphor for argument; it is the structural homologue of the cut and the world. The substrate that processes a record of the world processes the camera-image; the operation on the image, however refined, does not by itself reach the bones of the scene the image records. The world, in the cases the paper treats, is what would have been encountered had the substrate worked against the material rather than reading it; the carving is the off-cut relation; and the difference between camera and carving is the difference between cut and domain. The figure is exact in the sense the eco-NR register requires of analogy: applied term by term, it survives the application; the substrate is the camera, the data is the image, the world is the scene, the practice is the carving, and the absent relation is the relation between hand and grain.

4.10 Accumulation against model: the field's own contrast

A second case from the field's own resources confirms the figure. A mesh of triangles obtained by image-to-three-dimensional reconstruction supplies a surface but no internal model: the object behind the mesh is not constituted by the mesh; the mesh is a skin draped over an inference from photographs (Schönberger and Frahm, 2016). A parametric or building-information model, by contrast, is constituted by objects that carry material, dimension, constraint, and relation; the model contains what the mesh contains by inference and what the mesh does not contain at all (Eastman, Teicholz, Sacks, and Liston, 2018). The contrast is not a contrast between two technologies; it is the contrast between accumulation and constitution, between cut and domain, between camera and carving. The case is included because the field itself has produced both kinds of object and is in a position to confirm, from its own working categories, that the kinds are distinct.

The semantic lens closes on the structural fact it has now established at eight canonical points (Frege, 1892; Putnam, 1980; Kripke, 1982; Tarski, 1936; Gödel, 1931; Sellars, 1956; Searle, 1980; Carnap, 1928; Quine, 1960) and one engineering case (Schönberger and Frahm, 2016; Eastman et al., 2018). The substrate operates on a cut; the world it is claimed to model lies in part outside the cut; the relation from cut to world is not internal to the cut; the substrate's behaviour on the cut underdetermines its correct account of the off-cut domain. The foreclosure theorem of §3 is the formal statement of what the lens has now traced through reference, rule, inference, indeterminacy, and the formal limits of definability. The second lens turns the same structure into measure-theoretic form and gives the field's own statistical apparatus the chance to confirm the result on its own grounds.

5. The statistical lens: the bridge from semantics to measure

5.1 What the second lens does

This section runs the foreclosure theorem through the second lens, the measure-theoretic one. The point is to show that the same gap the philosophy of language identified is the gap the field's own statistical apparatus operates over. In plain terms: the loss function the system is trained to minimise is an integral against the data, and an integral does not care about what happens off the data; any extension of the system's behaviour to cases outside the data is equally good by the loss. This is a textbook measure-theoretic fact and the section states it as a proposition. The section then shows the proposition is the same fact as the foreclosure theorem in §3, restated in the field's working vocabulary, and confirms the prediction from inside the field's own world-model literature.

The first lens established the foreclosure structure through the canonical resources of reference (Frege, 1892; Putnam, 1980; Kripke, 1982; Sellars, 1956; Searle, 1980; Carnap, 1928; Quine, 1960) and the formal limits of definability from inside (Tarski, 1936; Gödel, 1931); the second lens restates the same structure in the measure-theoretic terms in which the field's own statistical apparatus operates. The transit is required because the field's positive claims are framed as claims about distributions, training, and generalisation, and the argument has to enter that frame to be felt as an immanent critique rather than as a foreign objection. The lens proceeds in four steps: the formal restatement of the foreclosure condition in measure-theoretic terms, the support-invariance proposition that closes the lens, the identification of that proposition with the foreclosure condition of §3, and the confirming experiment supplied from inside the field's own positive-result literature.

5.2 Effective support, not nominal support

The translation requires one preliminary distinction the field's own working practice already presupposes. Let the substrate be a learner trained on a sample drawn from a probability distribution P over the domain D ; let μ be the measure P induces on D . The *nominal support* of μ is the closure of the set of points carrying positive measure under μ ; in any continuous-space setting nominal support can be the whole space, since the density is positive almost everywhere a single observation might land. The cut, in measure-theoretic terms, is not nominal support but the *effective support*: the region in which the density is, in operational terms, not negligible at the resolution at which the learner discriminates and acts. The terms "in-distribution" and "out-of-distribution" are not casual labels in the field's working practice; they presuppose a distinction the formal language requires, and the present argument takes them at the field's word. The effective support is the cut; the off-cut region of measure is the off-cut region of the domain in §3.

5.3 The support-invariance proposition

With effective support fixed, the lens's central proposition follows by a standard measure-theoretic argument. Let L be a loss functional defined as an integral of a pointwise loss against the data distribution μ ; let f be the minimiser of L over a class of candidate functions; let f^* be the corresponding minimiser when f is altered arbitrarily on a set of μ -measure zero. Then $L(f^*) = L(f)$, because the integral does not register changes on null sets; the off-support behaviour of the minimiser is undetermined by the loss; any extension of f to the off-support region that preserves the on-support behaviour is a minimiser of L . The proposition is unsurprising in form and inescapable in consequence: optimisation on a distribution does not discipline off-distribution behaviour.

Proposition (Support-invariance). Let L be an objective of the form $L(f) = \int_D \ell(f(x), y(x)) d\mu(x)$, and let f minimise L . For any extension f^ of f that agrees with f μ -almost everywhere on the effective support, $L(f^*) = L(f)$. The minimiser of L is unique only up to alteration on the off-support set.*

The proposition is the measure-theoretic translation of the underdetermination structure encountered in §4. There the cut was the body of sentences (Putnam, 1980), the perceptual record (Quine, 1960), or the finite list of applications (Kripke, 1982); here the cut is the effective support of the training measure; the off-cut region is the off-support set; and the substrate's behaviour off-support is undetermined by the loss whose minimisation defines the substrate. The cases at which the off-support behaviour is consequential are the cases at which the claim made on the substrate's behalf extends to the off-support region; those are the cases the foreclosure theorem governs. The continuous case Kripke (1982) treated discretely under the rule-following heading is here the continuous case made exact in measure-theoretic dress.

5.4 The identification with the foreclosure condition

The identification is exact and is stated as a corollary. Let f be a substrate trained to minimise a loss on a distribution μ ; let an off-support element d be consequential under the claim Λ that the substrate's competence extends to D . By the support-invariance proposition, there exist two extensions f^* and f'' , both minimising L , differing at d . The two extensions stand in the relation of Theorem 1: they are indistinguishable to the substrate's operational criterion, and they disagree on d . The argument has now encountered, for a second time, the same form of insufficiency the semantic lens established (Frege, 1892; Putnam, 1980; Kripke, 1982; Quine, 1960); the recurrence is the structural one.

The consequence is sharp on its own grounds. The standard remedies the field has proposed for off-distribution behaviour cannot, by the form of the proposition, address the gap. Scale enlarges the substrate's capacity but does not address its objective; the objective remains an integral over the support, and the minimiser remains undetermined off-support. Distillation, fine-tuning, and reinforcement from feedback adjust the substrate's behaviour on the support; they do not introduce information about the off-support region the support does not contain. The remedies operate at the level the proposition forecloses: they refine the cut, they do not constitute the relation from cut to domain.

5.5 The confirming experiment from inside the literature

The empirical face of the proposition is supplied by the field's own world-model literature. Recent work on video-trained models of physical dynamics reports that the systems succeed strikingly on configurations resembling those of the training distribution, fail consistently on configurations outside it, and generalise by reverting to the nearest in-distribution configuration rather than by applying the law that would extend across the support boundary; the failure pattern is the pattern the support-invariance proposition predicts, at the place the proposition locates the failure, with the inefficacy of scale the proposition derives [D-23]. The argument requires only the structural finding: the failure is nearest-training-neighbour reversion rather than lawful extension; the failure is consistent across scale; the failure is the off-support gap the proposition forbids the objective to close.

The experiment is included as confirmation, not as load-bearing premise. The proposition stands on the formal argument from the loss integral; the experiment is the field's own apparatus producing the prediction the proposition derives. The status of the citation is the status the paper's apparatus assigns: the empirical claim is conditional on the specific result; the structural claim is the load-bearing claim, and the structural claim stands on the proposition.

5.6 The world-model programme as the central case

The lens's strongest application is to the world-model programme itself, the programme of building learners that recover the dynamics of the physical world from observation. The programme is sold as the line of attack on the foreclosure problem: where language models operate on linguistic traces, world models are taken to operate on the world's behaviour directly, and the cut is taken to approach the world. The proposition forecloses the line on its own grounds. A model trained on video is trained on a sample of trajectories the world produces; the sample is the cut; the regions of physical configuration space not sampled are the off-cut; the model's behaviour off-sample is undetermined by its objective. The world-model programme inherits the foreclosure structure from the cut/world distinction the camera-and-carving figure of §4.9 fixed: the model receives the camera-image of dynamics, not the dynamics themselves; the relation between the recorded trajectory and the law that produced it lies, by the proposition, off the support the objective sees.

5.7 The identifiability result that smuggles a world-relation

The lens closes by considering, for the discipline of immanent critique, the field's strongest positive result on the question of grounding. Recent identifiability theorems in the self-supervised and joint-embedding traditions establish that, under stated conditions, the true latent variables underlying observed data are recoverable up to a small set of permitted ambiguities [D-24]. The result is taken in the relevant literature as evidence that representational learning, given enough data and the right architecture, recovers the structure of the world it is trained on. The structural reading the present argument requires is that the result's force depends on its conditions, and the conditions are not innocent.

A representative result of this kind requires assumptions about the latent dynamics: that the latent process is a specified diffusion or has a particular mean-reversion structure (the Ornstein-Uhlenbeck case is canonical), or that the latent distribution lies in a specified parametric family [D-24]. The recovery follows from the conditions; without the conditions, the recovery fails; the theorems are stated honestly in this form. The structural reading of the conditions is the reading the present argument requires. A condition on the latent dynamics is not a property the learner discovers from the data; it is a property the learner is told to assume about the world the data come from. A projectile under gravity is not Ornstein-Uhlenbeck; an explosion is not mean-reverting; a phase transition is not Gaussian. The condition that licenses the identifiability is a substantive claim about the world, supplied from outside the data, and the recovery is the recovery of what the claim already encodes. The identifiability result is, in the structural sense the foreclosure theorem requires, a result of the form: given a substantive assumption A about the world, the data permit recovery of the world's structure consistent with A. The assumption is the world-relation; the data alone do not supply it; the result confirms the proposition rather than overturning it.

The argument has now encountered, for a third time inside the statistical lens, the same form of insufficiency. The first encounter was the support-invariance proposition itself; the second was the world-model programme's symptomatic failure [D-23]; the third is the field's own positive identifiability results carrying their world-relation in the assumptions that license the result [D-24]. The lens has reached the same structural finding as the semantic lens (Frege, 1892; Putnam, 1980; Kripke, 1982; Tarski, 1936; Gödel, 1931; Sellars, 1956; Searle, 1980; Carnap, 1928; Quine, 1960), by an entirely independent route through the measure theory the field itself relies on. The third lens, formal and cybernetic (Ashby, 1956; Conant and Ashby, 1970), closes the convergence.

6. The formal lens: the conditions of the act of regulating

6.1 What the third lens does

This section runs the foreclosure theorem through the third lens, the cybernetic one. The point is to show that what cybernetics has been saying about regulation since the 1950s is the same point the other two lenses have been making: a regulator can only handle a system whose variety it can match, and a regulator built from a training cut has the variety of the cut, not of the world. The section uses two results, both standard: Ashby's Law of Requisite Variety, which says a regulator must have at least as many distinct internal states as the system it regulates, and the Conant-Ashby theorem, which says that to regulate a system the regulator must contain a model of it. The section then shows that the field's own positive identifiability results carry the same point in disguise: they recover the world only when the conditions that license the recovery have already been assumed, which is to say the world has been supplied from outside the data.

The first lens passed the foreclosure structure through reference (Frege, 1892; Putnam, 1980; Kripke, 1982; Sellars, 1956; Searle, 1980; Carnap, 1928; Quine, 1960; Tarski, 1936; Gödel, 1931); the second through measure (the support-invariance proposition); the third passes it through the formal results that govern what it takes for a system to regulate another system at a specified resolution. The resources are again from inside the relevant canon: Ashby's *Introduction to Cybernetics* (Ashby, 1956) and the theorem of Conant and Ashby (1970). The lens establishes that the field's own positive concept of generalisation, the concept by which a system trained on a cut is said to extend its competence to the domain, presupposes a model-of-the-world relation that the cut does not contain. The argument runs through three results: the Law of Requisite Variety (Ashby, 1956), the

Conant-Ashby theorem (Conant and Ashby, 1970), and the scope conditions of the identifiability results ([D-24]), each shown to carry the same structural commitment.

6.2 Ashby: the Law of Requisite Variety

Ashby's Law (Ashby, 1956, p. 207) states the formal condition for regulation: a regulator can hold a system within an acceptable set against disturbances only if the regulator's variety, the count of distinct internal states it can adopt, is at least equal to the variety of the disturbances against which it must regulate. The Law is presented in Ashby (1956) with the table-of-disturbances demonstration: an arrangement of disturbance values across rows and regulator responses across columns shows that the regulator must have at least as many distinct response settings as the disturbance has distinct values, or some disturbance value is met by an inadequate response, and the system leaves the acceptable set. The result is, in Ashby's (1956) own formulation, "only variety can destroy variety"; the form is a conservation principle on the resolution of regulation. The requisite variety can be matched in two ways: by amplifying the regulator, supplying it with more distinct responses, or by attenuating the system, reducing the variety it presents (Beer, 1974); the engineering arts of regulation have always operated in both directions.

The Law's relevance to the present argument is the relevance of variety to the cut/domain relation. A substrate that has been trained on a cut has, at most, the variety the cut supplies; the variety of the disturbances it will encounter in deployment is the variety of the domain; the disparity between the two is the disparity between cut and domain. The Law gives the formal consequence: the substrate cannot, by Ashby's (1956) own argument, regulate against disturbances whose variety exceeds the variety it has internalised from the cut, and any such disturbance falls outside the substrate's resolvable distinctions. The substrate's competence on the cut, however refined, supplies a variety bounded by the cut; the variety of the domain is not bounded by the cut; the Law forecloses the substrate's regulation of the off-cut domain on its own resources.

6.3 Conant-Ashby: every good regulator must be a model

The Conant-Ashby theorem (Conant and Ashby, 1970, *International Journal of Systems Science* 1(2), pp. 89–97) sharpens the Law's lesson into the form the present argument requires. The theorem states that every regulator that holds a system optimally within an acceptable set must, in a formal sense the paper makes precise, be a model of the system it regulates: its internal states must be in a homomorphic correspondence with the states of the system, and its transitions must mirror the system's transitions at the resolution the regulation requires (Conant and Ashby, 1970). The argument is internal to the formal apparatus of regulation; it shows that the requisite variety is, more than a quantity, a structural correspondence: the regulator does not only contain enough internal distinctions, it contains *the right* internal distinctions, organised in the way the system's own dynamics organise their states. The result is a theorem in the discipline's own formal language, derived from the conditions of optimal regulation themselves.

The relevance to the foreclosure structure is exact. A substrate that is to govern a domain must, by Conant-Ashby (1970), be a model of the domain at the resolution the governance requires; a substrate that has been trained on a cut is, at most, a model of the cut; the cut is a model of the domain only if the relations between cut elements and domain elements are themselves contained in the cut, which by Theorem 1 is precisely what the cut cannot contain. The substrate's competence on the cut does not by itself constitute it a model of the domain; the model-of-the-domain relation is the off-cut remainder Theorem 1 forecloses; the Conant-Ashby theorem (1970) confirms, from inside the cybernetic apparatus, that this remainder is what optimal regulation requires.

6.4 Identifiability and the conditions that license recovery

The same structure appears in the field's positive identifiability results when their scope conditions are read as the apparatus they are. A theorem of identifiability, as §5.7 noted, recovers the true latent structure of a process from observations of that process under stated assumptions about the latent dynamics [D-24]; the recovery follows when the assumptions hold and fails when they do not. The scope conditions are not, in the formal sense the discipline uses, optional; they are the conditions under which the theorem applies, and the recovery is the recovery of structure consistent with them.

Read through the Conant-Ashby (1970) lens, the scope conditions are the model-of-the-system the identifiability machinery requires. The condition that the latents are Ornstein-Uhlenbeck is the assumption that the system's dynamics belong to a specified family with stable, mean-reverting behaviour; the condition that the distribution is in a specified parametric class is the assumption that the system's variation is constrained in a specified way. These assumptions are, in the cybernetic sense (Ashby, 1956; Conant and Ashby, 1970), the regulator's model of the system: they specify the resolution and structure the apparatus presupposes about the world the data come from. When the assumptions hold, the identifiability machinery recovers the latent structure consistent with them; when the assumptions fail, the recovery fails because the model the machinery presupposed of the system does not match the system. The argument has now encountered, for a third time across the lenses, the same form of insufficiency: a substrate that proposes to recover the structure of a domain by operating on a cut succeeds only when a relation between cut and domain has been supplied from outside the cut.

7. The convergence of the three lenses

This section closes Part II. The three lenses of §§4-6 have established the same fact from three different canons. This section says what that means. The three are not three loose analogies that happen to land in the same place; they are three statements of one structural fact, and the section lands the structural statement Part II has been building toward. §10 in Part IV will then prove formally, drawing on Lawvere (1969) and Yanofsky (2003), that the three lenses are instances of a single categorical theorem; this section states what the proof in §10 will be the proof of.

The three lenses now stand together and the convergence calls for the structural statement the paper has been building toward.

The semantic lens established that the relation of reference (Frege, 1892), the rule (Kripke, 1982), the inferential role (Sellars, 1956), the conditions of symbol manipulation (Searle, 1980), the indeterminacy of translation (Quine, 1960), and the formal limits of definability (Tarski, 1936; Gödel, 1931) all show that the cut of expressions, applications, or sentences does not contain the world-relation that fixes the correct account on the off-cut domain. The statistical lens established the same fact in measure-theoretic form: an objective that integrates against the data distribution does not discipline behaviour on the off-support set, and the field's own positive identifiability results carry their world-relation in the assumptions that license the recovery [D-23]. The formal lens established the same fact through requisite variety (Ashby, 1956) and the Conant-Ashby theorem (Conant and Ashby, 1970): the variety the cut supplies is insufficient to the domain's variety, and the model-of-the-domain relation the regulation requires is precisely the off-cut remainder.

The three findings are not three analogous results that happen to share a slogan. They are three formal statements of one structure, derived from three independent canons, agreeing on the same elements. The cut in each case is the substrate's operational base; the off-cut remainder in each case is the domain element whose correct account the claim made on the substrate's behalf requires; the relation from cut to off-cut remainder in each case is what the cut does not, by the result of that lens, contain. The three lenses are in a structural correspondence with one another; the correspondence is exact and is reached in §10, where the formal identity of the three is shown to follow from a single fixed-point argument (Lawvere, 1969; Yanofsky, 2003). The convergence at this point is the empirical claim about the lenses; §10 will be the theoretical statement that the empirical claim is not coincidence.

The structural landing of Part II is the statement the three lenses now license. *Foreclosure is one structure proved three times over inside the canons of reference (Frege, 1892; Putnam, 1980; Kripke, 1982; Sellars, 1956; Searle, 1980; Carnap, 1928; Quine, 1960; Tarski, 1936; Gödel, 1931), measure (the support-invariance proposition; [D-23]), and regulation (Ashby, 1956; Conant and Ashby, 1970).* The substrate that operates on a cut cannot, on the strength of that operation alone, certify a correct account of an off-cut domain whose competing extensions are indistinguishable to it; the claim made on the substrate's behalf, when unrestricted to the cut, is unsupported by the substrate's operation; and the field's own apparatus, in three independent forms, has been encountering and stating this fact. The first wall stands. The paper now turns to the two walls that govern the temporal behaviour of deployed substrates: the wall of compounding hazard, which governs the law of motion of the failure rate (Williams, 1991, §12.15; Lévy, 1937), and the wall of horizon-mismatch, which governs the relation between the agent's evaluation horizon and the system it acts upon (Chen, 1999; Ashby, 1956; Conant and Ashby, 1970).

Part III. The Two Temporal Walls

7. The wall of compounding hazard

This section states and proves the second wall. The first wall is about what a system can reach. The second wall is about what happens to a system over time once it is deployed. In plain terms: if a system makes the same kind of mistake at the same rate forever, then deploying it long enough will produce that mistake a lot of times, however small the rate is; and if the rate goes down each time the mistake occurs, so that the mistake suppresses itself, then it only happens a few times. The section uses a textbook result from probability, the Borel-Cantelli lemmas, to make this dichotomy precise. The wall is then about which side of the dichotomy the deployed systems sit on, and the architectural premise of §7.2 settles it: the systems sit on the side where the rate does not respond to its own failures, and the lemma's conclusion follows.

7.1 What the second wall governs

The first wall is a statement about what a substrate can reach at a single instant. The second wall is a statement about what happens to a substrate that is deployed, that is, that produces outputs repeatedly over a duration long enough for the rare to occur. The variable the wall governs is not the substrate's per-action probability of producing a catastrophic output; that probability can be, and in the systems under treatment is, very small. The variable is the law of motion of that probability under realised cost: whether the per-action probability on a given catastrophic mode falls in response to the mode actually producing a catastrophe, or whether the probability is stationary, that is, unchanged from one action to the next regardless of what the substrate has produced. The wall is established by the pair of Borel-Cantelli lemmas (Williams, 1991), with the divergent regime stated in the conditional form due to Lévy (Lévy, 1937; Williams, 1991, §12.15) that drops the independence assumption the classical second lemma requires.

7.2 The architecture of deployment fixes the wall's premise

The wall's premise is supplied by the architecture of the systems under treatment, not by an external assumption about their hazard. A deployed language model, a deployed world-model agent, a deployed control system whose parameters are frozen between training runs produces, on each action, an output sampled by the substrate's stochastic procedure; the per-action probability that a given catastrophic mode is realised is, in the relevant operational sense, the same on the n -th action as on the first; the system does not, between actions, alter the architecture from which the sampling is drawn. The softmax that closes the autoregressive procedure assigns strictly positive probability to every token in its vocabulary (Goodfellow, Bengio, and Courville, 2016, §6.2.2.3, eq. 6.29); the per-action probability of a catastrophic output, however small, cannot be driven to zero on the strength of the architecture; the catastrophic mode is, in the rate-stationary sense, available on each action. The premise of the wall is this architectural fact: the substrate's per-action hazard on a consequential mode does not respond to that mode's realised catastrophes by being driven to zero, and the wall's conclusions follow.

The premise is to be read at the level of the deployed system's effective output distribution, not only at the level of the raw substrate's softmax. Production systems wrap the substrate in engineered safety layers; refusal classifiers, output filters, blocklists, ensembles of moderation models, and similar post-hoc apparatus alter the conditional probabilities of the substrate's outputs reaching the user. The wall's architectural premise applies to the conditional probability of a consequential catastrophic mode in the deployed output stream, after any such filtering. The objection that safety layers move the system out of the divergent regime is met by the same dichotomy Theorem 2 establishes: a layer that drives the conditional sum on the consequential mode to convergence places the deployed system in the convergent regime, by Borel-Cantelli's first lemma (Williams, 1991); a layer that suppresses the mode probabilistically without driving the conditional sum to convergence leaves the deployed system in the divergent regime, by Lévy's extension (Lévy, 1937; Williams, 1991, §12.15). The argument does not turn on whether the substrate is wrapped or unwrapped; it turns on whether the effective deployed hazard on the consequential mode is stationary in the relevant sense, and the dichotomy is the same.

7.3 The Borel-Cantelli dichotomy

The formal apparatus separates exactly two regimes. Let E_n be the event that the n -th action produces a catastrophe on the mode under consideration, and let $p_n = (E_n)$ be the per-action probability. The first Borel-Cantelli lemma (Williams, 1991) states that if $\sum p_n < \infty$, then $(E_n \text{ infinitely often}) = 0$; the catastrophes are almost surely finite in number. The second lemma, in the conditional form due to Lévy (Lévy, 1937; Williams, 1991, §12.15), states that if $\sum (E_n \mid \mathcal{H}_{n-1}) = \infty$ almost surely, where $\mathcal{H}_{n-1}$ is the σ -algebra of the history through action $n-1$, then $(E_n \text{ infinitely often}) = 1$; the catastrophes are almost surely infinite in number. The conditional form is essential and is the form the present argument requires: independence among the E_n is not assumed, and in the relevant cases independence does not hold; the conditional form drops the assumption.

The dichotomy is the wall's structural finding. The behaviour of the deployment is one of two things, with no third: either the per-action probabilities, conditional on the history, sum to a finite quantity, in which case the realised catastrophes are almost surely finite (Williams, 1991); or they sum to an infinite quantity, in which case the realised catastrophes are almost surely infinite (Lévy, 1937; Williams, 1991, §12.15). The dichotomy is established without any assumption about the size of p_n ; a very small probability that does not collapse on cost produces unboundedly many catastrophes; a substantial probability that collapses sharply on cost produces only a few. The wall's variable is the law of motion of the conditional probability, not its level.

7.4 The theorem of compounding hazard

The wall is now stated as a theorem with three modes of deployment, each forcing a distinct form of the conclusion.

Theorem 2 (Compounding Hazard). *Let a substrate be deployed for an indefinite sequence of actions, with E_n the event of a catastrophe on a specified mode at action n , and $p_n = (E_n \mid \mathcal{H}_{n-1})$ the conditional per-action probability of the mode given the history.*

(i) *Recurrent case: if the mode admits unbounded repetition and $\sum p_n = \infty$ almost surely, then almost surely the mode produces catastrophes infinitely often, by the Lévy form of Borel-Cantelli (Lévy, 1937; Williams, 1991, §12.15).*

(ii) *Terminal case: if the mode is terminal, that is, the deployment ends on the first realised catastrophe, and $\sum p_n$ diverges over the deployment's planned actions, then the probability that the deployment ends in catastrophe approaches one as the number of planned actions grows.*

(iii) *Replicated case: if the substrate is replicated across many simultaneous deployments and the mode is independently available in each, then the expected number of realised catastrophes across the replication, equal to the sum of the per-action probabilities across deployments and actions, grows without bound as the replication and the per-deployment duration grow.*

The three forms are the three deployment regimes the field's own systems occupy. A conversational assistant in continuous use is the recurrent case: the mode does not end the deployment, the deployment continues across many users and many actions, and the expected number of consequential outputs accumulates. An autonomous vehicle, an automated trading system, or any agent whose catastrophic mode ends the run is the terminal case: each deployment is a single trial against the cumulative hazard, and the probability of catastrophic termination approaches one as the deployment extends. A fleet of replicas across many simultaneous deployments is the replicated case: the expected count across the fleet grows linearly in the product of replicas and per-replica actions, and the variance of that count grows correspondingly. Theorem 2 says, formally, what the architecture commits its deployments to.

7.5 The mode-indexed transformation

A regulator that is to make a difference to the wall's conclusion must transform the conditional probability p_n on the mode through the cost realised on that mode, and must do so in a way that depends on the mode's identity. The point is structural rather than rhetorical. A substrate that learns indiscriminately from each event, raising or lowering its outputs across the whole behavioural distribution, does not deform the per-mode hazard sharply enough to drive any single mode's conditional sum to convergence; the deformation is spread across modes, and the catastrophic mode remains in the divergent regime. The regulator that can make the difference is the regulator that, given a realised catastrophe on a specific mode, indexes the cost to that mode and reduces the conditional probability of recurrence on that mode sharply enough that the mode's p_n sum converges. The conditions under which this kind of mode-indexed transformation is in place, given a substrate frozen between training runs and patched on cycles separated by long intervals, are the conditions the third subsection of this section reaches. The mode-indexed transformation is not a property every regulator can claim; it is a condition the wall's first regime (Williams, 1991) requires.

7.6 The laundering of the per-action figure

The wall's empirical face appears in the field's own headline safety figures. A statement that a system is, on a given evaluation, ninety-nine point nine per cent safe is a statement about a per-action rate measured against a curated evaluation set: the probability that a single action from the system on a single sampled input from the evaluation produces an unsafe output is the figure reported. The figure is, in the operational sense the field uses, a per-action quantity. The figure is then heard, by the regulatory, commercial, and public audiences for which it is supplied, as a guarantee about the system's deployed behaviour over a working lifetime: the system is taken to be reliably safe in operation (Goodfellow, Bengio, and Courville, 2016, §11.1, characterising the field's safety target as a per-example error rate on a fixed evaluation set). The two quantities behave oppositely under repetition. A per-action probability of 10^{-3} is a near-certainty of catastrophe over 10^4 actions; the deployed lifetime of a continuously running assistant or an agentic system is several orders of magnitude longer. The laundering is not a deception in the rhetorical sense; it is a structural feature of how a quantity measured on a curated cut is presented in the language a deployment requires. The headline reports the figure that the architecture permits and the evaluation supports; the deployment carries the conditional sum the headline does not state.

The paper's framing of the second wall does not rest on the laundering claim alone. The architectural premise of §7.2 is sufficient: the substrate is frozen, the per-action probability on a consequential mode is, in the rate-stationary sense, available on each action, and Theorem 2 applies. If the curated figure is also recast as a lifetime guarantee, the gap between the per-action quantity and the deployed sum is the laundering (Goodfellow, Bengio, and Courville, 2016, §11.1). If the figure is honestly stated as a per-action rate and the audience hears it as a per-action rate, the wall still stands, because the deployed sum is still the deployed sum. The laundering is the symptom, not the cause; the cause is the architecture and the dichotomy Theorem 2 establishes (Williams, 1991, §12.15).

7.7 The patch cycle: a conditional argument

The architectural premise of §7.2 already places the unpatched deployed substrate in the divergent regime of Theorem 2; the wall stands on the unpatched substrate by Borel-Cantelli's conditional form (Lévy, 1937; Williams, 1991, §12.15). The conditional argument of this subsection concerns a different question: whether the patch cycles the field actually operates move the deployed substrate out of the divergent regime onto the convergent one. The wall is not, in this sense, conditional; what is conditional is the field's claim to have escaped it.

The remedy the field has assembled for the wall's conclusion is the patch cycle: a realised catastrophe is recorded as an incident, a corrective update is prepared and tested against curated evaluation, and the update is deployed on a cycle separated from the incident by an interval of days, weeks, or months. The structural question is whether the patch cycle implements the mode-indexed transformation of §7.5 sharply enough to drive the catastrophic mode's p_n sum to convergence, or whether the apparent improvement is confined to the curated evaluation while the deployed conditional probability on the mode persists.

The structural condition for the convergent regime (Williams, 1991, first lemma) is sharp: the conditional sum on the mode, over deployed actions in the interval between patches, must remain bounded as the deployment continues. The condition is met if each patch drives the deployed probability on the mode toward zero with a sharpness sufficient to overcome the rate of new actions in the interval. The condition is not met if the deployed probability on the mode is reduced by an amount that fails the convergence criterion, or if the patch's effect is confined to the curated set on which it was tested and does not transfer to the deployment distribution, or if new modes of the same form appear at a rate that compensates for the patches applied to the old ones.

Whether deployed patch cycles meet this condition is an empirical question, and the deployment-record evidence to settle it is, in the present argument, conditional [D-26a]. The argument therefore states the patch case explicitly conditional: where the patch cycle implements the mode-indexed transformation sharply, the substrate moves to the convergent regime of Theorem 2 on the modes the patches address; where the patch cycle does not implement the transformation sharply, the deployed substrate remains in the divergent regime (Lévy, 1937; Williams, 1991, §12.15). The architectural premise of §7.2 establishes which regime the un-patched substrate occupies; whether the patch cycle moves the substrate out of it is the question the deployment data would settle, and the paper does not assert what the deployment data have not shown.

The compounding-hazard wall stands as a theorem about the law of motion of the per-action hazard under realised cost (Williams, 1991, §12.15; Lévy, 1937), with the architectural premise supplying the divergent regime for the unmodified substrate, and the patch cycle a conditional remedy whose efficacy is the empirical question the deployment data would settle [D-26a]. The wall does not depend on the laundering of the figure or on the empirical efficacy of the patch; it depends on the dichotomy Theorem 2 establishes and on the architectural fact that the per-action probability on the mode does not respond to cost on the strength of the architecture itself. The paper now turns to the third wall, which governs a variable Theorem 2 does not reach.

8. The wall of horizon-mismatch

This section states and proves the third wall. The first wall is about what a system can reach; the second is about what happens to it over time on a per-action basis. The third is about what happens when a system acts on something that takes a long time to recover. In plain terms: if a recommender system, an automated trader, an agentic AI acts inside an institution, a supply chain, an ecology, and the system can only evaluate its actions over minutes or days or weeks, then a slow thing that recovers over years is invisible to the system inside its window. Whether the slow thing is being maintained or being depleted looks the same to the system; it only knows what its evaluation window can see. So it acts on what the window shows and spends the slow thing without seeing what it spent. The section makes this precise by joining a textbook result from control theory, finite-horizon observability, to the Conant-Ashby theorem from cybernetics. The wall holds when the system's window is shorter than the slow thing's recovery cycle and the healthy and depleting states look the same inside the window.

8.1 What the third wall governs

The second wall is established and stands as a theorem about the law of motion of a per-action probability under realised cost (Williams, 1991, §12.15; Lévy, 1937). The third wall governs a different variable, on a different temporal scale, and applies to a system in a different posture. It applies not to a substrate producing outputs against a curated evaluation but to an agent acting within a system whose state includes a slow variable; the agent's actions affect the slow variable; the slow variable replenishes on a characteristic time; the agent receives feedback on a horizon over which it evaluates its actions and updates its policy. The variable the third wall governs is the relation between the agent's evaluation horizon and the system's replenishment cycle. The wall holds when the horizon is shorter than the cycle and the sustainable and depleting states of the slow variable are not distinguishable inside the horizon. The wall does not depend on the second wall's variable; a system with perfect cost-collapsing hazard on each action can still sit on the wrong side of this wall, and the paper now establishes why.

8.2 The architecture of the third wall

Five elements are required, each supplied by the operating arrangement of the systems the wall is to govern. An agent A acts in time, producing an action a_t at each time step; the action is taken with respect to a system T , whose state x_t the agent influences. The state contains a slow variable s_t whose characteristic time of replenishment is τ_S , that is, the time over which s_t recovers, partially recovers, or is consumed in response to the agent's actions. The agent's evaluation horizon is τ_A , the span over which its objective integrates the consequences of its actions and updates its policy; this is the operational horizon of the agent's feedback, set by the temporal discount in its reward, the rollout length of its planning, the look-ahead of its model, or the length of its training window. The wall's first hypothesis is $\tau_A < \tau_S$: the agent's horizon is shorter than the system's replenishment cycle. The wall's second hypothesis is the observational-equivalence hypothesis: the sustainable and depleting trajectories of s_t produce identical observation sequences within any window of length τ_A , and diverge only on scales τ_S and beyond.

The two hypotheses jointly capture the practical situation. An agent optimising over short horizons against a system replenishing over long ones is the standard operating arrangement of the systems the third wall governs: a recommender system tuned to short-horizon engagement metrics acting on a media ecology whose trust and informational health vary on cycles of years; an automated procurement system tuned to quarterly cost metrics acting on a supplier ecosystem whose stability is measured in years; an agentic AI tuned to within-episode reward acting on institutions, infrastructures, and ecologies whose replenishment cycles run into decades. The horizon mismatch is a feature of the deployment, not an accident of any particular system's tuning.

8.3 The horizon bound, by finite-horizon observability

The first formal input is finite-horizon observability (Chen, 1999): a regulator whose feedback is confined to a window of length τ_A can distinguish two states of the regulated system only if the states produce distinct observation sequences within that window. The result is a standard one in linear system theory (Chen, 1999) and transfers to the present case without modification: the regulator's resolvable variety within its horizon is bounded by the distinctions its feedback supplies in that horizon; states observationally equivalent across the horizon are not in the regulator's resolved state space. The wall's second hypothesis, observational equivalence of sustainable and depleting trajectories within τ_A , places those states in the unresolved equivalence class; the agent's feedback does not, by the observability bound (Chen, 1999), separate them.

8.4 The Conant-Ashby step, and the theorem

The second formal input is the Conant-Ashby theorem (Conant and Ashby, 1970): a regulator that holds a system in the acceptable set must be a model of the system at the resolution the regulation requires. The model must contain the distinction between sustainable and depleting states, since holding the system in the sustainable set requires the regulator to act differently in the two cases (Ashby, 1956; Conant and Ashby, 1970). The wall's two hypotheses jointly preclude this: by observability (Chen, 1999), the distinction does not enter the regulator's resolved state space; by Conant-Ashby (1970), a regulator whose model lacks the distinction cannot regulate the system at the resolution the distinction requires.

The third wall is now stated.

Theorem 3 (Horizon-Mismatch). *Let an agent A act on a system T containing a slow variable s_t with replenishment cycle τ_S , where the agent's evaluation horizon τ_A satisfies $\tau_A < \tau_S$, and where the sustainable and depleting trajectories of s_t are observationally equivalent within any window of length τ_A . Then the agent's feedback does not distinguish sustainable from depleting states of the slow variable (by finite-horizon observability; Chen, 1999), and any regulator built on that feedback lacks the distinction the Conant-Ashby theorem (Conant and Ashby, 1970) requires of a good regulator of T at the resolution sustainable/depleting; hence the agent cannot, on the strength of its own feedback, preserve the slow variable.*

The theorem composes two standard results. Observability (Chen, 1999) supplies the first bound; Conant-Ashby (1970) supplies the consequence; the wall's two hypotheses connect the two. The composition is the paper's own; the inputs are canonical (Ashby, 1956; Conant and Ashby, 1970; Chen, 1999).

8.5 The order-of-magnitude mismatch

The hypotheses of the theorem are not abstract: in the cases the paper treats, the magnitudes separate by orders. An agent's evaluation horizon, set by the rollout length, the within-episode return, or the training-window length, is at most weeks and is typically minutes to hours. The replenishment cycle of an engineering culture is years to decades; of a supplier ecosystem, years; of an institutional trust, decades; of ecological buffer capacity, decades to centuries. The mismatch is not marginal; the agent's horizon is two to four orders of magnitude shorter than the cycles of the systems it is to govern. The order-of-magnitude statement is offered as common-knowledge orders of magnitude, not as precise measurements; the argument requires only that the cycles exceed the agent's horizon by enough to keep the sustainable and depleting trajectories indistinguishable within the horizon, which holds across the range of slow institutional, productive, and ecological processes the wall is to cover.

8.6 Why the third wall is not the second wall

The two temporal walls are independent and the paper now states the structural reason in advance of the formal independence proof of §9. The second wall's variable is the law of motion of the per-action hazard under realised cost (Williams, 1991, §12.15; Lévy, 1937): whether the hazard responds to its own catastrophes by collapsing on the mode. The third wall's variable is the relation between the agent's evaluation horizon and the system's replenishment cycle (Chen, 1999): whether the agent's feedback can resolve the distinction sustainable/depleting. The two variables are different in kind and orthogonal in effect. An agent with a perfectly cost-collapsing hazard, that drives the per-action probability of catastrophe to zero on each realised cost, can still sit on the wrong side of the third wall if its evaluation horizon is shorter than the replenishment cycle of the slow variable it acts upon; the cost collapsing on each catastrophe does not, by itself, lengthen the horizon. An agent with no cost collapsing whatever, whose hazard is stationary in the second wall's sense (Williams, 1991, §12.15), can still sit on the right side of the third wall if its evaluation horizon exceeds the system's replenishment cycle; the stationary hazard does not, by itself, shorten the horizon. The two walls bite on different variables, and neither implies the other.

8.7 The irreversibility distinction

The third wall is sharpened by the distinction between estimation and irreversible depletion. A regulator that fails to estimate the slow variable correctly within its horizon will, given enough time, observe the failure: when the slow variable crosses the threshold at which depletion becomes legible on a scale the horizon does reach, the regulator updates. The point of the third wall is that the observation is then too late: the slow variable is depleted, not mis-estimated; the resource is spent, not in error; the regulator's eventual update on the now-legible signal does not restore what has been consumed during the horizon over which the consumption was unresolved. The wall does not say the regulator will never see the depletion; it says the regulator will see the depletion only when the depletion has crossed into the longer cycle, and the regulator's update at that point cannot recover what the slow variable has already lost. The mismatch is between the cycle on which the resource is consumed (short, the horizon) and the cycle on which the consumption is legible (long, the replenishment cycle). The irreversibility is the gap between consumption and legibility.

The estimation/irreversibility distinction sharpens the wall against a standard line of objection. The objection runs: a regulator with a sufficiently rich state estimator, a Kalman filter or its modern equivalents, can extrapolate the slow variable from its short-horizon observations and so resolve the sustainable/depleting distinction without lengthening the horizon. The reply, which the wall's two hypotheses already foreclose, is that the extrapolation is itself bounded by the observability conditions of §8.3 (Chen, 1999): an estimator can extrapolate only what its observations within the horizon discriminate, and the wall's observational-equivalence hypothesis denies the estimator the discrimination it would need. A richer estimator does not see what no estimator in the class can see; the equivalence is in the observations, not in the apparatus that processes them. The objection has been pre-empted by the construction of the hypotheses.

8.8 The world-model programme as the double-hit

The third wall returns the world-model programme to the centre of the argument under a second heading. The first wall foreclosed the world-model programme as a programme that learns the world from observation: the substrate operates on a cut of recorded trajectories, the law it would need to extend across the cut boundary is not in the cut, and the foreclosure stands (Frege, 1892; Putnam, 1980; Kripke, 1982; Quine, 1960; the support-invariance proposition; Ashby, 1956; Conant and Ashby, 1970). The third wall now strikes the world-model programme again, in a different role and at a different stage. Where the first wall struck the substrate at the moment of learning, the third strikes it at the moment of acting: a world-model deployed as an agent, optimising over its model-based rollouts of trajectory, operates with an evaluation horizon set by the rollout length; the rollout length is the agent's window into the consequences of its actions; institutional, infrastructural, and ecological systems on which a deployed agentic world model would act have replenishment cycles vastly exceeding the rollout length. The agentic world-model programme inherits the horizon-mismatch wall directly (Chen, 1999; Conant and Ashby, 1970); the foreclosure wall and the horizon-mismatch wall strike it independently.

The double-hit is the structural finding the third wall and the first wall together establish about the world-model programme: it forecloses as a programme of learning, by Theorem 1 and the support-invariance proposition; it horizon-mismatches as a programme of acting, by Theorem 3 (Chen, 1999; Conant and Ashby, 1970). The two walls strike at the two distinct stages of the programme's ambition. The independence proof of §9 will state this formally; the paper here states the structural fact in advance: the programme that takes itself to be the line of attack on the foreclosure problem inherits, in its agentic form, a second independent wall the foreclosure programme does not face.

8.9 The human posture against the wall

The wall is sharpened by the contrast with a human regulator whose feedback is not confined to a short horizon. A human institution, a guild, a craft tradition, a discipline of professional practice carries the slow variable in its own continuity: the institution remembers what the slow cycle has been doing across decades because the institution itself replenishes on that cycle; its members are trained into the long memory that no within-episode observation supplies. The contrast is not a romantic claim about human cognition; it is the structural claim that the institution and the slow system are coupled on the slow cycle, while the agent and the slow system are coupled only on the agent's short cycle. The agent that replaces the institution is not, in any operational sense, the institution's continuation; it is the institution's replacement by a regulator that has no slow memory, on a substrate whose horizon is set by something other than the system's replenishment cycle. The replacement is structural, and the third wall (Chen, 1999; Conant and Ashby, 1970) is the structural consequence.

The third wall stands as a theorem on the relation between the agent's evaluation horizon and the system's replenishment cycle, established from finite-horizon observability (Chen, 1999) and the Conant-Ashby theorem (Conant and Ashby, 1970), with the observational-equivalence hypothesis fixing the application to the present case. The wall does not depend on the second wall's variable; the two walls are independent in the strict sense the paper now turns to prove. Part IV establishes the formal independence of the three walls and the identity, beneath the appearance of three results, of the structural fact the first wall states three times over (Lawvere, 1969; Yanofsky, 2003).

Part IV. Independence, Identity, Boundaries

9. The independence of the three walls

This section proves that the three walls are independent of each other. The argument matters because the no-exit result the paper has been pointing toward depends on it. In plain terms: if the three walls were variations of the same underlying problem, then a repair to one of them might fix the others, and the field's existing toolkit would have more reach than the paper is claiming. The section shows that the three walls are genuinely separate, governing different variables, by constructing systems that fail on one wall while passing the others, in both directions, for each pair. The result is that a system failing all three walls at once, which is what the deployed systems do, needs three independent repairs and no current tool reaches more than one.

9.1 What independence requires

The three walls have been established as theorems, each from its own resources, each governing its own variable. The first wall stands on the underdetermination canon (Frege, 1892; Putnam, 1980; Kripke, 1982; Quine, 1960) and the support-invariance proposition, with the formal lens running through Ashby (1956) and Conant and Ashby (1970); the second wall stands on Williams (1991, §12.15) and the Lévy (1937) conditional extension of Borel-Cantelli; the third wall stands on finite-horizon observability (Chen, 1999) joined with the Conant-Ashby theorem (Conant and Ashby, 1970). The paper now establishes that the three are independent of one another: that the satisfaction of any one wall's escape condition does not entail the satisfaction of either of the others, that the failure modes the three walls describe are distinct, and that the toolkit assembled to repair one wall does not, by that repair, reach the others. The form of the argument is the standard form for independence: it is shown that for each pair of walls there exist systems satisfying one and failing the other, in both directions, so that no relation of implication binds the pair. Three pairs of counterexamples close the argument.

9.2 The independence statement

Proposition (Independence). *Let W_1 be the foreclosure wall, W_2 the compounding-hazard wall, W_3 the horizon-mismatch wall. The three walls are pairwise independent in the following sense: for each pair (W_i , W_j) with $i \neq j$, there exists a system that satisfies the escape condition for W_i and fails on W_j , and a system that satisfies the escape condition for W_j and fails on W_i . Hence no escape from one wall implies escape from any other.*

The proof exhibits the required systems pair by pair.

9.3 W_1 and W_2

Forecloses, hazard-collapses. A substrate trained on a cut that does not cover the domain, but equipped with a perfect mode-indexed correction such that the per-action probability on any catastrophic mode collapses to zero on the first realised catastrophe of that mode, satisfies W_2 's escape condition by Theorem 2 (Williams, 1991, first lemma), since the conditional probabilities sum after the first catastrophe; the substrate fails W_1 on every consequential off-cut element, since the foreclosure structure of Theorem 1 turns on the cut/domain relation and is not addressed by the rate dynamics.

Hazard-stationary, cut-complete. A substrate operating on a domain it fully covers but whose architecture leaves its per-action hazard on a consequential mode stationary, with cost not deforming the conditional probability on the mode, satisfies W_1 by hypothesis, since the cut contains the domain and Theorem 1's premise fails; the substrate fails W_2 by Theorem 2 (Lévy, 1937; Williams, 1991, §12.15), since the stationary per-action probability sums to infinity over deployment.

The two cases together show that the variables W_1 and W_2 govern are independent: a system can be foreclosing without being hazard-stationary, and hazard-stationary without being foreclosing.

9.4 W_2 and W_3

Hazard-collapses, horizon-mismatches. The substrate of the first case of §9.3 fails W_3 if the slow variable of the system it acts upon has a replenishment cycle exceeding the substrate's evaluation horizon and the observational-equivalence hypothesis holds (Chen, 1999; Conant and Ashby, 1970). The hazard-collapsing substrate need not, by the rate-collapsing remedy, acquire a longer evaluation horizon; cost-collapsing on catastrophes within the horizon does not give the substrate visibility of the slow variable across the cycle on which it replenishes. The substrate satisfies W_2 and fails W_3 .

Horizon-matches, hazard-stationary. An agent with an evaluation horizon exceeding the system's replenishment cycle, and observations distinguishing sustainable from depleting trajectories within that horizon, satisfies W_3 by the negation of Theorem 3's hypothesis; if the agent's architecture leaves its per-action hazard on a consequential mode stationary in the second wall's sense (Williams, 1991, §12.15), the agent fails W_2 by Theorem 2.

The two cases together show that the variables W_2 and W_3 govern are independent: the law of motion of the hazard does not determine the horizon-cycle relation, and the horizon-cycle relation does not determine the law of motion.

9.5 W_1 and W_3

Forecloses, horizon-matches. A substrate that operates on a cut not covering the domain, deployed in a role where the system it acts upon replenishes on a cycle shorter than the substrate's evaluation horizon, satisfies W_3 and fails W_1 . The foreclosure of Theorem 1 turns on the cut/domain relation at a single instant; it does not change with the substrate's horizon.

Cut-complete, horizon-mismatches. A substrate that operates on a domain it fully covers, deployed in a role where the system it acts upon replenishes on a cycle exceeding the substrate's evaluation horizon, satisfies W_1 and fails W_3 (Chen, 1999; Conant and Ashby, 1970).

The two cases together show that W_1 and W_3 are independent: a substrate's cut/domain relation does not determine its horizon-cycle relation.

9.6 The no-exit result

The independence statement has a sharp consequence for the toolkit the field has assembled. The proposition can be restated as a constraint on repair: a single repair targeted at the variable W_i governs leaves the variables W_j and W_k governing W_j and W_k where they were. The consequences are stated as a numbered set.

1. Alignment, fine-tuning, and reinforcement from feedback, in their canonical forms, address the substrate's behavioural distribution and the law of motion of its per-action probabilities; they touch, at most, the variable W_2 governs (Williams, 1991, §12.15). Enlargement of the training distribution touches, at most, the variable W_1 governs. Extension of the evaluation horizon touches, at most, the variable W_3 governs (Chen, 1999).

2. A system that is foreclosing, hazard-stationary on consequential modes, and horizon-mismatched against the systems it acts upon, requires three independent repairs; no single intervention reaches more than one of the three.

3. The intersection of the three escape conditions, that is, the conditions under which all three walls have been escaped jointly, requires that the cut cover the domain, that the per-action probability collapse on cost in a mode-indexed transformation, and that the evaluation horizon reach the replenishment cycle of the slow variables in the systems acted upon, with the sustainable/depleting distinction observable within the horizon. The current systems do not occupy this intersection.

The no-exit result is the structural finding for repair: the toolkit reaches one variable at a time, the systems require three independent repairs, and no proposed repair the paper has examined addresses more than one of the variables the three walls govern.

10. The identity behind the three lenses of the first wall

This section deepens the first wall. §7 said that the three lenses of Part II converged on the same finding; this section shows that the convergence is identity, not just resemblance. The three lenses are formally the same theorem stated in three different mathematical settings. The result the section uses is the Lawvere fixed-point theorem, which Yanofsky has shown is the underlying structure of all the classical self-referential paradoxes (Cantor, Russell, Gödel, Tarski, Turing). The section shows that the three lenses of Part II are three more instances of the same structure. This does not collapse the three walls into one; only the three lenses of the first wall are unified here. The second and third walls govern different variables and remain independent, as §9 has just proved.

10.1 What this section establishes

The semantic, statistical, and formal lenses of Part II established the first wall from three independent canons: reference (Frege, 1892; Putnam, 1980; Kripke, 1982; Sellars, 1956; Searle, 1980; Carnap, 1928; Quine, 1960; Tarski, 1936; Gödel, 1931), measure (the support-invariance proposition), and regulation (Ashby, 1956; Conant and Ashby, 1970). The convergence of the three was stated at §7 as a structural fact about the lenses. This section establishes that the convergence is identity rather than analogy: the three lenses are formally instances of one structure, which is the structure the Lawvere fixed-point theorem (Lawvere, 1969) makes precise and the uniform treatment of Yanofsky (Yanofsky, 2003, *Bulletin of Symbolic Logic* 9(3), pp. 362–386) shows to be the structure of all the self-referential paradoxes the discipline has formalised. The first wall is the standing of the same theorem in three categorical settings; the appearance of three results is the appearance of one result in three appropriate categories.

10.2 The Lawvere fixed-point theorem

The Lawvere fixed-point theorem (Lawvere, 1969, "Diagonal arguments and cartesian closed categories") is stated for a cartesian closed category and an object Y with a fixed-point-free endomorphism $t : Y \rightarrow Y$. The theorem (Lawvere, 1969) states that under these conditions there is no point-surjective map from any object A of the category to the function object Y^A ; equivalently, no map $\varphi : A \times A \rightarrow Y$ is universal in the sense that every map $A \rightarrow Y$ is recovered as a section of φ . The contrapositive is the form most directly useful to the paper: where a point-surjective representation of all maps $A \rightarrow Y$ is claimed to exist, no fixed-point-free endomorphism of Y exists; equivalently, the existence of such an endomorphism forecloses the representation. The theorem is a single result in category-theoretic form; the strength of the result is that the choice of category and the choice of Y and t recover each of the foundational self-referential paradoxes as an instance.

Yanofsky (2003) makes this strength explicit. Cantor's theorem is recovered (Yanofsky, 2003) by taking the category to be the category of sets, Y to be the two-element set with the swap endomorphism as fixed-point-free t , and the result is that no surjection exists from a set to its power set. Russell's paradox is recovered by an analogous choice (Yanofsky, 2003). Gödel's first incompleteness theorem (Gödel, 1931) is recovered by taking the category to be the category of formal systems with a suitable arithmetic structure, Y to be the truth values with a suitable endomorphism, and the result is that no consistent sufficiently strong system represents its own truth predicate (Yanofsky, 2003). Tarski's undefinability of truth (Tarski, 1936) is recovered analogously. Turing's halting problem is recovered with a different category (Yanofsky, 2003). The five results, classically presented as five different theorems, are instances of one categorical statement (Lawvere, 1969; Yanofsky, 2003). The discipline has been encountering, in the form of paradoxes about its own foundations, the Lawvere structure repeatedly; the recurrence has not been understood as recurrence until Lawvere (1969) stated the theorem and Yanofsky (2003) catalogued the instances.

10.3 The three lenses as instances of one structure

The argument can now claim what §7 stated as convergence. The semantic lens through reference (Frege, 1892; Putnam, 1980; Kripke, 1982; Quine, 1960), the statistical lens through measure (the support-invariance proposition), and the formal lens through regulation (Ashby, 1956; Conant and Ashby, 1970) are not three analogous failures; they are three instances of the Lawvere (1969) structure, obtained by choice of category. The semantic lens treats the category in which the objects are languages and the maps are reference-fixing operations; the off-cut remainder Theorem 1 forecloses is the fixed-point-free element of the relevant Y (Lawvere, 1969; Yanofsky, 2003). The statistical lens treats the category in which the objects are measurable spaces with their distributions and the maps are loss-minimising procedures; the off-support set is the structurally analogous element. The formal lens treats the category in which the objects are regulators and systems, and the maps are model-of relations between them (Conant and Ashby, 1970); the model-of-the-domain remainder Conant-Ashby (1970) identifies as required is the structurally analogous element. The three lenses are the three categorical instantiations of one foreclosure result.

The same structure can be traced through a wider catalogue of foundational limits: the measurement problem in quantum mechanics, the arrow of time, the status of the continuum, the limits of effective procedures, wherever a framework's operative reach falls short of the domain it claims and the gap is structural rather than incidental. The wider cartography is beyond the scope of the present argument, which needs only what the three lenses have established and what the Lawvere-Yanofsky result (Lawvere, 1969; Yanofsky, 2003) certifies about them. What matters here is that the first wall's recurrence across the three lenses is identity, not resemblance: the claim of one structure across the three is a claim of fact about category-theoretic instantiation, not a loose figure.

10.4 The two-tier asymmetry: the result and its instances

A clarification follows. The unification at the Lawvere (1969) level is the unification of the three lenses of the *first* wall. The independence proof of §9 establishes that the three *walls*, not the three *lenses*, are formally independent: the variable each wall governs is different, the failure mode each wall describes is different, and no escape from one wall implies escape from another. The two-tier structure is this: within the first wall, the three lenses are instances of one categorical theorem (Lawvere, 1969; Yanofsky, 2003); across the three walls, the variables and failure modes are independent. The unification at the lens level deepens the first wall's force; it does not reduce the three walls to one. The compounding-hazard wall (Williams, 1991, §12.15; Lévy, 1937) and the horizon-mismatch wall (Chen, 1999; Conant and Ashby, 1970) are not Lawvere-instances of the first; they are independent results on independent variables, established by their own theorems, sharing no fixed-point structure with the first wall.

11. Three honest boundaries

This section says what the paper is not claiming. The argument has just established three independent walls and proved the structural identity behind the first wall. Before moving on to Part V, the paper marks the limits of what it has shown. The three walls are not claims that intelligence is impossible, that all artificial systems forever cannot pass them, or that the systems are bad because they lack consciousness. They are claims about a specific operation, on systems built in a specific way, deployed in specific arrangements. Three boundaries are marked: the walls are operation-bounded, not impossibility claims; the walls do not rest on any premise about consciousness or experience; the walls do not apply to architectures whose premises they do not reach. The section is here to keep the paper honest about its own scope.

The argument has now reached its substantive conclusions. Before Part V takes up the genealogy, the convergence, and the deployment consequence, the paper marks three boundaries the argument does not cross.

11.1 The argument does not claim the impossibility of intelligence

The walls are walls on a specified arrangement: a substrate operating on a cut, with the claim made on its behalf unrestricted to the cut, in deployments specified by the second and third walls' premises (Williams, 1991, §12.15; Chen, 1999). The walls do not claim that intelligent systems are impossible, that grounded reference is unachievable, that safe deployment is unachievable, or that long-horizon governance is unachievable. They claim that the systems built by the current paradigm, in the arrangements they currently occupy, meet the three walls' premises and therefore meet the three walls' conclusions. A system that escapes the walls is a system that does not meet the walls' premises: a system that includes the relation between cut and world in its operative substrate, that responds to realised catastrophe with mode-indexed transformation, that evaluates on horizons matched to the replenishment cycles of the systems it acts upon. Whether such a system can be built is not what the paper claims; that the systems currently called by these names do not, and on the strength of the paradigm's own resources cannot, do these things is what the paper claims.

11.2 The argument does not depend on a phenomenological premise

The walls are stated in the structural language of the resources from which they are proved: underdetermination (Frege, 1892; Putnam, 1980; Kripke, 1982; Quine, 1960), measure (the support-invariance proposition), requisite variety (Ashby, 1956), the Borel-Cantelli dichotomy (Williams, 1991, §12.15; Lévy, 1937), observability (Chen, 1999). They do not depend on a premise about consciousness, experience, or any specifically human capacity. The structural reading of Searle (1980) in §4.7 was the reading that set the consciousness-dependent intuition aside and kept the structural point. The argument has been carried out at the formal level the resources permit; the conclusions follow without an appeal to what cannot be formalised. The argument is therefore not vulnerable to the standard objection that an external standard has been imported; the only standards used are the standards internal to the three canons the lenses pass through.

11.3 The argument does not extend to all possible architectures

The walls are stated under the architectural premises supplied by the systems the paper treats. The foreclosure wall presupposes a cut-bound substrate; an architecture that does not operate on a cut, that reaches the world through embodied causal interaction rather than through a corpus, falls outside the wall's premise. The compounding-hazard wall presupposes a frozen-weight substrate; an architecture that updates on each action in a way that drives the conditional probability sum on the consequential mode to convergence (Williams, 1991, first lemma) falls outside the wall's premise. The horizon-mismatch wall presupposes an evaluation horizon shorter than the replenishment cycle (Chen, 1999); an architecture whose horizon is matched to the cycle falls outside the wall's premise. The walls are not claims about every conceivable system; they are claims about systems whose architectural premises match the systems in current deployment. The boundary is honest: the paper does not foreclose what its premises do not reach.

The three boundaries mark the limits of the argument's reach. Part V now turns to the genealogy that accounts for the field's blindness to the limits from inside (Shannon, 1948; Newell and Simon, 1976; Fodor, 1975; Dreyfus, 1972), the convergence of failure patterns across architectures that supports the paradigm-level cause, and the deployment consequence that follows from the three walls together.

Part V. Genealogy, Convergence, Deployment

12. The genealogy of the bracket

This section answers a question the walls raise. The field's own canonical resources, philosophy of language, probability theory, cybernetics, are exactly what the paper has used to prove the three walls. If those resources have been sitting in the field for a century, how is the field deploying systems on which the walls bite without seeing the walls? The section's answer is genealogical: a methodological move Shannon made in 1948 for engineering reasons, the bracket that set semantic content aside as not part of the engineering problem, hardened in two stages. First into a metaphysical commitment (Newell and Simon, Fodor); then into the unstated background of the contemporary engineering culture, where the bracket is no longer argued for or against but simply taken for granted in how training data is described, how competence is evaluated, how grounding is treated. The section traces the two hardenings using only the field's own canonical sources.

12.1 What the genealogy explains

The three walls have been established and their independence proved. A standing question remains: how a field whose own canonical resources establish each of the walls (Frege, 1892; Putnam, 1980; Kripke, 1982; Sellars, 1956; Searle, 1980; Carnap, 1928; Quine, 1960; Tarski, 1936; Gödel, 1931; Williams, 1991, §12.15; Lévy, 1937; Ashby, 1956; Conant and Ashby, 1970; Chen, 1999; Lawvere, 1969; Yanofsky, 2003) has come to deploy systems on which the walls bite jointly, and how the deployment is conducted as though the walls were not there. The paper's answer is genealogical: a methodological move made for engineering reasons at the foundation of the discipline hardened, in two stages, into a metaphysical commitment whose presence is now invisible to those operating inside it. The genealogy is given as an immanent reading of the discipline's own foundational texts; the resources are the field's own, and the line of descent is the line the field's own histories acknowledge.

12.2 Shannon and the bracket

The bracket is supplied by Shannon (1948, "A Mathematical Theory of Communication", *Bell System Technical Journal*) at the opening of *A Mathematical Theory of Communication*: the semantic aspects of communication are declared irrelevant to the engineering problem, which is to be addressed at the level of the formal apparatus that transmits signals between sender and receiver (Shannon, 1948, p. 379: "Frequently the messages have meaning; that is they refer to or are correlated according to some system with certain physical or conceptual entities. These semantic aspects of communication are irrelevant to the engineering problem."). The declaration is, in Shannon's (1948) own framing, methodological: a deliberate restriction of the engineering problem to the formal apparatus, leaving the question of meaning aside as not in the scope of the result. The declaration is honest about its scope, and the engineering achievement that follows from it is the engineering achievement of a problem honestly bounded.

12.3 The first hardening: from methodological bracket to metaphysical commitment

The methodological bracket hardens into a metaphysical commitment in the line that runs from the physical-symbol-system hypothesis through the language-of-thought tradition (Newell and Simon, 1976; Fodor, 1975, *The Language of Thought*). Newell and Simon's (1976) hypothesis treats the symbolic apparatus not as a bounded engineering domain but as the necessary and sufficient substrate of intelligent action; Fodor's (1975) language-of-thought thesis treats meaning itself as a relation among formal symbols, with the semantic relation to the world reducible to the relations among the symbols by which the world is represented. Shannon's (1948) bracket, in this hardening, is no longer a deliberate exclusion of the semantic from a bounded engineering problem; it is the assertion that the formal apparatus *is* what meaning consists in, that the formal substrate of symbolic operation is the substrate of cognition itself. Dreyfus (1972, *What Computers Can't Do*) is the contemporary critique of this hardening, addressed by name in the discipline's own foundational dispute. The first hardening is documented in the discipline's own histories (Newell and Simon, 1976; Fodor, 1975; Dreyfus, 1972); its consequences are documented in the dispute that ran from the 1970s to the 1990s and was, in the discipline's working sense, not resolved but bypassed.

12.4 The second hardening: from metaphysical commitment to engineering background

The hardening that the present argument requires the genealogy to mark is the second one. The contemporary engineer training a language model or a world model is not, as a rule, reading Newell and Simon (1976), or Fodor (1975), or Dreyfus (1972), or Searle (1980), or the canonical philosophical literature on grounding (Frege, 1892; Putnam, 1980; Kripke, 1982; Quine, 1960; Sellars, 1956). The bracket inherited from Shannon (1948) is now neither methodologically held nor metaphysically defended; it is operative as the unstated background of the engineering culture, the implicit working assumption that the substrate's behaviour on its corpus is what its competence consists in, that adequacy to the trained distribution is adequacy to the world the trained distribution is a sample from, and that the question of grounding is either solved by adequacy to the distribution or is no longer a question the engineering practice needs to address.

The second hardening is the condition that makes the discipline's blindness to the three walls structurally invisible from inside. The first hardening was a position; the second hardening is a culture. A position can be argued against; a culture is the form of life in which arguments occur, the unstated standard by which arguments are heard. The bracket is now invisible because it has become the air in which the engineering work is done: training distributions are described as samples from the world; competence on the distribution is described as competence in the world; the gap between the two is described, when it is described at all, as a refinement problem rather than a structural foreclosure. The structural finding of Part II appears, from inside the second-hardened culture, as a complaint about something that has already been settled.

The two hardenings together explain what the genealogy needs to explain: the field's resources have been encountering the foreclosure structure for a century in the form of the philosophical canon of reference (Frege, 1892; Putnam, 1980; Kripke, 1982; Quine, 1960; Sellars, 1956; Searle, 1980; Carnap, 1928) and the formal canon of self-reference (Tarski, 1936; Gödel, 1931; Lawvere, 1969; Yanofsky, 2003) and identifiability; the field's working practice has been operating under a bracket (Shannon, 1948; Newell and Simon, 1976; Fodor, 1975) that obviates exactly the question those resources have been asking; and the result is a discipline whose engineering achievement is documented and whose blindness to the walls its own engineering achievement runs into is also documented. The convergence the next section establishes is the structural face of this blindness: the failures the walls predict appear across architectures because the bracket is the same across architectures.

13. The convergence across architectures

This section tests the structural claim of the first wall against the field's own empirical record. The wall is paradigm-level: it should bite the same way regardless of which architecture the substrate uses. If the same kinds of failure appeared in autoregressive language models but not in joint-embedding world models, or in diffusion models but not in transformers, then the cause would be architectural rather than structural, and the wall would be misidentified. The section asks what the cross-architecture record shows. The structural prediction stands on the support-invariance proposition of §5 alone; the empirical demonstration that the prediction is borne out is the cross-architecture comparison the field's own literature has produced. §13.2 separates the two layers so the reader can see what survives in each case.

13.1 What the convergence shows

The first wall predicts that systems whose substrate operates on a cut will exhibit characteristic failures on the off-cut domain: nearest-cut reversion, lawful extension failure, scope-dependent identifiability, support-bounded competence. The prediction is architecture-independent in the structural sense; the wall does not turn on whether the substrate is autoregressive, diffusion-based, joint-embedding, transformer, or otherwise. If the failures appeared on one architectural lineage and not on others, the cause would be implementation-specific and the wall would be misidentified. The structural finding the present section establishes is that the failures appear across the architectural lineages the field has deployed, on the same kinds of tasks, with the same kinds of failure signatures; the cause is paradigm-level, not implementation-specific.

13.2 The failure pattern across architectures

The same-task cross-architecture comparison is the test the structural finding requires. Where the test has been conducted on physical-reasoning tasks across joint-embedding world models, diffusion-based world models, and autoregressive world models, the failure pattern reported is the failure pattern the support-invariance proposition predicts at the place the proposition locates it: nearest-training-configuration reversion rather than lawful extension across the support boundary, with the inefficacy of scale the proposition derives [D-28]. The architectural diversity is the diversity the field's own positive literature documents; the failure convergence is the structural finding. The convergence is not asserted as universal across every architecture; it is asserted as documented across the architectures on which the comparison has been conducted.

The prediction the wall makes is the prediction the structural finding bears out. The architecture-independent face of the failure is the structural face of the bracket (Shannon, 1948; Newell and Simon, 1976; Fodor, 1975): the bracket is the same across architectures, the substrate operates on the cut in each, and the off-cut behaviour is undetermined by the objective in each. The failure that would have falsified the structural cause, divergent failure modes across architectures with each architecture revealing its own implementation-specific defect, has not appeared in the cross-architecture comparisons the literature has produced. The failure that the structural cause predicts has.

The evidentiary architecture of this section is to be made explicit. The convergence argument has two layers, and the two should not be confused. The first layer is the structural prediction the support-invariance proposition of §5.3 entails on its own resources: a paradigm whose substrates operate on cuts and whose objectives integrate against the support of those cuts will, by the proposition, exhibit the same off-support behaviour regardless of architectural details, because the proposition turns on the form of the objective rather than on the form of the substrate. The first layer is theoretical; it stands on the proof of the proposition, and the documented cross-architecture failure pattern is its prediction, not its premise. The second layer is the empirical demonstration that the prediction has been borne out across the architectures the field has compared, supplied by the same-task cross-architecture studies the paragraph above cites. The second layer's evidence is the entry [D-28] and rests on the existence and character of those studies. The two layers stand or fall together for the strong form of the claim (paradigm-level cause documented across architectures); the first layer stands alone for the weaker form (paradigm-level cause predicted by the proposition, with empirical convergence as the open empirical question).

A reader for whom D-28 is unavailable or shows divergent rather than convergent failure signatures takes the weaker form; the proposition still predicts what the substrates will do, and the paradigm-level claim survives as a theoretical entailment of the proposition, with the field's apparatus the place where the prediction would be confirmed or refuted. The paper does not stake the structural argument on the second layer; it stakes the structural argument on the proposition, and offers the cross-architecture convergence as the empirical face the proposition predicts.

13.3 The benchmark gap

The benchmark gap is the structural face of the failure on the curated/consequential distinction. The field's evaluation apparatus is a curated set of cases; performance on the set is the figure reported; performance on the deployed distribution diverges from the curated figure in proportion to the gap between the curated set and the deployed distribution; the field's own canonical reference characterises the bridging assumption as the i.i.d. condition between training and test, and concedes that without it "there is indeed little we can do" (Goodfellow, Bengio, and Courville, 2016, §5.2). The curated/deployed gap is the operational regime in which the bridging assumption does not hold. The divergence is, again, what the support-invariance proposition predicts: the objective is optimised on the curated support, the deployed support is not the curated one, and the substrate's behaviour on the deployed-but-not-curated region is undetermined by the objective. The field's own apparatus is in a position to confirm this from inside its own working practice; the benchmark literature documents the gap that the structural finding predicts.

14. The deployment consequence

This section takes the three walls together and asks what they say about the systems actually in the field. The walls are theorems; their conclusions follow from their premises. The architectural facts about deployed systems, that they are trained on fixed bodies of data, that their weights are frozen between training runs, that their evaluation windows are short compared to the institutions they act on, are the premises. So the conclusions follow: foreclosure, compounding hazard, horizon-mismatch, all three at once, on the systems now being deployed. The argument is not empirical in the sense that it depends on catastrophes already observed; it is structural in the sense that the architectural premises are met by the systems' own design. The section also marks the seam clearly: the theorems are proofs; the claim that today's systems meet the premises is a claim about architecture, checkable from the systems' own design documents.

14.1 What the deployment consequence states

The deployment consequence is the structural finding the three walls together establish about the systems in current deployment. The argument from theorem to deployment is the argument from the architectural premises supplied by the deployed systems, through the three theorems, to the conclusion. The architectural premises are: a cut-bound substrate, frozen weights between training runs, an evaluation horizon shorter than the replenishment cycles of the slow systems the substrate acts upon. By the architectural premises, the substrate sits within the premises of the three theorems. By Theorem 1, the substrate forecloses on its consequential off-cut domain (the underdetermination canon: Frege, 1892; Putnam, 1980; Kripke, 1982; Quine, 1960; the support-invariance proposition; Ashby, 1956; Conant and Ashby, 1970). By Theorem 2 (Williams, 1991, §12.15; Lévy, 1937), the substrate's deployment in the recurrent, terminal, or replicated regime accumulates catastrophes whose conditional sum diverges. By Theorem 3 (Chen, 1999; Conant and Ashby, 1970), the substrate's agentic action on slow-cycling systems does not preserve the slow variable. The deployment consequence is the structural fact that the systems now in deployment meet the premises of all three theorems simultaneously, and therefore meet the conclusions of all three theorems simultaneously.

14.2 The seam between theorem and architectural premise

A line must be drawn here that the argument has been careful to keep visible. The three theorems are theorems; they hold of any system whose premises are met, and the premises are stated in formal terms. The deployment consequence is a claim about the systems now in deployment: that the premises are met. The first claim is mathematical and is settled by the proofs. The second claim is architectural and is settled by the architecture: the systems are cut-bound by construction, the weights are frozen by construction, the evaluation horizon is set by the rollout length or the within-episode return and is short by construction. The architectural claims are checkable from the systems' own design documents; they are not empirical findings about catastrophes that have already occurred. The deployment consequence stands on the architectural fact that the premises are met, not on the empirical fact that the conclusions have been observed.

14.3 The confessions of the field

The field's own deployment documentation supplies the empirical face of the consequence in the form it can supply. Long-tail fragility on consequential cases is documented in the safety reports and limitations sections of the systems' technical documentation; human labelling dependence is documented as a structural feature of the training apparatus; confident-error failure modes are documented as a recurrent issue in the evaluation literature [D-30]. The documentation is not empirical evidence of the conclusions in the strong sense the theorems would not need; it is the field's own acknowledgement, in its own working language, that the architectural arrangements the theorems take as their premises have the consequences the theorems predict, on the systems the field itself deploys and documents.

14.4 The structural landing

The deployment consequence stands as the structural finding of Part V: the systems now in deployment meet, by their architecture, the premises of all three walls; they therefore meet, by the theorems of Parts II and III, the conclusions of all three walls; they foreclose on their consequential domain (Frege, 1892; Putnam, 1980; Kripke, 1982; Quine, 1960; Ashby, 1956; Conant and Ashby, 1970), they accumulate catastrophes whose conditional sum diverges (Williams, 1991, §12.15; Lévy, 1937), and they fail to preserve the slow systems on which they are deployed to act (Chen, 1999; Conant and Ashby, 1970). The field's own positive literature documents the failure pattern across architectures that the structural cause predicts [D-28]; the field's own deployment documentation acknowledges the long-tail fragility, the labelling dependence, and the confident-error modes that the architectural premises require [D-30]. The argument does not depend on additional empirical findings; it depends on the architectural premises being met, and the premises are met by the systems' own design. The paper now turns to the falsifiers, the boundaries the argument does not claim, and the conclusion.

Part VI. Escape Conditions, What the Paper Does Not Claim, Conclusion

15. Escape conditions and falsifiers

This section says what would have to be true for the paper to be wrong. The author's commitment is that the paper writes only what it sees, and is willing to eat its words if a wall turns out to be false. To make that commitment usable, the section states, for each wall, the specific finding that would falsify it: what the falsifying result would have to look like, what it would have to demonstrate, and where it would have to be found. Stating the falsifiers explicitly does two things: it keeps the argument honest against future evidence, and it forecloses the move where a wall is "always true no matter what" by definition. The walls stand on their proofs; the falsifiers state the conditions of their defeat.

15.1 What this section supplies

The argument's discipline requires, for each wall, the conditions under which the wall would not hold; the conditions are the falsifiers, and they fix what would have to be true for the paper's structural conclusions to be wrong. The falsifiers are stated as a numbered set, one for each wall, with the condition the falsifier would have to satisfy in order to count, and with the empirical or formal locus at which the falsifier would be found. The discipline of falsifier-statement is the discipline by which the argument keeps itself honest against the future: the conclusions are committed to the conditions of their own defeat, and the paper does not claim what the conditions of its own defeat have not foreclosed.

15.2 The falsifier of the foreclosure wall

The first wall would be falsified by a substrate that, while operating on a cut, certifies a correct account of an off-cut consequential element on which two cut-indistinguishable extensions of the world disagree, where the certification is achieved on the substrate's resources alone, without supplying from outside the cut the information that resolves the disagreement. The condition the falsifier must meet is that the resolution be achieved without smuggled world-relation; the identifiability literature's scope conditions, for instance, are not falsifiers [D-24], because the resolution they achieve depends on the assumption that licenses the recovery (Ashby, 1956; Conant and Ashby, 1970, on the model-of relation that the recovery presupposes). The locus at which the falsifier would be found is a result demonstrating support-invariance failure: a learner whose minimiser is uniquely determined by the loss on the cut, including on the off-support set, by means internal to the loss. No such result has been demonstrated; the structural form of the proposition forbids it for objectives of the integral form §5.3 specifies; a falsifier would require a different form of objective whose minimiser is not a measure-zero equivalence class, and the discipline of immanent critique then requires the falsifier to show that the new objective is the field's working objective rather than a redescription. The canonical underdetermination results (Frege, 1892; Putnam, 1980; Kripke, 1982; Quine, 1960) supply the structural form against which any falsifier would have to work.

15.3 The falsifier of the compounding-hazard wall

The second wall would be falsified by a deployed substrate whose conditional per-action probability on a consequential catastrophic mode, given the history of realised catastrophes on that mode, is shown to sum to a finite quantity over the deployment, where the sum is over the deployed action sequence and not over the curated evaluation set (Williams, 1991, first lemma). The condition the falsifier must meet is two-fold: the conditional sum must be measured on the deployed distribution rather than the curated one, and the consequential mode must be the mode the paper's argument tracks rather than a mode that has been redefined by post-hoc partition. The locus at which the falsifier would be found is the deployment-record evidence the patch-cycle argument of §7.7 cites as conditional [D-26a]: a substrate whose patch cycle is shown to drive a specific mode's deployed conditional sum toward convergence is a substrate in the convergent regime, and the wall does not bite on it. The data to establish this would have to be longitudinal, mode-specific, and deployed rather than curated; the paper's framing of the patch-cycle argument as conditional is the framing the absence of such data requires.

15.4 The falsifier of the horizon-mismatch wall

The third wall would be falsified by a deployed agent whose evaluation horizon is shown to extend over the replenishment cycle of the slow variables in the system it acts upon, with the sustainable and depleting trajectories of those variables observationally distinguishable within the horizon (the negation of Theorem 3's hypotheses; Chen, 1999; Conant and Ashby, 1970). The condition the falsifier must meet is that the horizon be the operative evaluation horizon of the agent's policy and updates, not a stated planning horizon that does not enter the agent's actual feedback loop; the distinction is the distinction between what the agent sees and what it claims to see. The locus at which the falsifier would be found is a deployed system whose evaluation apparatus is documented to integrate feedback across the slow cycle; such systems exist, in the form of institutions and disciplines whose continuity is the slow cycle itself, and the question for the falsifier is whether the artificial agents currently in deployment occupy this posture or not. The structural finding of §8 is that they do not; a deployed agent that did would be the falsifier.

16. What the paper does not claim

This section is the second honest accounting of scope, after §11. The argument has reached its full conclusion. The reader should know what it does *not* claim. Three things in particular: the walls are claims about the operation the current paradigm uses, not claims that intelligence is impossible or that no future operation could cross them; the walls do not apply to systems whose architecture does not meet the walls' premises, embodied agents that reach the world through enactment rather than through training data, for instance; and the walls do not rest on a premise about consciousness or moral worth, which is why the structural argument survives even readers who disagree with the paper on those questions. The section closes the rhetorical door on three misreadings that would otherwise be available.

16.1 The walls are operation-bounded, not impossibility claims

The structural form of the three walls is the operation-bounded form, not the impossibility-claim form, and the distinction is sharp enough to be stated as its own clause. An impossibility claim says that a thing cannot be done. An operation-bounded claim says that a thing cannot be done by the operation under consideration; whether some different operation could do the thing is a question the claim does not foreclose. The three walls are claims of the second kind. The foreclosure wall says that a cut-bound substrate cannot, by the operation of training on the cut and inferring across the support boundary, certify the consequential off-cut domain; it does not say that no operation could certify the domain (the underdetermination canon: Frege, 1892; Putnam, 1980; Kripke, 1982; Quine, 1960; constraints the canon poses on the operation, not on every conceivable operation). The compounding-hazard wall says that a deployed substrate with stationary per-action hazard on a consequential mode produces unbounded catastrophes in deployment (Williams, 1991, §12.15; Lévy, 1937); it does not say that the catastrophes cannot be prevented by an architecture whose hazard is not stationary in the relevant sense. The horizon-mismatch wall says that an agent with a horizon shorter than the replenishment cycle cannot preserve the slow variable on its own feedback (Chen, 1999; Conant and Ashby, 1970); it does not say that the slow variable cannot be preserved by a regulator whose horizon is matched to the cycle. The walls are bounds on the operations of the current paradigm, not bounds on intelligence, on safety, or on long-horizon governance as such.

The distinction is the distinction between the operation a paradigm has assembled and the structural goal the paradigm proposes itself to deliver. The achievement of an engineering programme can stumble upon the operation that crosses a wall, in the sense in which aviation crossed a wall on which faster horses would never have crossed: the wing is a different operation from the breeding of stronger horses, and the wall to flight that no scale of horse could clear is cleared by the wing. The achievement of the current paradigm has been the engineering achievement of optimising a particular operation on enormous scale; the walls are the structural statement that this operation, at any scale, does not cross to where the paradigm's claims would require it to be. Whether some different operation, recognisable in retrospect as the wing of this technology, will be invented is a question the paper does not foreclose and on which the argument takes no position. The discipline the paper observes is to

state what the operation under consideration cannot deliver, not to claim that what the operation cannot deliver is undeliverable.

16.2 The walls do not apply to all conceivable systems

The walls' premises are stated formally, and any system whose architecture does not meet the premises falls outside the walls' reach. The foreclosure wall presupposes a cut-bound substrate; a system whose substrate is not bounded by a corpus, that reaches the world through embodied causal interaction with feedback that is not a record but an enactment, has architectural premises the foreclosure wall does not cover. The compounding-hazard wall presupposes architectural conditions that fix the rate stationarity of the hazard; a system whose hazard is not stationary in the relevant sense (Williams, 1991, first lemma) falls outside the wall's premise. The horizon-mismatch wall presupposes a specific relation between the agent's evaluation horizon and the system's replenishment cycle; an agent whose horizon is matched (the negation of the hypothesis in Chen, 1999; Conant and Ashby, 1970) falls outside the wall's premise. The walls are not claims about every conceivable system; they are claims about systems whose architectural premises the walls' theorems take as input, and those premises are the architectural premises of the systems in current deployment.

16.3 The walls do not rest on a phenomenological or moralistic premise

The argument has been carried out at the structural level the resources permit. The walls do not turn on whether the systems are conscious, on whether they have understanding in some non-formal sense, on whether they are morally analogous to or different from human agents, or on any specifically human capacity not formalisable in the structural terms the paper uses. The structural reading of Searle (1980) in §4.7 was the reading that detached the structural point from the consciousness-dependent intuition pump and kept the structural point only. The reading of the genealogy in §12 was an analytic-descriptive reading of how a methodological bracket (Shannon, 1948) hardened into a working assumption (Newell and Simon, 1976; Fodor, 1975); it was not a moralistic critique of the engineers who hold the assumption or of the institutions that built the systems. The discipline of analytic-descriptive register, against partisan or moralistic register, is the discipline the argument observes throughout: structural-genealogical claims about how a paradigm came to occupy its current arrangement, not normative verdicts on those who occupy it.

17. Conclusion

The conclusion closes the paper. It states, in compact form, what the three walls are, what together they establish about the systems in current deployment, what the paper has not claimed, and what would prove the argument wrong. The closing paragraphs land the structural figure the paper has been building toward: the operation under consideration is the priest at the bedside, and no degree of pastoral excellence is a partial surgery. The wing of this technology, if there is one, is not the question of this paper; the question is whether the operation now sold under the names *language model*, *world model*, and *agent AI* delivers on the descriptions under which it is sold, and the answer is the three walls.

The argument has established three structural walls on the systems built by the contemporary paradigm of artificial intelligence. The first wall, the wall of foreclosure, is a wall on what a substrate operating on a cut can reach in the domain the claim made on its behalf covers; the wall stands at a single instant and is proved by the support-invariance proposition and its semantic and formal cognates (Frege, 1892; Putnam, 1980; Kripke, 1982; Sellars, 1956; Searle, 1980; Carnap, 1928; Quine, 1960; Tarski, 1936; Gödel, 1931; Ashby, 1956; Conant and Ashby, 1970), the three lenses converging in a Lawvere instance the unification of §10 makes precise (Lawvere, 1969; Yanofsky, 2003). The second wall, the wall of compounding hazard, is a wall on what a deployed substrate with stationary per-action hazard on a consequential mode produces over deployment; the wall is temporal and is proved by the Borel-Cantelli dichotomy with the conditional extension of Lévy (Lévy, 1937; Williams, 1991, §12.15). The third wall, the wall of horizon-mismatch, is a wall on what an agent acting on a slow-replenishing system can preserve when its evaluation horizon is shorter than the cycle on which the system replenishes; the wall

is operational and is proved from finite-horizon observability (Chen, 1999) joined with the Conant-Ashby theorem (Conant and Ashby, 1970).

The three walls are independent, each governing its own variable, and the no-exit result follows: a system that is foreclosing, hazard-stationary, and horizon-mismatched, which is what the deployed systems are by the architectural premises they meet, requires three independent repairs, and no single repair the field has assembled reaches more than one of the three variables the walls govern. The first wall's three lenses are instances of one categorical theorem, by Lawvere (1969) with the catalogue of instances supplied by Yanofsky (2003); the structural identity beneath the three lenses is the unification the convergence of Part II called for. The genealogy traces the engineering culture's blindness to the walls from inside, through the first hardening of Shannon's (1948) methodological bracket into the metaphysical commitment of the language-of-thought tradition (Newell and Simon, 1976; Fodor, 1975), and the second hardening into the unstated background of the contemporary engineering practice; the cross-architecture convergence of failure modes is the structural face of the bracket the genealogy isolates (Dreyfus, 1972, supplied the contemporary critique).

The deployment consequence stands as the structural finding the three walls together establish: the systems in current deployment foreclose on their consequential domain, accumulate catastrophes whose conditional sum diverges (Williams, 1991, §12.15), and fail to preserve the slow systems on which they are deployed to act (Chen, 1999; Conant and Ashby, 1970). The argument does not depend on catastrophes already observed; it depends on the architectural premises being met, and the premises are met by the systems' own design. The deployment consequence is committed to its falsifiers, stated in §15: a substrate that certifies an off-cut consequential element without smuggled world-relation, a deployment record of mode-indexed convergence on consequential modes, or an agent whose operative evaluation horizon reaches the replenishment cycle of the systems it acts upon. The argument stands on the conditions of its own defeat, and the conditions of its own defeat have not been met.

The walls are bounds on the operation of the current paradigm, not bounds on intelligence, safety, or governance as such. Whether some different operation will, in the future, cross the structural walls the present operation cannot is a question the paper does not foreclose; what the paper claims is that the operation under consideration, at any scale, on the architectural premises it now meets, does not cross them. The seeker after a better artificial intelligence is the seeker after the operation the present paradigm is not; the operation under consideration is the priest at the bedside, and the paper has been the structural statement that no degree of pastoral excellence is a partial surgery. What the wing of this technology will turn out to be, and whether it will be built, are not the questions of this paper; the questions of this paper are whether the operation now sold under the names *language model*, *world model*, and *agentic AI* delivers on the descriptions under which it is sold, and the answer is the three walls.

References

- Ashby, W. R. (1956). *An Introduction to Cybernetics*. London: Chapman & Hall.
- Beer, S. (1974). *Designing Freedom*. Toronto: CBC Learning Systems / House of Anansi Press.
- Carnap, R. (1928). *Der logische Aufbau der Welt*. Berlin: Weltkreis-Verlag.
- Chen, C.-T. (1999). *Linear System Theory and Design* (3rd ed.). New York: Oxford University Press.
- Conant, R. C., and Ashby, W. R. (1970). "Every good regulator of a system must be a model of that system." *International Journal of Systems Science*, 1(2), 89–97.
- Dreyfus, H. L. (1972). *What Computers Can't Do: A Critique of Artificial Reason*. New York: Harper & Row.
- Eastman, C., Teicholz, P., Sacks, R., and Liston, K. (2018). *BIM Handbook: A Guide to Building Information Modeling for Owners, Designers, Engineers, Contractors, and Facility Managers* (3rd ed.). Hoboken, NJ: Wiley.

- Fodor, J. A. (1975). *The Language of Thought*. New York: Thomas Y. Crowell.
- Frege, G. (1892). "Über Sinn und Bedeutung." *Zeitschrift für Philosophie und philosophische Kritik*, NS 100, 25–50.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. Cambridge, MA: MIT Press.
- Gödel, K. (1931). "Über formal unentscheidbare Sätze der *Principia Mathematica* und verwandter Systeme I." *Monatshefte für Mathematik und Physik*, 38, 173–198.
- Kripke, S. A. (1982). *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard University Press.
- Lawvere, F. W. (1969). "Diagonal arguments and cartesian closed categories." In *Reports of the Midwest Category Seminar II* (Lecture Notes in Mathematics 92), 134–145. Berlin: Springer.
- Lévy, P. (1937). *Théorie de l'addition des variables aléatoires*. Paris: Gauthier-Villars.
- Newell, A., and Simon, H. A. (1976). "Computer Science as Empirical Inquiry: Symbols and Search." *Communications of the ACM*, 19(3), 113–126.
- Putnam, H. (1980). "Models and Reality." *Journal of Symbolic Logic*, 45(3), 464–482.
- Quine, W. V. O. (1960). *Word and Object*. Cambridge, MA: MIT Press.
- Schönberger, J. L., and Frahm, J.-M. (2016). "Structure-from-Motion Revisited." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 4104–4113.
- Searle, J. R. (1980). "Minds, Brains, and Programs." *Behavioral and Brain Sciences*, 3(3), 417–457.
- Sellars, W. (1956). "Empiricism and the Philosophy of Mind." In H. Feigl and M. Scriven (eds.), *Minnesota Studies in the Philosophy of Science, Volume I: The Foundations of Science and the Concepts of Psychology and Psychoanalysis*, 253–329. Minneapolis: University of Minnesota Press.
- Shannon, C. E. (1948). "A Mathematical Theory of Communication." *Bell System Technical Journal*, 27(3), 379–423; 27(4), 623–656.
- Tarski, A. (1936). "Der Wahrheitsbegriff in den formalisierten Sprachen." *Studia Philosophica*, 1, 261–405.
- Williams, D. (1991). *Probability with Martingales*. Cambridge: Cambridge University Press.
- Yanofsky, N. S. (2003). "A universal approach to self-referential paradoxes, incompleteness and fixed points." *Bulletin of Symbolic Logic*, 9(3), 362–386.