

JEREMY RODGERS

Shadow Theory and Consciousness

Awareness, Perspectival Realization,
and the Source-to-Experience Problem

An integrated physical, mathematical,
phenomenological, and philosophical monograph

Independent Researcher

Website: everythingequation.com

SPC-2: candidate psychophysical constitution

Version 2 — publication edition

20 September 2026

Preface and scope

This monograph asks how awareness is manifested as a bounded perspective and a lived world. Its concern is both formal and human: what distinguishes experience from the processes that support it, how a personal identity is constructed, and what can reasonably be said when perception, memory, or the familiar sense of self changes.

The philosophical interpretation distinguishes reality prior to an operational cut, denoted U , from a working source S and its readout T . Awareness is proposed as the knowing aspect of U , without making it a cosmic person or an observer located inside the brain. A physical vessel conditions a localized subject and its lived scene. The scene, the subject, and awareness are not interchangeable. Neither an autobiographical self nor adult human language is required for every possible form of experience.

The formal account is SPC-2—a candidate psychophysical constitution. A certified realization R^* supplies physical organization and an explicit selection doctrine. A0 states the ontological commitment; A1–A3 assign perspectives, phenomenal relational structure, and experiential continuation. The finite completion theorem establishes conditional determinacy and covariance. Its mathematics does not prove the ontology or certify a real organism merely because its nervous system is active.

For biological readers, the distinction between explicit conscious access and implicit generative processing helps explain how a scene can be organized without every contributing process appearing as a separate content. Dreaming, anesthesia, sensory isolation, and dissociation make the differences between a continuing vessel, present experience, and later memory especially important. These are interpretive applications, not additional premises of the completion theorem. A truly experience-free interval contains neither a lived scene nor a currently instantiated localized experiential subject. Darkness would itself be a scene. The organism and person can remain physically continuous while a genuine qualification gap ends the experiential episode under A3; renewed qualification begins a new episode.

The physical chapters establish how events, records, resources, and retained causal use can be modeled without conscious observation producing the event. The two measurement constitutions retain their separate assumptions. The philosophical chapters explain the proposed meaning of the psychophysical laws and their limits. Together they support a *candidate internal constitutive resolution of the source-to-manifestation problem, under declared assumptions*. Selecting and validating R^* for actual vessels, testing phenomenal adequacy, and obtaining independent verification remain distinct obligations.

Abstract

We develop a source-aspect account in which awareness is the knowing aspect of U , reality prior to the operational source/readout distinction, and physical organization determines localized manifestation. The mathematical working source S is complete only relative to a nominated readout problem; it does not represent U exhaustively. Awareness, localized subject, lived scene, personal identity, conscious access, agency and metacognition are distinguished. Explicit and implicit processing provide an interpretive account of richer biological vessels, without imposing a human cognitive architecture on the finite admission law. Contemplative inquiry and altered experience motivate these distinctions without serving as empirical proof of the ontology. Separate observer-independent quantum-measurement constitutions establish actual histories and records under their own physical assumptions; awareness is not an event trigger.

The central proposal is the Shadow Psychophysical Constitution SPC-2. A certified finite realization fixes physical components, native timing, internal routes, endogenous operations, boundary contracts, resource accounting and process provenance. Minimal joint-output dependence detects causal distinctions absent from single-component marginals. Candidate histories and native instruments are constructed before admission. A compatible covering-return test and nontrivial native prediction determine qualifying recurrent cores under A1.

A2 assigns phenomenal relational organization to the complete endogenous predictive object over all finite continuations. Nested finite-horizon quotients realize its inverse limit on a finite intrinsic carrier. A rational observable-span calculation determines all-future equivalence; an explicit, checked predictive-fibre congruence condition yields actual-class transition instruments. External report channels descend separately when they are constant on predictive fibres at fixed apparatus context. Consequently merely changing a report device or an analyst's forecast cutoff cannot change the assigned content when the native organization is preserved. A3 defines nonbranching process continuation through growth and gradual replacement while treating copying, splitting, merging and genuine qualification gaps explicitly.

The completion theorem gives total admission, a partition of admitted support and episodes, pointed phenomenal structure, grounded report laws and covariance on this certified domain. Finite countermodels delimit its hypotheses. A symmetry obstruction proves that a bare transition kernel need not select a unique primitive decomposition, even among prime binary factorizations. Exact individuation remains sensitive to weak physical coupling and to the supplied realization doctrine. Physical conservativity prevents transcript-only discrimination of the aspect interpretation from an otherwise identical physical reduct. The result is a candidate internal constitutive resolution of the source-to-manifestation problem, under declared assumptions. Actual-vessel realization, phenomenal adequacy and the source-awareness premise remain distinct questions.

Conventions and principal distinctions

$\mathbf{U}$ denotes reality prior to the operational source/readout distinction. The different typeface distinguishes this ontological notation from locally defined evolution operators such as U_t in the mathematics. S is the working source or structure-side cut; T is the readout or record-side cut. The schematic map is $p : S \rightarrow T$. The descent theorems use X for that readout codomain; this is a local notation convention, not another ontological level. Symbols such as T for a time horizon or T_u for a transition kernel retain their expressly defined mathematical meanings.

The older ontological notation Ω_{src} is interpreted here through $\mathbf{U}$, not through a source state in S . Neither symbol names the closure operator $\mathcal{O}$ in $L = \mathcal{O}\Delta\partial[L]$. No theorem in this book treats $\mathbf{U}$ as a set, a Hilbert space, a field, or an argument of a computable transition law.

For probability laws P, Q on the same measurable space,

$$d_{\text{TV}}(P, Q) = \sup_A |P(A) - Q(A)|.$$

For density operators ρ, σ , $D(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1$. Complete-path distance, endpoint distance, quantum-state distance, and faithfulness to a past event are different claims. A bound in one space is transferred to another only through a specified common channel or map.

An *exact theorem* is a mathematical implication on a declared domain. A *physical model* supplies laws of motion, preparation, and interaction. A *psychophysical constitutive law* assigns experiential significance to a physical structure. A *phenomenological report* describes what a subject reports as experienced; its metaphysical interpretation is a separate claim. In particular, a mathematically complete assignment need not be an empirically correct assignment.

The phrase *candidate internal resolution* has one controlled meaning in this monograph: all maps and laws required for the nominated finite constitutive model are specified, and its conclusions follow under their hypotheses. It does not mean that the actual-vessel realization problem, empirical adequacy, universal source selection, or external verification has already been settled.

Terminology

Term	Meaning in this monograph
U	Unsplit ontological prior: reality before an operational source/readout distinction. Not a countable individual or a mathematical state space.
S	Working source, the structure-side description complete relative to the nominated problem.
T	Readout, the restricted record-side description in which physical histories, archives and vessels are represented. Called X in the descent theorem.
Awareness	The proposed knowing aspect of U; not a person, a scene, a report, or an extra force.
Localized subject	An experiential perspective instantiated only during qualification under A1, with a lived scene under A2. A3 relates qualifying occurrences into episodes; a genuine gap ends the episode.
Lived scene	What it is like now: the current phenomenal presentation. Scene does not require vision, spatial imagery, reflective attention, or a narrative self.
Unconscious/ generative organization	Processes contributing to an awareness-bearing vessel without appearing as separately presented contents or possessing a second scene.
Self-model/person	Bodily, practical and autobiographical organization; biological and person-level continuity can persist without a current scene or experiential subject.
Memory of a scene	Retained or reconstructed representation of prior experience; distinct from its occurrence and the continuing organism. Failed recall does not establish that no scene occurred.
Reportability	Availability to an actual reporting channel. It depends on access, memory, motor or linguistic resources, and channel conditions.
Agency	Capacity for organized selection and action; not a synonym for experiential presence.
Metacognition	Representation or assessment of one's cognitive processes; not required for every admitted perspective.
Vessel / R^*	A physically realized organization / its certified specification of components, native instruments, boundaries, resources and process provenance.

The argument and how to read it

The proposal is that a localized experience is the manifestation of intrinsic awareness through a particular realized organization. A vessel shapes what can appear and how its contents can change. Language, an autobiographical self, and reflective self-recognition are further capabilities of some vessels; they are not prerequisites for all experience in this account.

The central result has the following form:

$$\begin{aligned} R^* &: \text{physical organization and native selection doctrine} \\ + \quad &\text{A0–A3: presence, admission, content, continuation} \\ \implies &\text{determinate constitutive assignment.} \end{aligned}$$

The conclusion determines which cores receive perspectives, the relational structure assigned to each actual state, the continuation of episodes, and their physically grounded reports. The arrow is a conditional mathematical implication. The philosophical case concerns the premises and their interpretation; empirical work concerns whether a particular realization and assignment describe the intended vessel.

Question	Contribution to the argument	Reading route
What is manifesting?	Awareness, contents, selfhood, and nonduality; first-person motivation and its interpretation.	Chapters 1–4
What can a vessel physically do?	Observer-independent events, material records, finite resources, and retained causal use.	Chapters 5–9
How is its organization described?	Witnesses, regulation, spectral response, and coherent–dissipative evolution.	Chapters 10–13
What makes a localized perspective?	R^* , predictive profiles, A0–A3, the completion theorem, and explicit finite examples.	Chapters 14–20
What would assess the account?	Comparative applications, independent phenomenal constraints, rival interpretations, and realization obligations.	Chapters 21–24

The logical dependencies are narrower than the breadth of the book. The finite completion theorem requires a specified finite R^* realization and the predictive constructions; it does not require that every vessel be quantum, possess the five O1 response profiles, or instantiate a particular CSCF operator. Those chapters supply physical constructions and mathematical resources with which realizations can be investigated. The two-bit model shows that the central constitution can be evaluated in a finite classical operational setting.

The physical chapters nevertheless have a definite purpose. They distinguish an abstract information variable from a record that is actually written, retained, and reused, and they show how records can arise without an observer supplying the event. This makes the subsequent experiential attribution a stated law concerning a realized process. It also preserves the independence of the two measurement constitutions: each supplies its own physical assumptions and conclusions.

Contents

Preface and scope	i
Abstract	ii
Conventions and principal distinctions	iii
The argument and how to read it	v
 I Awareness, Experience, and Perspective	 1
1 The source-to-experience problem	2
1.1 A question about presence, not merely performance	2
1.2 The unsplit and its operational descriptions	2
1.3 Awareness as a primitive and manifestation as organization	3
1.4 Nonduality and the reality of the readout	4
1.5 Four explanatory obligations	4
1.6 What is distinctive about the proposed account	5
 2 Awareness, contents, and the personal model	 6
2.1 A vocabulary that does not settle the question in advance	6
2.2 The red object and the interpretation of seeing	6
2.3 Memory, scaffolding, and continuity	7
2.4 Conscious access, generative processing, and the lived scene	7
2.5 Vessel, scene, and memory	8
2.6 Minimal presence and higher-order reflection	9
 3 Contemplative inquiry and its evidential role	 10
3.1 Reading <i>I Am That</i> as a philosophical source	10
3.2 Identification is not the deletion of cognition	10
3.3 Witnessing, nonduality, and the limits of the translation	11
3.4 First-person investigation and public evidence	12
 4 Altered identity and unfamiliar experience	 13
4.1 An autobiographical example	13
4.2 Reality, veridicality, and experiential occurrence	13
4.3 Altered weighting of self and world	14

4.4	What remains constant?	14
II	Source, Physical Events, and Material Records	15
5	Source/readout foundations	16
5.1	Nominated sources and restricted descriptions	16
5.2	Predictive closure	17
5.3	Multiple apertures and non-reconstructibility	17
5.4	The Everything Equation and constitutive selection	17
6	Quantum currents and observer-independent events	18
6.1	The coherent source	18
6.2	Predictive current coordinates	19
6.3	The independence of physical events from conscious observation	19
6.4	Complete experiments and comparison spaces	20
7	A massive configuration and record constitution	21
7.1	Guidance, equilibrium, and actual histories	21
7.2	An exact moving Gaussian pointer	21
7.3	Capture, pending states, and null instruments	22
7.4	Historical archives and reset receivers	24
7.5	Autonomous realization and historical error	24
7.6	What the constitution establishes	25
8	A deterministic pilot medium and a complete path limit	26
8.1	The declared microscopic inventory	26
8.2	Bond ownership and conservative export	26
8.3	Spatial preparation produces timing	27
8.4	From signed queues to the Bell path	28
8.5	A static finite Hamiltonian for the whole programme	28
8.6	The observer-independent conclusion	29
9	Resources, continued operation, and agency	30
9.1	Finite records are physical resources	30
9.2	Operational capacity	30
9.3	Viability and the orientation of repair	31
9.4	Agency and stable decisions	31
9.5	Physical stability and philosophical relevance	32
III	The Mathematics of Realized Vessels	33
10	Internal representation and self-referential records	34
10.1	Representing a specified target	34
10.2	Centralizers and their scope	35
10.3	Positive witnesses and finite calibration	35

10.4	Local jets: finite existence without a universal order bound	36
10.5	Self-reference without phenomenal promotion	36
11	Regulation, unity, and the functional self	37
11.1	A typed regulation architecture	37
11.2	Access and valuation are different	37
11.3	Weak unity, strong unity, and assembly	38
11.4	Functional selfhood and episode identity	38
12	Finite incidence, recurrent response, and physical access	39
12.1	Source-sensitive response carriers	39
12.2	Five exact profiles	40
12.3	Interaction-relative sector invariance	40
12.4	Actuation, source drift, and seed-specific mobility	41
12.5	Robust spectral access	42
13	Fixed points, canonicalization, and coherent–dissipative structure	43
13.1	A local contraction theorem	43
13.2	Observable certification	44
13.3	One operator, two analytic regimes	44
13.4	A shared spectral carrier with a general coherent symbol	45
13.5	What a common generator does not determine	45
13.6	The correct role of vessel mathematics	46
IV	A Finite Psychophysical Constitution	47
14	Why a manifestation law is additional	48
14.1	Common source does not determine a bridge	48
14.2	Regularity and symmetry leave genuine freedom	48
14.3	Report laws and latent descriptions	49
14.4	Quantitative aperture adequacy	49
15	The realized domain and its boundaries	50
15.1	Physical data, constitutive selection, and analysis	50
15.2	A physical interface doctrine before candidate selection	50
15.3	Joint causal dependence	51
15.4	Native recurrence and admission	52
15.5	Limits of canonical selection	53
16	Endogenous predictive structure across horizons	56
16.1	An intrinsic carrier and a native operational contract	56
16.2	All finite futures, with no privileged truncation	57
16.3	The transition closure that prediction alone does not prove	58
16.4	An evaluable all-future object	59
16.5	External reports and exact descent	60

16.6	The exact one-step specialization	60
16.7	Why successor-state geometry is a specialization	61
16.8	Stability of finite predictive contrasts	61
17	The Shadow Psychophysical Constitution	63
17.1	Awareness and explicit laws of manifestation	63
17.2	A1: where perspectives are assigned	63
17.3	A2: phenomenal organization	65
17.4	Process provenance and temporal identity	66
17.5	Reports, phenomenal identification, and access	68
17.6	The constitutive choice	68
18	Completion, invariance, and empirical conservativity	69
18.1	The certified finite domain	69
18.2	Redescription and operational presentation	70
18.3	Independent composition	71
18.4	Physical conservativity	72
18.5	An aspect and an independently variable force	72
19	A fully worked finite vessel	73
19.1	Physical dynamics and realization	73
19.2	The exact transition matrix	74
19.3	Delayed observations and the native contrast geometry	74
19.4	External present-state reports have a different metric	74
19.5	Reports, symmetry, and gauge	75
20	Splitting, merging, and the stability of subject boundaries	76
20.1	Two cores and a controlled merger	76
20.2	Arbitrarily weak coupling gives a discontinuity	77
20.3	Copying, reset, and continuation	77
20.4	Finite partitions and universal unity	77
V	Interpretation, Comparison, and Empirical Responsibility	79
21	Biological and comparative applications	80
21.1	What an application must establish	80
21.2	Development and the absence of adult prerequisites	80
21.3	Sleep, anesthesia, and the lived scene	81
21.4	Candidate cortical correlates of scene-forming coupling	82
21.5	Disorders of consciousness and incomplete access	83
21.6	Blindsight and divided processing	84
21.7	Animals and nonhuman presentation	84
21.8	Plants, fungi, and minimal systems	85
22	Artificial systems, embodiment, and recursive cognition	86

22.1	The relevant implementation is active	86
22.2	Embodiment changes the available world	86
22.3	Memory transfer and the apparent continuity of a model	87
22.4	Implicit candidates and explicit certificates	87
22.5	Moral caution without premature attribution	88
23	Evidence, identification, and discriminating tests	89
23.1	One joint experiment, several kinds of claim	89
23.2	Descriptor adequacy before ontological inference	89
23.3	Finite witness calibration	90
23.4	What A2 would have to match	90
23.5	Conservativity and the limits of a report test	91
23.6	A concrete research sequence	91
24	The scope of the proposed resolution	92
24.1	A conditional law, not an unexplained transition	92
24.2	The philosophical price of structural completeness	92
24.3	What the contemplative interpretation contributes	93
24.4	Relation to Shadow Theory	93
24.5	Relation to existing consciousness programmes	94
24.6	Remaining realization obligations	94
25	Conclusion	95
	Technical Appendices	97
A	Contact, kinetic, and complete-path estimates	97
A.1	Finite contact gas and its Poisson comparison	97
A.2	Balanced packet removal	98
A.3	Signed queues and the exact inventory	98
A.4	A companion queue bounded away from empty origins	99
A.5	Removing the empty-origin cutoff	100
A.6	Existence and nonexplosion of the Bell comparison	100
A.7	From bulk flux to complete tagged paths	101
B	Predictive equivalence: algorithms and infinite-horizon limits	103
B.1	Finite linear representation of transcript probabilities	103
B.2	Why a finite hidden system can have infinitely many predictive states	104
B.3	Conditioning and continuation	104
B.4	A controlled measurable extension	105
C	Record witnesses and quantitative limitations	106
C.1	Finite jet determination without an order bound	106
C.2	Positive witnesses and physical implementation	106
C.3	Approximate capacity of a finite record	107

C.4	An exact probability countermodel and its repair condition	107
C.5	Physical archive corruption and wave error	107
D	Exact finite matrices and reproducibility	109
D.1	Boundary data for the five O1 profiles	109
D.2	Signed relabelling	110
D.3	The recurrent two-bit calculation	111
D.4	What the executable package checks	111
E	Reconstruction, symmetry, and the meaning of a source	113
E.1	A selected representative is not an inverse	113
E.2	Equivariant sections	113
E.3	Topological and physical qualifications	113
F	Notation, claim status, and construction dependencies	115
F.1	Principal notation	115
F.2	Claim-status register	116
F.3	Dependencies that cannot be omitted	117
F.4	Reproducibility and provenance	117

Part I

Awareness, Experience, and Perspective

Chapter 1

The source-to-experience problem

1.1 A question about presence, not merely performance

A person can distinguish a red object, act on that distinction, remember an encounter, and explain the optical and neural processes involved in seeing it. These are not identical achievements. Discrimination concerns a difference in response; explanation concerns a model of the process; memory concerns a relation to a previous event; experiential presence concerns there being something presented at all. A theory may ultimately identify some of these features, but the identification must be argued or declared. It cannot be obtained by moving without notice between the words “information,” “awareness,” and “experience.”

The Hard Problem is the explanatory question about experience that remains when the functional organization has been described. Its original modern formulation does not require a commitment to reductive materialism. Chalmers explicitly considered fundamental experience, structural coherence, organizational invariance, and a double-aspect approach [10]. Accordingly, taking awareness as primitive is a legitimate strategy within that discussion, not a discovery that invalidates the question by exposing a premise every earlier theory accepted.

The Shadow formulation changes the proposed explanatory direction. Instead of asking how a categorically nonexperiential entity produces an experiential one, it asks how one underlying reality admits restricted perspectives, what physical structures delimit them, and how their qualitative organization is fixed. This is a *source-to-manifestation* problem. It retains the hard explanatory burden in a different location. The assertion that awareness is fundamental does not tell us whether a particular process has one perspective, several, or none. It does not determine which differences within that process correspond to differences in experience. Those are the tasks assigned to the psychophysical constitution developed in this book.

An analogy with a physical primitive is useful but limited. A theory need not derive every primitive from something more elementary. Nevertheless, a primitive becomes explanatory through laws that constrain its instances. Declaring that an entity has charge is useful because a larger theory specifies its transformations and effects. Declaring source awareness without a realization law leaves the relevant assignments open. The present account therefore states its primitives and its assignment rules separately, and then asks exactly which consequences the resulting theory has.

1.2 The unsplit and its operational descriptions

U denotes reality prior to the operational source/readout distinction. “Prior” describes an explanatory order, not a time before the universe or a location beyond it. No numerical category such as one or many is applied at this level. This is a refusal of category, not a proof

that reality is a mystical numerical One. Nor is U a manifold, vacuum, empty space, field, operator, CSCF construction, soul, or knowing deity.

The working source S is a structure-side cut: a description sufficiently complete for the nominated readout problem. The readout T is its record-side cut, partial relative to that problem and capable of representing physical records, archives, and the vessels encountered in lived physics. Source and readout are complementary restricted descriptions under an operational cut. Their mathematical roles can be asymmetric without making them independent substances. The formal relation $p : S \rightarrow T$ concerns those descriptions; it is not a map from U into a world that stands outside reality.

A readout is a restricted description or manifestation of a nominated source domain. The readout belongs to the same physical account as the apparatus, its records, and its users. An experienced world may be model-mediated without being causally detached from its environment. Likewise, a representation may be incomplete without being false in every respect.

Two descriptions of one underlying event need not determine one another. A complete physical source state might support a coarse temperature readout and a separate chemical readout. The fact that both arise from the same source does not make temperature sufficient for chemistry. The same logical caution applies to physical and phenomenal descriptions. A common-source ontology makes their coexistence intelligible, but it does not establish the function connecting them. A precise fibre-sufficiency condition or a constitutive identification is still needed.

The hypothesis pursued here is that experiential presence is an intrinsic aspect of reality, while localized experiential organization depends on a realized vessel. A vessel is not simply an object with a name, such as “brain,” “computer,” or “organism.” It is a specified organization with components, interactions, accessible distinctions, time scales, and retained resources. The same macroscopic noun can conceal importantly different organizations. Conversely, different materials can realize the same nominated organization. Whether the differences that remain are phenomenally relevant is a question for the theory, not a fact fixed by terminology.

1.3 Awareness as a primitive and manifestation as organization

Awareness is the knowing aspect of U . This is an ontological proposal, not the attribution of thought, intention, personality, or an omniscient point of view to the universe. “Knowing” here names the proposed condition of manifestation; it does not describe a cosmic agent acquiring information. Within this interpretation awareness does not require a vessel in order to exist as capacity. A localized subject and a lived scene, by contrast, are assigned only through the certified physical organization specified by the constitution.

$$\boxed{\text{awareness} \neq \text{localized subject} \neq \text{lived scene.}}$$

A scene is what is presented; a subject is a localized perspective; awareness is not another item presented within that scene. None is a small observer hidden behind the eyes. The capacity language does not add a dormant spectator waiting for a brain, and it does not introduce a separately acting physical force.

The person is not the fundamental knower. Here that statement distinguishes the knowing aspect from the body image, name, memories, preferences and self-description through which a person ordinarily understands experience. It does not commit the theory to a permanent personal Self. Contemplative traditions differ on precisely that language, and some resist treating awareness as any kind of enduring self. The present proposal should be assessed on its stated distinctions rather than on an assumed agreement among traditions.

Physical development remains essential to manifestation. What develops is an organization with particular capacities, limits and history, not a quantity of primordial awareness manufactured from nothing. Growth, learning, injury and artificial embodiment can alter admission, content and continuity differently. A0 supplies this interpretation; A1 delimits perspectives and A2 specifies their phenomenal organization. The distinction motivates the laws without deriving them.

1.4 Nonduality and the reality of the readout

The intended nondual interpretation concerns ontological dependence. A vessel and its experienced world belong to reality; they are not detached copies existing outside it. Their distinctions can be real and consequential while remaining distinctions within one underlying reality. Calling a self-model constructed therefore does not make the person fictitious, and calling a readout incomplete does not deprive it of physical reality.

Source/readout non-equivalence concerns a different question: what can a representation determine? A representation may preserve a target exactly while identifying other physically distinct source conditions. The mathematical source domain used to demonstrate this is a nominated model, not an exhaustive description of U. The philosophical claim that no local perspective exhausts U accompanies the source-aspect account; its mathematical applications must still state which distinctions are lost and by which apertures.

On this interpretation, the cessation of one vessel's organization need not be pictured as a personal entity travelling back to source. A3 specifies when that localized episode ends. The persistence of underlying reality is a separate ontological claim, and does not establish continuation of the terminated subject. Likewise, a common source does not collapse distinct localized perspectives into one global observer.

The priority of awareness proposed here is ontological rather than a dated event before the formation of a brain. A0 does not posit a measurable episode of contentless experience occurring before every vessel. It specifies the knowing aspect of U; A1 still has to delimit the localized perspectives. The lens metaphor concerns this dependence of form on organization. It does not by itself establish a signal transmitted from elsewhere.

1.5 Four explanatory obligations

A source-aspect theory faces at least four obligations. *Admission* determines whether an experiential episode is assigned. *Individuation* determines which components belong to one perspective and which belong to different perspectives. *Temporal organization* determines whether an episode continues, terminates, divides, or is succeeded by another. *Qualitative organization* determines the contrasts, similarities, and transitions within its content. Evidence introduces a fifth obligation: the relation between these assignments and physically produced reports or other observations.

These obligations can be bundled into a single symbol, but doing so does not reduce their informational content. A formula such as

$$\Phi_V = \mathfrak{M}(R^*, \Sigma_V, x_V)$$

becomes a theory only when the dependence is specified. U is not a numerical argument of this map; A0 supplies an interpretation rather than an additional computable state variable. Here Σ_V denotes retained learning and representational scaffolding, while x_V denotes a realized state. Neither should be treated as an inexhaustible explanatory container. The domain, allowed transformations, and operational consequences of each variable must be given.

A finite construction makes the explanatory commitments inspectable: it requires a subject count, an assignment in every admitted case, and explicit handling of reports. Small counterexamples can expose freedoms that remain hidden in an informal description of a complex brain.

1.6 What is distinctive about the proposed account

The account combines target-relative source/readout sufficiency, physically specified records and interventions, and a complete predictive-profile assignment. Preparations, apparatus, receivers, and nulls remain part of the realization; a symbolic operator becomes a physical intervention only through an admitted coupling. The profile supplies relational phenomenal organization under a named constitutive law.

These commitments use established quotient, operator, and graph mathematics. The contribution is the constitution and its exact dependencies. The conservativity theorem distinguishes report predictions from an aspect interpretation that preserves their probabilities.

Chapter 2

Awareness, contents, and the personal model

2.1 A vocabulary that does not settle the question in advance

This book uses *primitive awareness* for the knowing aspect of U, and *lived scene* for the phenomenal presentation of a localized subject at a time. The word scene is deliberately broader than a visual field: feeling, sound, bodily presence or an unfamiliar nonhuman presentation need not resemble a picture. It uses *experience* for an episode of presentation, *content* for what is presented or discriminable within an episode, and *personal model* for the organization of body ownership, memories, practical commitments, and narrative self-description. *Access* is the causal availability of information to a nominated process. *Agency* is an operational capacity for model-based selection and action. *Metacognition* is the representation or assessment of cognitive processes.

These definitions are conventions for disciplined comparison, not a claim that all traditions use the same words. Nisargadatta's distinction between primordial awareness and content-relative consciousness differs from uses of "awareness" in parts of cognitive science. The text therefore compares roles rather than merely matching terms. The statement that two traditions both mention a witness is weak evidence of agreement unless their witnesses have the same specified properties.

The distinctions do not imply that awareness is a little observer inside the head. A picture inside a brain, inspected by another subject, would reproduce the question at a smaller scale. The source-aspect account instead associates presentation with a whole specified process. When a person says that they notice a thought, the thought and the noticing can be different aspects or components of that process. Their representational distinction does not establish that the noticing exists outside the physical system.

2.2 The red object and the interpretation of seeing

Consider a person who reports a red cup. A physical account can distinguish the environmental cause, the sensory transformations, a representation used to discriminate and act, a linguistic category, and a reflective explanation of perception. A second person might make the same color discrimination while lacking the reflective explanation. Their difference concerns conceptual and metacognitive organization. It does not, by itself, rank their experiential presence or their moral worth.

The phrase "the cup is constructed" is ambiguous. It may mean that the experienced cup

is a vessel-conditioned representation. It may also be taken to mean that the external cause was produced by the act of seeing. The former does not entail the latter. An internal model can be newly instantiated without creating its represented object. It can also misrepresent that object. A mistaken representation remains a real state of the representer even when its content is inaccurate.

For the formal account, the relevant source state includes not only the external stimulus but also the vessel’s preparation, retained history, and current condition. A stimulus-only map generally cannot settle a perception question whose answer depends on attention, learning, or context. Thus the appropriate abstract scheme is

$$s \mapsto x_V = p_V(s), \quad (x_V, \Sigma_V) \mapsto m_V,$$

where s is the complete nominated state and m_V is an internal model. A phenomenal assignment, if supplied, is another map or aspect relation. It is not inserted into this scheme merely by calling m_V a “rendering.”

This distinction also clarifies objectivity. Agreement across apertures can constrain a shared cause, but agreement is not a literal escape from every aperture. Two observations may share a calibration error or a prior interpretation. Objectivity is strengthened by interventions, independent channels, and invariance under relevant changes of procedure. It is not established solely by the number of people who repeat a description. Source/readout mathematics supplies an exact language for these sufficiency and comparison questions.

2.3 Memory, scaffolding, and continuity

A narrative person includes a relation to a past: particular memories, habitual expectations, learned categories, and a body-centered history. The account does not identify this entire structure with experiential presence. Nor does it imply that a narrative model is unreal because it is constructed. A constructed policy or autobiographical representation can have enduring causal and ethical importance.

The learning variable Σ_V may be distributed across physical structures rather than stored as a verbatim transcript. The distinction between storing a sequence and retaining a structure sufficient to reconstruct selected features is mathematically straightforward. If a code c retains only a function g of an original history, then every successful reconstruction factors through g . Features absent from g cannot be recovered uniquely without additional assumptions. A mnemonic scene, a summary, or a machine-generated state capsule can be analyzed in this way without assigning experience to the code.

Two notions of continuity must remain separate. *Informational continuity* means that relevant state or history relations are preserved. *Subject continuity* is a psychophysical identity relation over episodes. A copied memory may preserve informational continuity while belonging to another carrier. Conversely, a persisting carrier may lose some memories. The temporal law in Part IV will state what counts as continuation within SPC-2; no theorem about copying alone settles numerical subjective identity.

2.4 Conscious access, generative processing, and the lived scene

For a mature biological vessel, it is useful to distinguish foreground conscious access from the implicit organization that helps sustain it. In the first mode, material is attended, bound into an accessible presentation, and available for selected thought or action. It may be reportable

when the relevant channel works. This describes an important mode of consciousness; it does not define every possible experience by focal attention, verbal access, or reflective appraisal.

The generative mode includes implicit prediction, mnemonic reconstruction, procedural organization and unattended processing. These processes can belong to an awareness-bearing vessel without being separately presented in its scene. There is no second viewer and no hidden phenomenal scene assigned to “the unconscious.” A process can help organize how a room, a memory or a bodily feeling appears without its own operations becoming objects of experience. In this sense, restricted mode means part of the awareness-bearing organization, not currently appearing as a separately experienced content.

The lived scene is what people ordinarily mean by “what it is like right now.” In the proposed biological interpretation, lived consciousness depends on coupling between explicit/attended and unconscious/generative organization within a qualifying vessel. If that coupling is sufficiently disrupted, no scene is produced; when it returns, lived manifestation resumes. This is the Shadow interpretive model for mature biological organization, not an established universal neural mechanism or an additional condition in A1. “Explicit/attended” here describes scene-forming organization, not a requirement for focal attention to every content, reflective self-recognition, or human-style access in every conscious animal.

A2 is the candidate structural law for the lived scene; it does not generate or prove awareness. The distinction between a contributing process and a presented content must therefore be used carefully. Under A2 the complete endogenous predictive structure determines phenomenal organization, including relations that need not be verbally accessible or explicitly represented as objects. The present gloss does not discard selected predictive distinctions, shrink A2’s object, or introduce an unformalized switch that withholds experience from an otherwise qualifying core. Whether the strong structural law adequately handles unconscious processing remains a question for independent phenomenal comparison.

2.5 Vessel, scene, and memory

The distinctions needed here are:

awareness $\neq$ localized experiential subject
 $\neq$ biological/person continuity
 $\neq$ current scene $\neq$ memory of a scene.

In particular, continuity of a physical vessel, occurrence of a current lived scene, and memory of a scene answer different questions. A continuing organism can retain biological and personal continuity despite a gap in experience or recollection. A scene can occur without later successful recall. Conversely, an interval remembered as blank does not establish an unremembered experience. Forgetting a dream does not prove there was no scene; remembering nothing does not prove there was one.

no scene $\neq$ no vessel continuity.

Here a genuinely scene-less interval means no current lived experience and no currently instantiated localized experiential subject. A felt blankness, silence or darkness would itself be a scene. Awareness as the knowing aspect of U is not said to disappear; physical regulation may continue without a subject secretly observing it.

A localized experiential subject exists only while the relevant qualifying organization supports a lived scene. Physical process provenance can cross an experience-free interval, but A3 relates qualifying experiential occurrences: a genuine qualification gap ends the episode, and renewed qualification begins a new episode. Biological and person-level continuity

can therefore persist without numerical continuity of the experiential subject across the gap. Neither loss of responsiveness nor failed recall establishes such a gap; an altered or disconnected scene remains experience.

2.6 Minimal presence and higher-order reflection

The motivating philosophical position allows experience without a sophisticated self-description. This is why language, explicit philosophical insight, and adult autobiographical competence are not included as necessary clauses of the admission law. It is also why recursive self-description is not treated as the source of awareness. A system's capacity to describe its own description may enrich the organization of its perspective, but the inference from an extra representational level to experiential presence requires its own premise.

The distinction is particularly important for comparative questions. A nonhuman animal need not organize its world through human color words, transport practices, or personal narratives. A hypothetical artificial system could have a different input and action space again. These differences make a human-centered phenomenological inventory incomplete as a universal definition. They do not provide a shortcut to a positive attribution in every case. A comparative theory needs an independently specified realization and a justified assignment rule.

Under SPC-2, an episode is assigned by a constitutive condition on recurrent causal organization and available predictive distinctions. This condition is neither asserted to follow from the preceding philosophical discussion nor advertised as an established empirical criterion. The philosophy motivates a non-anthropocentric question. The law answers that question inside a model, at the cost of a substantial additional commitment.

Chapter 3

Contemplative inquiry and its evidential role

3.1 Reading *I Am That* as a philosophical source

I Am That is a translated and edited collection of dialogues. The edition consulted contains 101 talks, Maurice Frydman’s translator’s note dated October 1973, and Sudhakar S. Dikshit’s editorial note dated July 1981. Dialogue numbers provide stable locators across editions and are used for the discussions below [51].

The dialogues offer a disciplined conceptual challenge to identification with body, thought, and personal history. They are not a controlled study of many independently sampled contemplatives. The questioners’ objections are part of the source. They ask about memory, the reality of the world, suggestion, continuity, and the coherence of the teacher’s terminology. A scientific engagement should preserve these objections rather than extract only passages that resemble a preferred theory.

Dialogue 11 distinguishes awareness from consciousness directed toward contents. Dialogue 5 associates manifested consciousness with a vehicle and instrument. Dialogue 1 separates the changing experiencer–experience relation from a proposed common ground, while Dialogue 10 observes that witnessing still presupposes something witnessed. These passages motivate the separation of presence, content, personal identity, and reflexive witnessing. They do not specify a neural mechanism, a Hilbert-space operator, or a graph-theoretic threshold.

The useful correspondence is therefore architectural. A vessel can condition a mode of presentation without being identified with the proposed source of presence. The person can remain a practical organization without exhausting what is meant by awareness. The correspondence should not be strengthened into the claim that the book independently predicted the exact constitution of this monograph.

3.2 Identification is not the deletion of cognition

A particularly important passage occurs in Dialogue 6. The personal organization continues, but its identification with the witness changes. Dialogue 12 similarly describes practical activity, speech, and habituated response continuing without the same sense of personal ownership. These are not descriptions of a brain with every learned representation removed. They motivate a distinction between a self-model’s functional availability and identification with that model.

The distinction suggests comparing what changes in identification with what persists in

the vessel. A narrative self-model may remain available for practical action while its role as an exhaustive self-description changes. When that change alters the core's endogenous distinctions or their continuation laws, the proposed phenomenal structure changes. When only an external report or interpretation channel changes, a different statement about the episode need not imply a different episode. Admission, content, report, and continuity therefore remain separate questions.

In the philosophical interpretation adopted here, awakening is not the creation of a new entity and not the person becoming U. It is a thinning of identification between awareness and the self-model. The lived scene, thought, memory and bodily activity may continue. What changes is the conviction that this body and autobiographical story exhaust the fundamental knower. This does not posit a permanent personal Self or establish a common doctrine across contemplative traditions. *I Am That* is an important philosophical and phenomenological interlocutor, not scientific confirmation of that interpretation.

Meditation is thus not modeled as automatic deletion of Σ_V . The more precise hypothesis is that practice changes selected patterns of attention, interpretation, and identification. The mathematical theory can ask which target functions vary and which remain stable under such interventions. It cannot infer from the word “quiet” that dissipation is zero, that a physical spectrum has a particular gap, or that an optimal fixed point has been reached.

There is no contradiction in a person using a sophisticated conceptual model to investigate a mode of experience reported as less conceptually elaborated. The investigation and the reported episode are different stages of the process. A later verbal account is already a transformation involving memory and language. The model must include that transformation rather than treat the final sentence as an unmediated sample from a metaphysical source.

3.3 Witnessing, nonduality, and the limits of the translation

A simple ladder from source to awareness, awareness to subject, and subject to object can be pedagogically helpful, but it is not an unambiguous doctrine of the book. In Dialogue 10, witnessing itself remains dual because it has an object. Other dialogues point beyond even the sense of being a witness. In Dialogue 13, metaphors of an opening and of illumination are qualified: an opening is not the light, and illumination is not straightforwardly creation of the things illuminated.

The present theory uses these distinctions as constraints on careless explanation. An aperture is not an additional substance through which a beam of mental energy has been measured. A subject is not an immutable inner agent required to inspect every representation. An observer–observed relation can be an organization within one source model. These interpretations are compatible with a nondual philosophical orientation, but they do not mathematically establish universal numerical identity of all experiencers.

Several stronger claims in the dialogues lie outside the demonstrated content of this book: personal survival or direct knowledge beyond bodily death, timeless being, and assertions about the world's dependence on consciousness. A0 proposes awareness as an aspect of U; it does not establish any person's survival or access to a state beyond a vessel. Interpreting “world” as “experienced world” sometimes makes a useful reconciliation with observer-independent physics. It is nevertheless an interpretive choice, not literal agreement with every passage. The physical measurement models below neither need nor prove survival, rebirth, or the cessation of an external source when it is not perceived.

The mathematical statement that a noninjective readout loses some information is also weaker than the metaphysical claim that ultimate reality cannot be objectified at all. A particular source property may descend perfectly through a readout. The limitation must

be stated relative to the complete family of admitted representations. This preserves the philosophical intuition of restricted perspective without making an invalid universal inference from one many-to-one map.

3.4 First-person investigation and public evidence

Varela's neurophenomenology proposed reciprocal constraints between disciplined accounts of experience and cognitive science [53]. First-person descriptions are relevant because the target concerns experience. Their scientific use requires a procedure that separates reported content, interpretation, and physical conditions. Lutz and colleagues provide an example of combining trained first-person classifications with neural measurements in a visual task [27]. This supports methodological integration, not the inference that any specific contemplative metaphysics has been experimentally confirmed.

Gamma and Metzinger's questionnaire work investigates reports described as pure awareness using multiple phenomenological dimensions [18]. Such a study makes it possible to compare patterns of reported selfhood, content, spatiality, and cognition. It does not isolate an unconditioned ontological substance. The possibility of mixed or intermediate profiles is methodologically important: a theory should not force every report into a binary choice between ordinary ego and completely contentless awareness.

The present programme therefore uses a joint evidence model rather than an orthogonal sum of supposedly independent evidence types. Let n be a physical measurement, r a report, and i a protocol. The object of comparison is $P(n, r \mid i)$ together with the assumptions connecting reports to the claimed experience. Shared instructions, expectations, and memory constraints can correlate the evidence. A model must not count the same causal information twice merely because one variable is called first-person and another third-person.

This approach places philosophy and science in contact without confusing their standards. Philosophy clarifies what an explanation would have to explain and which inferences are conceptually licensed. Mathematics clarifies consistency, sufficiency, and invariance. Physical and phenomenological observations constrain whether a specific realization describes the intended case. None of these roles can be removed simply by declaring the others complete.

Chapter 4

Altered identity and unfamiliar experience

4.1 An autobiographical example

An autobiographical report motivating this account concerns a remote salvia-associated episode. Ordinary bodily location, personal identity, and access to the familiar life narrative were reported as absent. The remembered presentation was described as kaleidoscopic and difficult to characterize, without the ordinary reflective judgment that an altered state was occurring. Only after return to the usual mode of cognition did comparison with the prior life become possible.

The example is included as one retrospective report, not a controlled experiment and not a recommendation for inducing altered states. Its philosophical importance lies in the reported dissociation between experiential presentation and ordinary personal organization. Its epistemic limits are equally important. Kaleidoscopic presentation is still content. Later recall cannot establish that all encoding mechanisms had ceased. The report does not prove that the subject became an external object, accessed a universal field, or occupied a state numerically identical with every other instance of awareness.

Controlled salvinorin-A work has examined dissociation, unusual experiences, and memory effects [28]. The relevance here is the need to distinguish occurrence, contemporaneous interpretation, and later recollection. Neither dismissing the report as “only a hallucination” nor treating its felt certainty as metaphysical verification adequately captures that distinction.

4.2 Reality, veridicality, and experiential occurrence

An experience can occur while misrepresenting its environment. This does not make the occurrence fictitious. Conversely, a vivid report that something occurred does not settle the physical interpretation of its contents. A representation of a room and a representation of an unfamiliar landscape can both be states of a physical vessel; they need not be equally accurate maps of its current surroundings.

The source-aspect account separates three predicates. A process may be *physically realized*; it may be assigned an *experiential perspective* under a psychophysical law; and its content may be *veridical* relative to a nominated environmental target. These predicates need not have the same extension. In particular, a causal role for an internal model does not establish that its represented object exists externally. Likewise, phenomenal presence cannot be inferred solely from the fact that the model influences behavior.

This separation is essential to the intended philosophical synthesis. The claim that consciousness participates in reality does not require it to create every external event. An internal map is itself a new physical or informational object within a source model. Its construction can affect subsequent action. Whether that model-building process is experienced, and what its presentation is like, are additional assignments. The distinction avoids both a crude observer-created universe and the equally crude inference that a constructed percept is therefore causally irrelevant.

4.3 Altered weighting of self and world

Sensory isolation and flotation suggest a simple interpretive contrast: when external constraint is reduced, continuing generative activity can carry more of the scene's content. This is a possible description of such experience, not a claim that reduced input inevitably causes hallucination or that flotation reveals an otherwise inaccessible reality. A scene remains conditioned by the vessel even when familiar environmental anchors are weak.

Classic psychedelic and dissociative states should not be treated as a single mechanism. One influential predictive-processing account of classic serotonergic psychedelics proposes altered constraint by high-level expectations [9]; that proposal is a model, not a demonstrated explanation of every altered state. The salvinorin-A findings cited above concern a different pharmacological case. Across these examples, the interpretive question is how ordinary self/world binding, attention and mnemonic scaffolding can be reorganized while some experience remains.

A weakened identification with the self-model need not remove content or reflective capacity in the same way in every episode. Dissociation can involve changes in ownership, bodily location, agency, familiarity or autobiographical access; these should be described separately rather than collapsed into "ego loss." Neither unusual vividness nor a feeling of revelation establishes that a scene gives direct access to U. An internally generated presentation is real as an occurrence without being externally veridical.

4.4 What remains constant?

The motivating intuition is that experiential presence can remain while contents, identity, body model, and interpretation change dramatically. A formal theory must distinguish two readings. The first is a report-level regularity: subjects describe experience in conditions with reduced or transformed self-related content. The second is a metaphysical identity claim: exactly the same underlying awareness is numerically present throughout all such conditions or all subjects. The first can motivate controlled investigation. It does not prove the second.

Within SPC-2, A0 expresses the source-aspect commitment, while A1 and the temporal rule specify localized episodes. A report of unchanged presence across altered content is compatible with the law when a qualifying core continues. It is not evidence that every source detail remained unchanged. A0 does not promise that every physically continuous organism has one uninterrupted episode through every sleep, anesthetic, or dissociative condition. That would require a concrete realization analysis and evidence, not an inference from a metaphor of a constant light.

The resulting discipline preserves the human insight without making the monograph depend on an inaccessible certainty. A person may reasonably regard a contemplative or altered-state episode as existentially significant. A public theory must still expose the inferential steps between that significance and its universal claims. The distinction is not hostile to first-person inquiry; it is what makes different kinds of inquiry capable of correcting one another.

Part II

Source, Physical Events, and Material Records

Chapter 5

Source/readout foundations

5.1 Nominated sources and restricted descriptions

A mathematical source domain is a specified class of states or histories relevant to a question. It is the working structure-side cut, complete relative to that problem, not $\mathbf{U}$. The schematic record-side codomain T of the philosophical discussion is denoted X in the following theorems. These results concern operational descriptions and make no set-theoretic claim about the unsplit prior. Let S be such a domain after declared descriptive redundancies have been removed. A readout is a surjection $p : S \rightarrow X$. A target $a : S \rightarrow A$ is a property or relation that the investigation asks the readout to retain. This formulation precedes a probability measure, a Hamiltonian, or a phenomenal assignment [41, 37].

Theorem 5.1 (Descent and coarsest target completion). *A map $\bar{a} : X \rightarrow A$ satisfying $a = \bar{a} \circ p$ exists if and only if a is constant on every fibre of p . It is unique when it exists. For a family of targets (a_α) , let*

$$e(s) = (p(s), (a_\alpha(s))_\alpha), \quad E = e(S).$$

If a surjection $f : S \rightarrow F$ retains p and every a_α , then there is a unique surjection $h : F \rightarrow E$ with $e = h \circ f$.

Proof. A factorized target takes the same value at any two states with the same p value. Conversely, for $x \in X$ choose s with $p(s) = x$ and define $\bar{a}(x) = a(s)$. Fibre constancy makes the choice immaterial, and surjectivity gives uniqueness. For the family, write $p = u \circ f$ and $a_\alpha = v_\alpha \circ f$. Then $h(y) = (u(y), (v_\alpha(y))_\alpha)$ has image in E because f is onto. It is surjective because e is, and unique because every point of F is attained by f . $\square$

The theorem orders information. It does not select the dynamics acting on that information. For example, with $S = \{0, 1\}^2$ and $p(s_1, s_2) = s_1$, retaining the second coordinate completes the representation to S . There are still many deterministic maps and stochastic kernels on S . A rule selecting one of them requires new data. This distinction will recur in both quantum event selection and phenomenal assignment.

A perceived object is not literally a set-theoretic fibre. The quotient expresses an identification performed by a specified representation: source conditions within one fibre are indistinguishable to that representation. A noisy observation requires a further distinction. Equal observed samples need not mean equal source laws. For a channel $Q(\cdot | s)$, operational equivalence is equality of the relevant output distributions, not accidental equality of one draw.

5.2 Predictive closure

Proposition 5.2 (Deterministic readout dynamics). *For a source evolution $F_t : S \rightarrow S$, an autonomous reduced evolution $U_t : X \rightarrow X$ with $pF_t = U_t p$ exists exactly when pF_t is constant on p -fibres. If this holds for all times and F is a semigroup, the reduced maps form a semigroup.*

Proof. Apply Theorem 5.1 to $a = pF_t$. For composition, $U_{t+s}p = pF_{t+s} = pF_t F_s = U_t U_s p$; surjectivity of p proves the result. $\square$

For stochastic systems the analogous condition concerns whole transition laws. Suppose S is finite and $T_u(s, s')$ is a controlled transition kernel. The p -projection has a common transition law for every complete initial distribution precisely when, for all x, x', u ,

$$\sum_{s': p(s')=x'} T_u(s, s') \text{ is independent of } s \in p^{-1}(x). \quad (5.1)$$

Necessity follows by starting at two point masses in one fibre. Sufficiency follows by conditioning on the current fibre and summing. The same reasoning works for finite adaptive protocols when the controls depend only on the retained reduced history.

This is a condition on a supplied kernel, not a way of choosing that kernel. If it fails, the readout is insufficient for autonomous prediction and must be enlarged, or a non-Markov history description must be used. The failure need not indicate a new force: it can reflect omitted ordinary physical variables. Conversely, successful reduced prediction does not prove that the representation exhausts the source.

5.3 Multiple apertures and non-reconstructibility

A noninjective p cannot recover every source property. It can recover exactly the properties constant on its fibres. Two noninjective readouts may jointly be injective: the two coordinate projections of $\{0, 1\}^2$ recover the entire pair. Therefore the philosophical doctrine that no internal observer exhausts U is stronger than the noninjectivity of one particular channel. It concerns the admitted collection of representations and the status of the source domain itself.

Target-relative sufficiency permits exact knowledge of some source distinctions through a restricted representation. Others require further structure. The consciousness programme asks the corresponding question: whether a vessel descriptor retains the distinctions needed for an independently nominated phenomenological or report contrast.

5.4 The Everything Equation and constitutive selection

The broader Shadow programme expresses lawhood by a fixed-point relation, conventionally written $L = \Omega \Delta \partial[L]$. Here the operator version is

$$L = \mathcal{O} \circ \Delta \circ \partial(L). \quad (5.2)$$

A contraction satisfying the fixed-point hypotheses selects a unique point, but the relation does not select its own domain or maps. The operator $\mathcal{O}$ belongs to the supplied mathematical construction. It is not U , and the fixed-point equation does not turn the ontological prior into an operator.

Accordingly, SPC-2 is a compatible constitutive extension, not a deduction from Equation (5.2). Its subject-admission and qualitative laws are additional commitments.

Chapter 6

Quantum currents and observer-independent events

The physical question in this chapter is what specifies an event before anyone notices it. Source/readout mathematics first identifies information sufficient to track weights and currents. The two following chapters then supply distinct constitutions for actual histories and material records. Their role is to establish physically specified processes to which an experiential interpretation can subsequently be applied.

6.1 The coherent source

Let $\mathcal{H}$ be finite dimensional, $H(t) = H(t)^\dagger$, and (P_n) a finite orthogonal resolution of the identity. The sectors can include apparatus and an inaccessible reference. For

$$i\hbar\dot{\Psi} = H\Psi, \quad \Psi_n = P_n\Psi,$$

define $w_n = \|\Psi_n\|^2$ and current into n from m by

$$J_{nm} = \frac{2}{\hbar} \operatorname{Im} \langle \Psi_n, H\Psi_m \rangle. \quad (6.1)$$

Self-adjointness and direct differentiation give

$$J_{nm} = -J_{mn}, \quad \dot{w}_n = \sum_{m \neq n} J_{nm}, \quad |J_{nm}| \leq \frac{2\|H\|}{\hbar} \sqrt{w_n w_m}. \quad (6.2)$$

The diagonal term makes no contribution because its expectation is real. The last estimate follows from Cauchy–Schwarz. It controls possible behavior near empty sectors without making a choice of actual events.

An actual jump history requires a law on paths, not only the continuity equation. If N_{nm} counts transitions $m \rightarrow n$, its predictable intensity λ_{nm} is defined relative to a specified filtration by requiring

$$N_{nm}(t) - \int_0^t \mathbf{1}_{\{X_{s-}=m\}} \lambda_{nm}(s) \, ds$$

to be a local martingale. A signed current density and a random atomic counting measure are different objects. Their identification cannot be made path by path by changing notation.

Bell's minimal rates are

$$\lambda_{nm}^B(t) = \frac{[J_{nm}(t)]_+}{w_m(t)} \quad \text{when } w_m(t) > 0. \quad (6.3)$$

The construction is established in the Bell-process literature [3, 15]. Its selection from a physical constitution is a further problem. For example, directed expected traffic of the form

$$F_{nm} = [J_{nm}]_+ + s_{nm}, \quad s_{nm} = s_{mn} \geq 0,$$

has the same net current for every symmetric surplus s . Equality of populations therefore does not identify the event traffic, waiting-time structure, or full path law.

6.2 Predictive current coordinates

A source/readout completion can determine what is needed to evaluate a proposed event law without selecting it. Let $(H_a)_a$ be a finite family of admitted control Hamiltonians, and define Hermitian current observables

$$C_{nm}^a = \frac{P_n H_a P_m - P_m H_a P_n}{i\hbar}, \quad \mathcal{L}_a^* O = \frac{i}{\hbar} [H_a, O].$$

Let W be the smallest real subspace of Hermitian operators containing all P_n, C_{nm}^a and invariant under every $\mathcal{L}_a^*$. It can be computed by repeatedly adjoining images of a basis. Its dimension is at most $(\dim \mathcal{H})^2$.

Theorem 6.1 (Predictive-current representation). *For a basis $O_1, \dots, O_d$ of W , let $\xi(\rho) = (\text{tr } \rho O_j)_j$. Two density operators have equal ξ if and only if all future sector weights and currents agree under every finite sequence of admitted constant-control pulses. The image of ξ is the coarsest representation sufficient for this target, and it evolves autonomously during each pulse.*

Proof. Invariance of W implies invariance under $e^{t\mathcal{L}_a^*}$ and finite products of these exponentials. Equal coordinates therefore give all nominated future expectations. Conversely, equality for all nonnegative pulse durations gives equality of all one-sided mixed derivatives at zero. The resulting ordered words in the generators applied to P_n, C_{nm}^a span W , so the coordinates agree. Coarsest sufficiency follows from Theorem 5.1. In a basis write $\mathcal{L}_a^* O_i = \sum_j A_{ij}^a O_j$; then $\dot{\xi} = A^a \xi$. $\square$

The theorem is about informational sufficiency. It says nothing about whether ξ is experienced. It also does not determine which event generator, preparation measure, or physically exposed actuator should accompany the coherent source. Those additional distinctions are retained in the two measurement constitutions below.

6.3 The independence of physical events from conscious observation

The mature measurement programme contains two different constitutive completions [37, 35, 34]. One uses actual continuous configurations guided by the complete wave and an initial equilibrium law. The other uses a deterministic pilot medium with an initial spatial ensemble and a controlled complete-path limit to Bell dynamics. Neither includes conscious awareness as the trigger of an event.

This is a conditional physical result. It establishes that the specified models define events and records without an awareness variable. It does not prove that these models are the unique laws of nature, nor that consciousness has no relation to the processes they describe. The relevance is conceptual and technical: a psychophysical aspect can be added to the realized history without being invoked to complete an otherwise undefined measurement.

6.4 Complete experiments and comparison spaces

A complete experiment includes the source, unknown input, inaccessible reference, apparatus, fuel, working display, archives, reset receivers, null and failure branches, control programme, and all future return interactions. An approximate operation leaves its actual output in the model. It does not replace a daughter state by an ideal one merely because a label has been declared.

Lemma 6.2 (Composition and conditioning). *If successive complete Markov kernels differ uniformly by at most ϵ_j in total variation, their complete retained-history laws differ by at most $\sum_j \epsilon_j$. The same telescoping conclusion holds for quantum instruments in half-diamond distance with the required reference retained. If $\|\sigma - \tau\|_1 \leq \delta$, $p = \text{tr } \sigma > 0$, and $q = \text{tr } \tau > 0$, then*

$$D(\sigma/p, \tau/q) \leq \min\{1, \delta/p\}.$$

If $d_{\text{TV}}(P, Q) \leq \epsilon$ and both conditional laws on E exist, with $P(E) = p > 0$, then

$$d_{\text{TV}}(P(\cdot | E), Q(\cdot | E)) \leq \min\{1, 2\epsilon/p\}.$$

Proof. Replace kernels one at a time. The common remaining suffix contracts the chosen distance, and the retained prefix is part of the input on which the uniform bound holds. For normalization, $\|\sigma/p - \tau/q\|_1 \leq \delta/p + |p - q|/p \leq 2\delta/p$. The same triangle argument applied to restrictions of probability measures proves the conditional estimate. $\square$

The positivity of conditioning probabilities is not optional. A formal cap by one cannot define conditioning on a null event. Likewise, closeness of final output probabilities does not by itself imply closeness of the actual historical record. These distinctions are crucial when physical reports are later used as evidence concerning experience.

Chapter 7

A massive configuration and record constitution

7.1 Guidance, equilibrium, and actual histories

Fix a smooth finite programme with wave function

$$\Psi(q, t) \in L^2(\mathbb{R}^n; \mathcal{H}_I \otimes \mathcal{H}_R)$$

and Hamiltonian

$$H = - \sum_j \frac{\hbar^2}{2m_j} \partial_{q_j}^2 + V(q, t),$$

where V is Hermitian and the chosen domain ensures the required wave regularity and an almost-everywhere global guidance flow over the promised horizon. Internal keys and the reference are components of the wave, not extra discrete actual occupancies. Define

$$\rho = \Psi^\dagger \Psi, \quad j_j = \frac{\hbar}{m_j} \text{Im}(\Psi^\dagger \partial_{q_j} \Psi), \quad \dot{Q}_j = j_j(Q, t) / \rho(Q, t).$$

The complete initial configuration has density $\rho(\cdot, 0)$. These are constitutive premises. They belong to the configuration-guided tradition [7, 16]; they are not inferred from source incompleteness.

The Schrödinger equation gives $\partial_t \rho + \nabla \cdot j = 0$. Transport by the guidance velocity obeys the same continuity equation, hence preserves the initial equilibrium density under the stated flow hypotheses. This is equivariance. It supplies ordinary position distributions at all times while the actual history remains the guided path. The mathematical existence assumptions are kept explicit rather than being hidden behind the formal quotient j/ρ at nodes.

7.2 An exact moving Gaussian pointer

Let A be a two-valued internal control projector, and take an oscillator pointer of mass M and frequency ω . Write

$$\sigma^2 = \frac{\hbar}{2M\omega}, \quad b(t) = L[10s^3 - 15s^4 + 6s^5], \quad s = t/T,$$

for $0 \leq t \leq T$, and set $c = b + \ddot{b}/\omega^2$. The controlled pointer Hamiltonian is

$$H_Y = \frac{p_Y^2}{2M} + \frac{M\omega^2}{2}(Y - c(t)A)^2. \quad (7.1)$$

The endpoint conditions make the translated packet stationary at the start and end. Extended by constants, the displayed b is C^2 ; its associated control c is continuous and piecewise smooth. A monotone C^∞ profile flat at both endpoints can be substituted when the autonomous programme requires that regularity, with the same formulas for its corresponding b and c . In the $A = 0$ branch the ground-state packet remains centered at zero. In the $A = 1$ branch a Gaussian centered at $b(t)$ with the usual linear phase is an exact solution; substitution reduces the center equation to $\ddot{b} + \omega^2 b = \omega^2 c$.

Let g_σ be the centered normal density and let the orthogonal control-branch weights be p_0, p_1 , $p_0 + p_1 = 1$. The physical pointer density and current are

$$\rho_Y(y, t) = p_0 g_\sigma(y) + p_1 g_\sigma(y - b(t)), \quad j_Y(y, t) = p_1 \dot{b}(t) g_\sigma(y - b(t)). \quad (7.2)$$

Orthogonality of the internal control keys removes cross terms in these marginal expressions. Finite overlap of the spatial packets is not replaced by an exact disjoint-support assumption.

For this exact example, Y is the sole actual coordinate participating in the write. The control, source, and reference belong to the internal Hilbert space; any other actual coordinates factor into a common spectator wave. Thus the full guidance component for Y is j_Y/ρ_Y . In a general correlated many-coordinate wave, marginalizing a current need not yield the actual componentwise guidance velocity. The following path theorem uses the closed pointer realization just specified.

Theorem 7.1 (Quantile flow and first passage). *Let*

$$F_t(y) = p_0 \Phi(y/\sigma) + p_1 \Phi((y - b(t))/\sigma),$$

where Φ is the standard normal distribution function. If U is uniform on $(0, 1)$ under initial equilibrium, the pointer trajectory is

$$Y_t = F_t^{-1}(U).$$

For a threshold $h = L/2$, put $\delta = \bar{\Phi}(L/(2\sigma))$. Apart from the initial tail of probability δ , the threshold first-passage density is

$$f_\tau(t) = p_1 \dot{b}(t) g_\sigma(h - b(t)), \quad 0 < t < T.$$

The mass crossing after time zero is $p_1(1 - 2\delta)$, and the final below-threshold probability is $p_0(1 - \delta) + p_1\delta$.

Proof. The continuity equation gives $\partial_t F_t(y) = -j_Y(y, t)$. Along a guidance trajectory, $\frac{d}{dt} F_t(Y_t) = -j_Y + \rho_Y \dot{Y}_t = 0$. The initial $F_0(Y_0)$ is uniform, proving the quantile formula. Since $\dot{b} \geq 0$, the guidance velocity is nonnegative. Survival below h is therefore $F_t(h)$ after accounting for the initial upper tail. Differentiating it yields the density. The endpoint values are $F_0(h) = 1 - \delta$ and $F_T(h) = p_0(1 - \delta) + p_1\delta$, whose difference is the stated crossing mass. $\square$

The law describes physical time, including an initial-tail event and a retained null branch. It is not a rule that waits for a person to look at the pointer. Whether its eventual presentation is experienced is a separate psychophysical question.

7.3 Capture, pending states, and null instruments

A detector programme must retain unsuccessful and unfinished branches. Let the resource basis include r, p_a, c_a, l_a , denoting ready, pending, captured, and lost states for $a = 0, 1$. For

orthogonal source projectors P_a , the unitary preparation can produce

$$\begin{aligned} \Psi = & \cos \theta \psi |r\rangle \\ & + \sin \theta \sum_a P_a \psi \left[\cos \varphi |p_a\rangle + \sin \varphi \sqrt{\eta} |c_a\rangle + \sin \varphi \sqrt{1-\eta} |l_a\rangle \right]. \end{aligned} \quad (7.3)$$

The resource keys are orthogonal wave components; they are not additional actual occupancies in the configuration-guided constitution. Their squared norms give the following ideal status weights:

$$\cos^2 \theta, \quad \sin^2 \theta \cos^2 \varphi, \quad \eta \sin^2 \theta \sin^2 \varphi, \quad (1-\eta) \sin^2 \theta \sin^2 \varphi.$$

A spatial readout must still realize a discrimination of these keys. First define the ideal reference instrument. When “null” means every noncapture status, tracing the resource-key carrier after the corresponding ideal projection gives the unnormalized source instrument

$$\mathcal{N}(\rho) = \cos^2 \theta \rho + \sin^2 \theta [\cos^2 \varphi + (1-\eta) \sin^2 \varphi] \sum_a P_a \rho P_a. \quad (7.4)$$

The coherent ready component survives. Replacing every null daughter by a fully dephased or fully unchanged ideal state would be incorrect.

Write the source Kraus blocks as

$$\begin{aligned} V_r &= \cos \theta I, & V_{p_a} &= \sin \theta \cos \varphi P_a, \\ V_{c_a} &= \sin \theta \sin \varphi \sqrt{\eta} P_a, & V_{l_a} &= \sin \theta \sin \varphi \sqrt{1-\eta} P_a. \end{aligned}$$

They obey $\sum_\alpha V_\alpha^\dagger V_\alpha = I$. If an admitted controlled spatial writer retains these keys and presents outcome b with calibrated probability $M(b | \alpha)$ on key α , its reduced source instrument is

$$\mathcal{I}_b^{\text{phys}}(\rho) = \sum_\alpha M(b | \alpha) V_\alpha \rho V_\alpha^\dagger.$$

The ideal null map uses $M(N | \alpha) = 1$ on noncapture keys and zero otherwise. Finite-overlap pointers generally give a different, explicitly noisy matrix M . If every column differs from its ideal column by at most ϵ in total variation, the status-flag plus source/reference instrument differs by at most ϵ in half-diamond distance. Indeed, after adjoining a reference each conditional block is positive; summing its trace times the classical column discrepancy gives this bound uniformly over input states. The reference is retained in this comparison, but the resource-key carrier has been traced. If those keys or the pointer later return to interact, their full conditional state must instead be propagated; this reduced bound is not a complete-history estimate.

For the ideal reference instrument with $\theta = \pi/3$, $\varphi = \pi/4$, and $\eta = 2/3$, the status probabilities are $(1/4, 3/8, 1/4, 1/8)$. If

$$\rho = \begin{pmatrix} 2/3 & 1/3 \\ 1/3 & 1/3 \end{pmatrix},$$

then $\mathcal{N}(\rho) = \rho/4 + \text{diag}(\rho)/2$, $P(N) = 3/4$, and

$$P(N, X+) = 11/24, \quad P(X+ | N) = 11/18.$$

This exact finite calculation shows why a full physical continuation is more informative than a record label alone. The surviving coherent component changes the later noncommuting measurement.

7.4 Historical archives and reset receivers

Suppose an archive coordinate y has a wave of the form

$$\Psi(y, z) = \sum_k g_k(y) |k\rangle \Xi_k(z),$$

where the g_k are real held packets and the internal keys $|k\rangle$ are orthogonal. If subsequent evolution preserves each key and the held real packet, then $\text{Im}(\Psi^\dagger \partial_y \Psi) = 0$ pointwise. The guidance coordinate of the archive is fixed. This is a physical retention theorem, not merely membership in a modular centralizer.

A reversible reset moves the old working state into a receiver. For a working register W and a prepared blank R , a SWAP gives

$$\rho_{WE} \otimes |0\rangle\langle 0|_R \mapsto |0\rangle\langle 0|_W \otimes \rho_{RE},$$

with the old correlations retained. The active register is restored, but usable global capacity has not been created. A later interaction with R must use its actual correlated state. If reset is realized by a moving trap, the potential, transport, and receiver remain part of the programme [35].

Spatial feedback is equally concrete. A bounded function $g(y)$ can control an ordinary internal operator B through $g(Y)B$. Let G_{id} be the intended branch-constant control on the retained key space, commuting with B , and let $\Psi_{\text{id}}(s)$ be the ideal programme wave. Suppose

$$\sup_{0 \leq s \leq t} \|(g(Y) - G_{\text{id}})\Psi_{\text{id}}(s)\|^2 \leq \epsilon_g.$$

Duhamel's formula bounds the wave error by $t\|B\|\sqrt{\epsilon_g}/\hbar$. The mean-square mismatch includes both the smooth transition region and tails on the wrong plateau; transition-region probability alone is insufficient. A separate archive coordinate can remain protected while the working display participates in feedback. The separation must be designed; it does not follow from the word "record."

7.5 Autonomous realization and historical error

A driven finite programme can be approximated by a massive controller with coordinate x :

$$H_{\text{aut}} = \frac{P_x^2}{2M_c} + H_{\text{osc}} + H_{\text{const}} + \sum_j f_j(x) B_j.$$

A packet centered on $x_0 + vt$ samples the desired couplings. For an initial width s_c , its free width is

$$s_t = \sqrt{s_c^2 + \left(\frac{\hbar t}{2M_c s_c}\right)^2}.$$

Proposition 7.2 (A clock approximation on a declared regularity domain). *Let $\mathcal{X}_0 = \mathcal{H}_{\text{sys}}$ and let $\mathcal{X}_2$ be a specified system Sobolev or weighted graph domain whose norm controls the archive traces used below. Clock coordinates are treated in $L^2(\mathbb{R}_x)$, not differentiated in $\mathcal{X}_2$. Assume the autonomous propagator is uniformly bounded by C_r on $L^2(\mathbb{R}_x; \mathcal{X}_r)$ over the promised interval, and that the driven solution ψ_t satisfies*

$$\sup_{t \leq T} \sum_j \text{Lip}(f_j) \|B_j \psi_t\|_{\mathcal{X}_r} \leq M_r, \quad r = 0, 2.$$

For a freely moving Gaussian clock χ_t of mean $x_0 + vt$ and width s_t , with product initial state, the actual autonomous wave satisfies

$$\epsilon_r := \sup_{t \leq T} \|\Psi_t - \chi_t \otimes \psi_t\|_{L_x^2 \mathcal{X}_r} \leq \frac{C_r M_r}{\hbar} \int_0^T s_t dt \leq \frac{C_r M_r T}{\hbar} \left(s_c + \frac{\hbar T}{2M_c s_c} \right).$$

Proof. The product wave solves the equation with couplings $f_j(x_0 + vt)B_j$. Its residual in the autonomous equation is $\sum_j [f_j(x) - f_j(x_0 + vt)]\chi_t \otimes B_j\psi_t$. The Lipschitz bound and the Gaussian variance give residual norm at most $M_r s_t$. Duhamel's formula and the stated propagation bound give the first inequality; $s_t \leq s_c + \hbar t/(2M_c s_c)$ gives the second. $\square$

For bounded couplings the $r = 0$ propagation bound is supplied by unitarity. The $r = 2$ estimate requires the stated domain preservation and commutator bounds; unbounded oscillator forces require the corresponding weighted domain. Choosing $s_c \propto \sqrt{\hbar T/(2M_c)}$ gives the claimed fixed-programme limit when these constants are uniform. Differentiating the fast clock phase itself would not furnish that uniform archive estimate.

Why retain the derivative estimate? A small L^2 wave difference alone need not control a boundary-crossing current. For a held archive with nominal zero current, a trace estimate on the relevant surfaces bounds the probability of corruption over an interval I by a term of the form

$$\frac{\hbar}{m} C_{\text{tr}}^2 |I| \epsilon_2 (2B + \epsilon_2),$$

where B controls the required nominal Sobolev norm. The complete retained-output error and the historical-crossing error must both be included. This provides a physical reason for distinguishing an accurate endpoint distribution from a faithful past record.

7.6 What the constitution establishes

The massive model specifies the actual configuration, initial equilibrium, material writing, pending and null branches, copies, receivers, feedback, and a finite autonomous approximation. Its measurement chain is specified on the closed-pointer domain, with a controlled autonomous approximation under the regularity and uniformity hypotheses above. It does not derive guidance or equilibrium from the source/readout premise. It also does not establish a phenomenal predicate. Its role in this monograph is to show how an ordinary physical record process can be fully defined before the psychophysical interpretation is supplied.

An equivariant diffusion gives an instructive contrast. The stochastic law

$$dQ_t = (j/\rho + D \nabla \log \rho) dt + \sqrt{2D} dW_t$$

has the same density continuity equation for constant $D > 0$, because its additional drift and diffusion cancel in the Fokker–Planck equation. It can nevertheless have different paths and archive behavior. Equal one-time densities do not identify the realized microscopic process. The same caution will apply to comparing outward reports of different candidate vessels.

Chapter 8

A deterministic pilot medium and a complete path limit

8.1 The declared microscopic inventory

The pilot completion uses a different ontology [34, 37]. Its complete finite ordinary configuration graph includes source, apparatus, fuel, loss products, displays, archives, receivers, clock, and any actual basis coordinate of a reference. One predesignated carrier is the actual ordinary configuration. Other carrier positions and packet banks are pilot degrees of freedom, not independent copies of the unknown quantum input.

Four constitutive commitments are essential. The canonical field has a numerical Hermitian energy with primitive bond couplings. A declared exporter and complete binary reaction catalogue govern the pilot medium. Additional direct ordinary forces depending on readable pilot coordinates are absent from that catalogue. Finally, deterministic gas flight and contact are supplied with an initially independent spatial ensemble and a carrier preparation law. These are physical assumptions; in particular, the access restriction is not a rule allowing a coupling and then refusing to examine its consequences.

The microscopic theory is hybrid: a canonical coherent field, continuous free flight, and explicit contact and export rules. Finite contacts have reversible logical lifts with retained blank receivers. The autonomous ordinary circuit is a finite Hermitian Hamiltonian. A smooth realization of an isolated contact module does not establish that the entire hybrid source follows from one universal smooth Hamiltonian. That distinction is retained throughout.

8.2 Bond ownership and conservative export

For a bond $e = (r, q)$, introduce a connection χ_e and conjugate coordinate Π_e . A canonical action is

$$S = \int \left[\frac{i\hbar}{2} (\Psi^\dagger \dot{\Psi} - \dot{\Psi}^\dagger \Psi) + \hbar \sum_e \Pi_e \dot{\chi}_e - h(\Psi, \chi, t) \right] dt, \quad (8.1)$$
$$h = \sum_r \Psi_r^* H_{rr} \Psi_r + \sum_{e=(r,q)} (e^{i\chi_e} \Psi_q^* H_{qr} \Psi_r + \text{c.c.}).$$

Since h does not depend on Π , an initial $\chi = 0$ remains zero. The equations give Schrödinger evolution and

$$\dot{\Pi}_e = -\hbar^{-1} \partial_{\chi_e} h = J_e.$$

Within real quadratic energies additive over primitive vertices and binary bonds, endpoint phase covariance forces the displayed $e^{i\chi}$ dependence once the ordinary H_{qr} is fixed. A gauge-invariant cycle term could contribute an extra divergence-free torque without changing that ordinary matrix at $\chi = 0$. Primitive bond additivity therefore matters.

Let $u_e = N(\Pi_e - \Pi_e(0)) - k_e$, with initially $u_e = k_e = 0$. At a first hit $u_e = \pm 1$, export one signed packet, advance k_e by that sign, and reset the residue to zero. Then

$$\sup_{t \leq T} \left| \frac{k_e(t)}{N} - \int_0^t J_e(s) ds \right| \leq \frac{1}{N}, \quad \#\text{exports}_e \leq N \int_0^T |J_e| dt. \quad (8.2)$$

The first inequality is the residue bound. Each completed excursion costs at least $1/N$ of total action variation, proving the second. A finite preallocated packet bank therefore suffices on a promised finite horizon.

Assign charge e_r to a carrier at r and $b_e = e_q - e_r$ to a positive packet. The binary charge-conserving service $e_i + b_e = e_j$ forces $i = r, j = q$. Negative packets reverse that service. Opposite packets may recombine to neutral products. With carrier counts n , net packet count Z , and coherent weights w , the inventory

$$\mathcal{C} = n + BZ + Bu - Nw$$

is conserved. The continuous terms cancel by $\dot{u} = NJ$ and $\dot{w} = BJ$; each discrete reaction preserves the same expression. This is an exact inventory law for the specified hybrid dynamics.

8.3 Spatial preparation produces timing

Allocate a candidate channel to every possible packet-carrier pair and to every pair of packet slots eligible for recombination. Their fixed frequencies depend on declared geometry and response coefficients, not on the instantaneous Bell quotient. Let the total candidate frequency be R_N . Prepare M_N particles with independent uniform longitudinal positions in $[-L_N, 0]$, common speed v , and independent transverse marks, where $L_N = M_N v / R_N > vT$.

Each particle crosses the contact plane once at the deterministic time fixed by its initial position. The probability that it arrives by T is $p = R_N T / M_N$. Conditional on arrival, the time is uniform on $[0, T]$ and its channel mark has the assigned distribution. The finite marked history is thus compared with a marked Poisson history by a binomial-Poisson coupling. A safe bound is

$$d_{TV}(\text{Law } \mathcal{G}_N, \text{Law } \mathcal{G}_P) \leq 2(R_N T)^2 / M_N. \quad (8.3)$$

Using a common causal reaction map cannot increase this distance. Complete receiver labels must be transported by a common decoration kernel; particle identities are not silently replaced by independent labels.

This is a derivation of timing from a declared spatial ensemble, not a preparation of desired future Bell event times. The ensemble assumption remains essential. Correlated incoming positions can generate another contact history. Once the complete initial microstate is fixed, the finite microdynamics is deterministic.

Fast but finite opposite-packet recombination suppresses balanced service. The comparison proof couples the two-species device to a signed-queue process and bounds the service that occurs before opposite packets annihilate. By choosing the recombination scale and gas size after fixing the graph, packet budgets, and horizon, both comparison errors tend to zero. The physical bank remains finite at each parameter value; virtual banks used to define a common Poisson comparison are mathematical comparison devices, not uncounted physical resources.

8.4 From signed queues to the Bell path

Write $x_r = n_r/N$, $m_e = \mu_N Z_e/N$, and, for $e = (r, q)$,

$$\Phi_e^+ = \kappa_e x_r [m_e]_+, \quad \Phi_e^- = \kappa_e x_q [-m_e]_+.$$

A tagged carrier at the eligible origin sees rate Φ_{qr}/x_r . This denominator follows from equal packet-carrier response and the number of eligible carriers. It is not inserted as a Bell rule into the contact chemistry.

The exact balance between census, queued charge, and coherent weights is

$$x(t) + BZ(t)/N - w(t) = x(0) - w(0) + Be^N(t), \quad \|e^N\|_\infty \leq C/N. \quad (8.4)$$

For a fixed finite graph and horizon, bounded-variation currents and $\mu_N \rightarrow \infty$, $\mu_N/N \rightarrow 0$ give integrated directional-flux and census convergence. The proof uses a positive population cutoff for a companion queue, scalar monotonicity, a martingale estimate, and a separate low-population charge bound. A full proof of the implication used here is given in Appendix A; it is important because replacing low weights by a fixed positive lower bound would omit nodes and dark intervals.

Let

$$\epsilon_F = \mathbb{E} \int_0^T \sum_e \left(|\Phi_e^+ - [J_e]_+| + |\Phi_e^- - [-J_e]_+| \right) dt, \quad \epsilon_x = \mathbb{E} \sup_{t \leq T} \|x(t) - w(t)\|_1.$$

For a tag initialized with $\nu \leq Cw(0)$, a coupling with the minimal Bell path gives, along suitable regular levels $\varepsilon \downarrow 0$,

$$d_{\text{TV}}(\text{Law } X^N, \text{Law } X^B) \leq b_C(\varepsilon) + \frac{2\epsilon_x + 2\epsilon_F}{\varepsilon} + \frac{C_G J_* T}{\varepsilon^2} \epsilon_x, \quad (8.5)$$

where $b_C(\varepsilon)$ bounds visits to low-weight sectors. It tends to zero along levels controlled by the variation of w . The target has finite expected jump count and does not remain at a sector as its weight vanishes. The resulting convergence is on complete physical-time paths, including reversals and null intervals.

For a fixed finite clock circuit, the coherent weight envelopes are explicit and unimodal. The source proof yields a conservative uniform rate $O(N^{-1/70})$ with

$$\mu_N = N^{1/2}, \quad \delta = N^{-1/7}, \quad \varepsilon = N^{-1/35},$$

and appropriately faster gas and recombination scales [34]. This is not a rate uniform over growing circuits, increasing reference dimensions represented as new graph coordinates, or unlimited storage times.

8.5 A static finite Hamiltonian for the whole programme

Let $U_0, \dots, U_{\ell-1}$ be material gates, including all receivers and an isolated reference where required. Put $V_0 = I$, $V_n = U_{n-1} \cdots U_0$, and $c_n = \sqrt{(n+1)(\ell-n)}$. The static clock Hamiltonian is

$$H_F = \hbar \nu \sum_{n=0}^{\ell-1} c_n \left(|n+1\rangle \langle n| \otimes U_n + |n\rangle \langle n+1| \otimes U_n^\dagger \right). \quad (8.6)$$

This familiar circuit-to-Hamiltonian architecture is related to reversible computation and engineered transfer chains [17, 11]. The particular propagation identity needed here follows directly.

Theorem 8.1 (Finite autonomous propagation). *From $|0\rangle \otimes \psi$, the wave at time t is*

$$\Psi(t) = \sum_{n=0}^{\ell} (-i)^n \sqrt{\binom{\ell}{n}} \cos^{\ell-n}(\nu t) \sin^n(\nu t) |n\rangle \otimes V_n \psi.$$

At $T = \pi/(2\nu)$ the complete programme has been applied. Moreover $H_F + \hbar\nu\ell I \geq 0$.

Proof. Conjugate by $D = \sum_n |n\rangle\langle n| \otimes V_n$. The gate factors disappear, leaving a weighted clock chain tensor the identity. That clock chain is the restriction of $\hbar\nu \sum_{j=1}^{\ell} X_j$ to the permutation-symmetric subspace of ℓ qubits. Expanding $(\cos \nu t |0\rangle - i \sin \nu t |1\rangle)^{\otimes \ell}$ gives the coefficients. Its spectrum lies in $[-\hbar\nu\ell, \hbar\nu\ell]$, proving the lower bound. $\square$

Theorem 8.2 (A historical archive at a monomial cut). *Suppose U_m is monomial in the complete material basis: $(U_m)_{yx} = e^{i\vartheta_x} \mathbf{1}_{\{y=\pi(x)\}}$. In the Bell process of Equation (8.6) in equilibrium, the clock cut $m \rightarrow m+1$ is crossed exactly once before T , almost surely. At that crossing the actual material label changes by π . If the permutation copies a working label into a previously blank archive and every later gate preserves the archive label, the archive contains the actual working label at that crossing for the rest of the first pass.*

Proof. Write the scalar clock coefficient as $(-i)^n a_n(t)$ with $a_n > 0$ inside the first pass. Fine current across the cut is

$$2\nu c_m a_{m+1} a_m \operatorname{Re} \left[(V_{m+1}\psi)_y^* (U_m)_{yx} (V_m\psi)_x \right].$$

For the monomial gate its nonzero real factor is $|(V_m\psi)_x|^2$, so every fine current is forward. The path begins below and ends above the cut with probability one, hence crosses it exactly once. Its permitted edge implements π . Later edges preserve the archive by hypothesis. Integrating the fine current gives the pre-crossing label law $|(V_m\psi)_x|^2$. $\square$

The monomial hypothesis cannot be dropped. A general unitary may have backward fine currents despite a forward coarse clock flux. Nor does the archive theorem extend indefinitely: the finite clock eventually reverses and can unwrite the programme. The first-pass domain is the physical retention claim, and the same complete-path comparison transfers it to the finite pilot approximation with its stated error.

8.6 The observer-independent conclusion

The source action, export, contact geometry, initial ensemble, recombination, census tracking, and autonomous record programme form a complete chain under the declared pilot constitution. No stage waits for subjective noticing. Its dependence on new force and preparation laws is explicit, and finite members are controlled approximations to the limiting Bell path rather than exact finite-resource realizations of that path.

This conclusion is sufficient for the consciousness monograph's purpose. A psychophysical assignment can concern the same realized process without being invoked to make the physical history well defined. The existence of two distinct constitutions with different microscopic commitments also cautions against deriving source identity from outward operational similarity.

Chapter 9

Resources, continued operation, and agency

9.1 Finite records are physical resources

A record is not merely a scalar assigned to a state. It requires a physical carrier, distinguishable alternatives, a writing interaction, a retention interval, and an allowed way of reading or reusing it. Exact copying of a known orthogonal alphabet can be unitary: $|a\rangle|0\rangle \mapsto |a\rangle|a\rangle$. This is not a universal cloner of unknown states. Superpositions evolve to correlated states rather than two independent copies.

A reversible implementation of a many-to-one logical step must retain enough information to distinguish its inputs. Bennett’s reversible computation and Landauer’s analysis of logical irreversibility provide the standard background [4, 23]. Here the finite consequence can be stated directly. If n independent binary historical alternatives must remain perfectly distinguishable in a retained carrier of dimension D , then $2^n \leq D$. The bound follows from orthogonality of perfectly distinguishable records. It is not a theorem that a finite system cannot recur, preserve a fixed record, or reuse a reversible working register.

A finite ready bank with N cells can support at most N operations that each consume a distinct ready cell and never restore global readiness. SWAP reset restores a local register by spending a receiver. The availability of another environment does not restore the bank unless an admitted coupling actually transfers a compatible resource. “Open” is a physical interaction condition, not an accounting word that creates capacity.

9.2 Operational capacity

Fix a protocol class $\mathfrak{P}$, retained obligations, and a finite horizon. For a complete state s , define $C_{\mathfrak{P}}(s)$ to be the largest number of protocol completions that one admissible strategy can guarantee on that horizon without violating the obligations. The number is zero if no completion is reachable. Other risk tolerances lead to other capacities; they must be specified rather than identified with this worst-case definition.

The complete state includes the remaining physical time and the retained obligations; every successor carries the decremented budget. Equivalently, a completion of duration $d(s, u, s')$ obeys

$$C_{\mathfrak{P},\tau}(s) \geq 1 + \inf_{s' \in \mathcal{R}(s,u)} C_{\mathfrak{P},\tau-d(s,u,s')}(s'),$$

provided every admitted successor fits the remaining horizon. In the abbreviated notation below this clock update is part of s' .

Capacity depends on resources, reachability, and compatibility together. Two states may have the same number of apparently blank cells while only one admits the coupling needed to use them. A locally ready marginal can also be correlated with a returning memory, so it is not the same preparation as an independent blank. Capacity therefore is not universally the expectation of a single source operator, nor a thermodynamic free energy without a specified Hamiltonian, bath, temperature, and work convention.

For a deterministic admissible preprocessing F that can be prefixed without violating obligations and with any time or resource consumption included in its successor state, $C_{\mathfrak{P}}(s) \geq C_{\mathfrak{P}}(F(s))$. For a guaranteed consuming step with successor set $\mathcal{R}(s, u)$, a valid dynamic-programming lower bound is

$$C_{\mathfrak{P}}(s) \geq 1 + \inf_{s' \in \mathcal{R}(s, u)} C_{\mathfrak{P}}(s').$$

Both statements follow by concatenating strategies. Additivity under composition needs independently usable resources and separable obligations; correlations or shared actuators can invalidate it. This is why a resource count should not be promoted to a universal consciousness magnitude.

9.3 Viability and the orientation of repair

Let G be a finite preservation region and T_u controlled transition kernels. Define

$$V_0 = G, \quad V_{n+1} = \{s \in V_n : \exists u \text{ supp } T_u(s, \cdot) \subseteq V_n\}.$$

The decreasing sequence stabilizes in finitely many removals. Its terminal set is the greatest subset of G from which an admissible policy can preserve the region indefinitely in the finite model. A selecting policy is obtained by choosing one witnessing control at each surviving state. This is a theorem about available policies.

It does not say the actual controller chooses that policy. Let (d, f) be a defect bit and repair-fuel bit. Admit idle and repair, with repair taking $(1, 1)$ to $(0, 0)$ and idle leaving it unchanged. Both selectors coexist with the same resource inventory and possible continuations. One is repair-oriented and the other is not. Adding the same awareness aspect to both does not select between them. Minimal Repair Orientation, when claimed, is a policy law or a consequence of additional physical dynamics, not an automatic consequence of recurrence, viability, or awareness.

The distinction is philosophically significant. Treating awareness as primitive does not attribute desire for survival, rich valuation, or benevolent action to it. Such properties may be realized in a vessel, but they require their own organization. This prevents a source-aspect ontology from silently turning into a theory of agency or ethics.

9.4 Agency and stable decisions

The agency manuscripts define a functional target: usable world modeling, counterfactual evaluation, and reliable action selection [31]. This target is independent of a phenomenal predicate. A controller can implement a specified comparison procedure without that procedure constituting evidence of experiential presence.

A quantum no-go must specify its input encoding and allowed resources. Unitary linearity rules out a generic nonlinear winner-selection rule on arbitrary unknown amplitude-encoded inputs. It does not rule out reversible evaluation of a classical decision function on orthogonally

encoded data. A reversible circuit can retain the input and computational garbage while writing a definite decision label for each basis input. Superposed inputs then yield the corresponding correlated superposition. The physical interpretation of an actual outcome belongs to the event constitution, not to a general slogan that quantum systems cannot decide.

Similarly, choosing a diagonal conditional expectation builds a record algebra into the construction. Showing that its image is diagonal does not derive the basis. A constant contraction $T(\rho) = I/2$ has one fixed density operator but does not select one of its two classical outcome labels. Banach uniqueness is uniqueness of a mathematical fixed state, not actualization of one sample. These distinctions allow useful fixed-point stability results without inheriting an invalid universal derivation of decisions or Born weights.

9.5 Physical stability and philosophical relevance

The constructive lesson is neither that agency is illusory nor that all agency is conscious. Model formation and action can be physically causal processes with stable records and real consequences. The consciousness theory asks whether and how those processes are accompanied by experiential presence under a specified aspect law. It can therefore admit experience without advanced deliberation and advanced deliberation without an independently established experiential attribution.

This separation also protects the contemplative discussion from a false dichotomy. Reduced identification with an ego need not abolish practical action, and practical competence need not certify enlightenment. The material theory identifies what can be implemented and retained. The philosophical account concerns how self-description and presence are related. Their integration requires a bridge, but not the collapse of one vocabulary into the other.

Part III

The Mathematics of Realized Vessels

Chapter 10

Internal representation and self-referential records

A vessel can retain information about its own operation without possessing an elaborate narrative identity. This chapter asks when a nominated internal quantity has an observable representative, whether that representative is unique, and how robustly it can be recovered. These questions concern the physical organization available for manifestation and the evidence used to describe it.

10.1 Representing a specified target

Self-Referential Record Closure (SRC) concerns the representation of nominated self-related quantities in a declared observable carrier [36]. It is valuable precisely when the target, state family, algebra, and actual record dynamics are kept separate. Let W be a finite-dimensional real vector space of Hermitian operators on $\mathcal{H}$, let $\rho_1, \dots, \rho_m$ be admissible preparations, and define

$$\mathcal{E} : W \rightarrow \mathbb{R}^m, \quad \mathcal{E}(A)_j = \text{tr}(\rho_j A).$$

A vector y specifies the desired target expectations.

Theorem 10.1 (Witness representation and ambiguity). *A witness $A \in W$ for y exists exactly when $y \in \text{im } \mathcal{E}$. When a witness A_0 exists, the complete set of witnesses is $A_0 + \ker \mathcal{E}$. It is a singleton exactly when the state family separates W .*

Proof. Existence is the definition of the image of a linear map. Two solutions differ by an element of its kernel, and adding a kernel element preserves the target values. Injectivity is exactly separation of all observables in the nominated carrier. $\square$

For a convex family of finite-dimensional states, an affine target extends to its affine span and admits a Hermitian representative after incorporating the identity. A general function on a curved state manifold is not automatically affine. The distinction matters for phenomenal coordinates: labeling a nonlinear target an internal expectation does not establish that an operator represents it.

The effective witness is naturally an equivalence class modulo $\ker \mathcal{E}$. This is a vector-space quotient. It need not be an algebra quotient. At $\rho_* = I/2$, $\text{tr}(\rho_* Z) = 0$ but $\text{tr}(\rho_* Z^2) = 1$. Thus null expectation is not preserved under multiplication. An algebra generated by witness representatives is meaningful only after the representatives and the relevant multiplication structure have been fixed. Without separation, different representatives can generate different raw algebras while agreeing on all nominated expectations.

10.2 Centralizers and their scope

For a faithful density matrix $\rho_* = \sum_{\alpha} \lambda_{\alpha} P_{\alpha}$, the finite centralizer is

$$\mathcal{B}(\mathcal{H})^{\rho_*} = \{A : [A, \rho_*] = 0\} = \bigoplus_{\alpha} \mathcal{B}(P_{\alpha} \mathcal{H}).$$

The real dimension of its Hermitian part is $\sum_{\alpha} (\text{rank } P_{\alpha})^2$. In finite dimensions this follows by examining matrix blocks between unequal eigenvalues. It provides a natural carrier for certain stationary or modularly invariant targets. It does not automatically supply physical retention.

Indeed, take $\rho_* = I/2$ and an actual Hamiltonian $H = X$. Every observable belongs to the centralizer, but Z changes under the physical Heisenberg evolution because $[X, Z] \neq 0$. Modular invariance is not the same as invariance under the admitted material dynamics. A physical record theorem must specify those dynamics, its interval, and an error criterion. The retention result for the held massive archive in Chapter 7 is such a theorem; centralizer membership alone is not.

For a fixed Hermitian witness A , the von Neumann algebra $W^*(A)$ is the smallest such algebra containing it. For several noncommuting witnesses the generated algebra need not be abelian. Neither its minimality nor its dimension identifies the number of experiential subjects. The algebra concerns the representation of specified observables. Subject attribution enters through the separate law in Part IV.

10.3 Positive witnesses and finite calibration

Suppose the physical use requires $A \geq 0$. The feasible set becomes the intersection of the positive cone with the affine witness set. Nonemptiness is an additional feasibility question. When one calibrated preparation obeys $\rho_0 \geq cI$ with $c > 0$ and has target g_0 , every feasible positive witness satisfies

$$c \text{tr } A \leq \text{tr}(\rho_0 A) = g_0, \quad \|A\| \leq \text{tr } A \leq g_0/c.$$

The feasible set is then bounded and closed, hence compact in finite dimensions. Extreme witnesses exist when this set is nonempty. This is a useful route to constrained reconstruction, not a proof that the positivity requirement is automatically satisfied.

Choose a Hilbert–Schmidt orthonormal basis $G_1, \dots, G_d$ of W . The evaluation matrix is $A_{jk} = \text{tr}(\rho_j G_k)$, and a witness has coefficients a satisfying $Aa = y$. Full column rank is a finite calibration condition. A numerical fit is evidence for this condition only to the extent that preparation errors, conditioning, and numerical tolerances are controlled. Rank that depends on a tiny singular value is not robust identifiability.

Theorem 10.2 (Robust witness identification). *Assume A has smallest singular value $\sigma > 0$, $y = Aa$, and measured quantities are $\hat{A} = A + E$, $\hat{y} = y + e$, with $\|E\|_2 < \sigma$. The least-squares coefficients $\hat{a} = \hat{A}^+ \hat{y}$ satisfy*

$$\|\hat{a} - a\|_2 \leq \frac{\|e\|_2 + \|E\|_2 \|a\|_2}{\sigma - \|E\|_2}. \quad (10.1)$$

Proof. For any unit vector v , $\|\hat{A}v\| \geq \|Av\| - \|Ev\| \geq \sigma - \|E\|_2$. Thus $\hat{A}$ has full column rank and $\|\hat{A}^+\| \leq (\sigma - \|E\|_2)^{-1}$. Since $\hat{A}^+ \hat{A} = I$, $\hat{a} - a = \hat{A}^+(e - Ea)$. Taking norms proves the bound. $\square$

For a density operator ρ , the resulting expectation error is at most $\|\hat{a} - a\|_2$, because $\|\rho\|_{\text{HS}} \leq 1$. If $D(\rho_j, \hat{\rho}_j) \leq \delta_j$, then

$$|E_{jk}| \leq 2\delta_j \|G_k\|_\infty, \quad \|E\|_2 \leq 2 \sqrt{\sum_j \delta_j^2} \sqrt{\sum_k \|G_k\|_\infty^2}.$$

A physical preparation theorem can therefore supply an input to a witness-identification theorem. A complete-path TV estimate cannot replace a preparation trace-distance estimate unless a common physical mapping justifies the transfer.

10.4 Local jets: finite existence without a universal order bound

A smooth or analytic family of states can generate local calibration data by differentiation. In a fixed coordinate chart, derivatives of $g(s) = \text{tr}(\rho_s A)$ are linear functionals of A . Full jets transform with lower-order terms under a change of chart; higher coordinate derivatives should not be treated as independent tensors without a connection or an appropriate jet formalism.

Let W be finite dimensional and F_k the common kernel of all derivative evaluations through order k at a point. The chain $F_0 \supseteq F_1 \supseteq \dots$ has at most $\dim W$ strict decreases. It eventually stabilizes, but the index of its final decrease need not be bounded by $\dim W$. A plateau is not a stopping certificate.

Example 10.3 (Arbitrarily delayed local information). For any integer $m \geq 1$, set

$$\rho_s = I/2 + (s^m/4)Z, \quad |s| < 1.$$

This is a faithful analytic qubit family. The witness Z has expectation $s^m/2$. Every derivative of order less than m vanishes at zero, while the m th derivative does not. The Hermitian carrier dimension remains four as m grows.

If all expectation functions are analytic on a connected domain, the intersection of all jet kernels equals the global null space: zero Taylor series gives local vanishing, and analytic continuation gives global vanishing. Finite dimensionality then ensures that some finite set of derivative functionals spans the required information. It does not provide a universal maximum derivative order or an effective stopping rule without further polynomial, frequency, or differential-equation bounds. This precise version preserves the usefulness of local reconstruction without an unjustified finite-order guarantee.

10.5 Self-reference without phenomenal promotion

A record becomes self-referential in the operational sense when its target concerns its own carrier or ongoing process and the record participates in later dynamics. It need not be secret from an external observer. A duplicate or prediction of its value does not remove its internal causal role. Conversely, a hidden internal variable with no effective read or control path is not a self-witness merely because it is inaccessible to outsiders.

The SRC results therefore describe a particular family of vessel capabilities. They support the construction of stable internal descriptions and the testing of their sufficiency. They do not deduce that those descriptions are experienced. This is not a defect to conceal; it is the boundary that makes a subsequent explicit psychophysical law intelligible.

Chapter 11

Regulation, unity, and the functional self

11.1 A typed regulation architecture

The Conscious Record Realization framework (CRR) organizes preservation, access, valuation, and control as distinct data [42]. In a finite deterministic specialization, let P be a state/record space, $\epsilon : P \rightarrow E$ a defect readout, and $Z \subseteq E$ the acceptable region. The preserved region is the pullback

$$K = \epsilon^{-1}(Z).$$

Let $a : P \rightarrow O$ be an access map, $\nu : E \rightarrow V$ a valuation, $\kappa : O \times V \rightarrow U$ a selector, and $\alpha : P \times U \rightarrow P$ an admitted update. Their closed loop is

$$F(p) = \alpha(p, \kappa(a(p), \nu(\epsilon(p)))). \quad (11.1)$$

The corresponding relational version existentially composes the same typed arrows. Correct typing constructs a loop; it does not prove its preservation, recovery, or phenomenal properties.

The pullback has the expected universal property: $f : Y \rightarrow P$ factors through K exactly when ϵf factors through Z . This is an elementary but useful separation of a preservation obligation from the mechanism that realizes it. An invariant subspace or conserved quantity alone supplies neither effective access nor an operative preference about defects.

11.2 Access and valuation are different

Endogenous access means that the readout belongs to the process's actual causal diagram and can affect its later operation. It is not defined by secrecy. Two complete-state conditions can witness effective access when changing the access value, while holding other admitted parents fixed, changes a later law. If such independent interventions are physically impossible, the claim must instead be stated for the admissible joint interventions. A marginally inert statistic does not justify deleting a correlated causal parent.

Valuation is not determined by the preservation region. On $E = \{0, 1, 2\}$ with acceptable state 0, two value maps can agree that 0 is good while ranking defects 1 and 2 oppositely. The physical region $\epsilon^{-1}(0)$ remains unchanged. A binary defect indicator is therefore a valid specialization, not a derivation of rich preferences. The repair-policy distinction in Chapter 9 is another instance of the same principle.

Within CRR, the term “access awareness” names an operational type. In this monograph it must not be confused with A0's primitive awareness or with experiential presence. The

translation is deliberate: an effective regulation type is one possible physical organization through which a proposed perspective can be expressed, but the type name does not discharge the phenomenal obligation.

11.3 Weak unity, strong unity, and assembly

Several locally realized modules need not form one realized global process. Their interfaces can be incompatible; their controls can conflict; their valuations may have no admitted common refinement. A global assembly therefore requires more than juxtaposition. In CRR, weak unity concerns a full-support global access/valuation/control loop. Strong unity adds an indecomposability requirement and appropriate invariant structure. A phenomenal claim adds a separate binding or presentation obligation.

These notions are useful for analyzing distributed systems. Hardware separation does not automatically prevent integration, and residence in one body or machine does not automatically establish it. The relevant question is whether the nominated interfaces and global loop exist and satisfy their conditions. This is a source-exact assembly question, not an inference from spatial proximity.

The strongly connected components used by SPC-2 will not be identified with CRR strong unity as a theorem. They serve a different role. A minimal recurrent core may have no rich valuation, repair policy, or reflection architecture. If a CRR-enriched subject is claimed, those extra obligations must be realized separately. Conversely, a globally regulated system need not receive a phenomenal attribution before a psychophysical law has been adopted.

11.4 Functional selfhood and episode identity

An invariant functional self can be represented by a subobject $I \hookrightarrow P$ with $F(I) \subseteq I$. More elaborate versions preserve an object up to an allowed transformation rather than pointwise. Such a self-invariant is derived from the supplied dynamics and chosen invariant doctrine. It does not imply a permanently unchanged ego or a unique experiential identity.

A self-model can be inaccurate, transient, or causally idle. A memory can record a history without generating any self-invariant. Reflection requires a second-order representation of some part of the loop that is itself available for correction. These distinctions preserve the philosophical possibility of experience without elaborate self-description, while allowing a technical theory of increasingly sophisticated self-related capacities.

The appropriate classification is generally a partial order of capabilities. One system may have strong retention and weak reflection; another may have broad access but unstable long-term control. A scalar ranking requires a chosen weighting. There is no universal consciousness score hidden in the mere existence of a capability lattice. No conclusion about moral worth follows from such a weighting.

Chapter 12

Finite incidence, recurrent response, and physical access

This chapter provides a finite example of why the provenance of a physical structure matters in addition to its spectrum. The calculations retain which parts of a carrier arose from a designated source region and which were added. They illustrate the realization discipline used later in R^* : an invariant numerical summary need not retain every distinction relevant to a proposed physical or phenomenal target.

12.1 Source-sensitive response carriers

A finite source-response construction begins with a rooted partial order, not an empirically fitted spectrum. The retained root is $0 < 1$, and the fresh loci are 2, 3. The order complex contains every nonempty chain, including compositional higher simplices. With a counting-normalized orthonormal simplex basis, the signed boundary operator d obeys $d^2 = 0$. Define

$$D = d + d^\dagger, \quad K = D^2 = dd^\dagger + d^\dagger d.$$

This is finite Hodge mathematics. Its source status depends on the admitted response principles and Hilbertization, as the constructive O1 source explicitly states [39].

Separate chains lying entirely in the root support from mixed-interface and fresh chains. Let P_R project onto root-supported chains and $P_F = I - P_R$ onto the remainder. The decomposition is assigned by provenance before the spectral calculation. It is not chosen to produce a desired gap.

For an eigenprojection E_λ of K , define $B_\lambda = P_R E_\lambda P_F$. The positive RSM-active spectrum is

$$\{\lambda > 0 : B_\lambda \neq 0\}.$$

This differs from the ordinary spectral gap because it depends on the geometry of the retained/fresh sectors. A zero-eigenvalue cross block may exist while being excluded by the word “positive” in this definition.

Theorem 12.1 (Finite return identity). *For finite $s, t \geq 0$,*

$$P_F e^{-sK} E_\lambda P_R E_\lambda e^{-tK} P_F = e^{-\lambda(s+t)} B_\lambda^\dagger B_\lambda.$$

It is positive and nonzero exactly when $B_\lambda \neq 0$.

Proof. Since E_λ is the whole orthogonal eigenprojection, $e^{-tK} E_\lambda = e^{-t\lambda} E_\lambda$. Move only these commuting factors, not P_R , through the exponentials. The remaining product is $P_F E_\lambda P_R E_\lambda P_F = B_\lambda^\dagger B_\lambda$. A matrix $B^\dagger B$ vanishes exactly when $B = 0$. $\square$

Degenerate eigenspaces cause no ambiguity because an arbitrary eigenvector is not substituted for E_λ . In an infinite-dimensional model the same identity applies to a genuine finite-eigenvalue spectral atom, but singleton projectors can miss the entire continuous spectrum. The physical interpretation also requires an instrument realizing the intervening operations and a retained record. An operator product is not, by itself, a physical recursive experiment.

12.2 Five exact profiles

The source relations and independently recomputed results are shown in Table 12.1. The common relation $0 < 1$ is implicit in each row.

Profile	Additional relations	$\dim \mathcal{H}$	Spectrum of K
$V2_{02}$	$0 < 2, 0 < 3$	7	$0^1, 1^4, 4^2$
$V2_{34}$	$2 < 1, 3 < 1$	7	$0^1, 1^4, 4^2$
$V2_{04}$	$0 < 2, 0 < 3, 1 < 2, 1 < 3$	11	$0^1, 2^4, 4^6$
$V2_{11}$	$0 < 2, 0 < 3, 2 < 1, 3 < 1$	11	$0^1, 2^4, 4^6$
$V2_{29}$	$2 < 0, 2 < 1, 3 < 0, 3 < 1$	11	$0^1, 2^4, 4^6$

Table 12.1: Finite local source-response carriers. Exponents denote multiplicity, not powers.

In every case P_R has rank three, supported by the two root vertices and root edge. For the seven-dimensional profiles, $\text{rank } B_0 = 1$, $\text{rank } B_1 = 2$, and $\text{rank } B_4 = 2$. For the eleven-dimensional profiles, the corresponding ranks are 1, 0, 1 at eigenvalues 0, 2, 4. The positive active gap is therefore one in the former pair and four in the latter three.

For exact calculation, the eigenprojectors can be obtained without choosing eigenvector bases:

$$E_\lambda = \prod_{\mu \in \text{Spec } K, \mu \neq \lambda} \frac{K - \mu I}{\lambda - \mu}. \quad (12.1)$$

The verification package supplies all boundary matrices, D, K, P_R , projectors, and cross ranks. It checks all 24 signed vertex relabellings for each of the five profiles. A relabelling must transport the root support and simplex orientation; an unsigned permutation that changes orientation is not a valid matrix comparison.

Equal spectra do not identify source objects. Different differentials, provenance projections, and admitted interactions can share K or its spectrum. Even for a fixed diagonal K , choosing a commuting projector gives no cross block, while a rotated projector can give one. The calculation therefore supports a source-sensitive response classification, not a universal consciousness threshold.

12.3 Interaction-relative sector invariance

Let $P = P_R$, $J = i[K, P]$, and in the eleven-dimensional examples set $\Pi = J^2/4$. Exact calculation gives

$$\Pi^2 = \Pi, \quad \text{rank } \Pi = 2, \quad [K, \Pi] = [P, \Pi] = 0, \quad \text{rank}[D, \Pi] = 2.$$

The block $\Pi D(I - \Pi)$ has rank one and singular value two. Thus a D -generated actuator would mix part of the two support sectors while conserving K , but an action algebra generated by K, P, J preserves them.

The P -commutation has a general explanation. Relative to $P\mathcal{H} \oplus (I - P)\mathcal{H}$, write

$$K = \begin{pmatrix} A & B \\ B^\dagger & C \end{pmatrix}, \quad J = \begin{pmatrix} 0 & -iB \\ iB^\dagger & 0 \end{pmatrix}.$$

Then $J^2 = \text{diag}(BB^\dagger, B^\dagger B)$, so P commutes with J^2 and its support. The additional condition needed for the full $K/P/J$ obstruction is preservation of that support by K . In the specified profiles it holds.

If every Kraus operator of an admitted instrument commutes with Π , an input supported in one extreme sector remains there on every nonzero conditional branch. Such a trace-preserving channel preserves $\text{tr } \rho \Pi$ unconditionally. Conditioning a mixed-sector input can change its normalized sector weight, but cannot create support absent from an extreme input. This is an interaction-relative conservation law, not a fundamental superselection of the whole matrix algebra.

12.4 Actuation, source drift, and seed-specific mobility

Mathematical membership of D in an operator algebra does not expose a D port or supply a Hamiltonian coupling. Even the conditions of self-adjointness, source covariance, and K conservation leave a family $h(K)D$. Selecting the primitive D as an actuator is a possible realization law, not a consequence of symmetry alone.

For an actual Hermitian actuator A and state $\rho = \Pi \rho \Pi$, leakage has the expansion

$$\text{tr}[(I - \Pi)e^{-itA}\rho e^{itA}] = t^2 \text{tr}[\rho A(I - \Pi)A] + O(t^3).$$

The first derivative vanishes. A nonzero commutator guarantees some possible mixing, not first-order escape of every extreme seed. In the type-11 crossing channel there are unit vectors $u \in \text{im } \Pi$, $v \in \ker \Pi$ with $Du = 2v$, $Dv = 2u$, hence $e^{-itD}u = \cos(2t)u - i\sin(2t)v$. An additional active zero-mode line is D -dark. These facts are about reachability in a supplied action grammar.

Similarly, an isolated tap generated by $\eta J \otimes X$ and a tap with continuing source drift are different operations. In the active two-level sector, the latter has, after a scalar shift,

$$H = 2Z \otimes I + 2\eta Y \otimes X, \quad H^2 = 4(1 + \eta^2)I.$$

Its memory-one effect is

$$\frac{\eta^2}{1 + \eta^2} \sin^2(2t\sqrt{1 + \eta^2})\Pi.$$

It is not a perfect Π measurement at finite η . Switching off or refocusing drift requires an admitted control and resources. An imperfect nontrivial tap can still provide useful conditional information, but a family of ensemble evaluations is not exact probability estimation from a single unknown specimen.

12.5 Robust spectral access

Let a contour Γ isolate a spectral cluster of Hermitian K at distance $a > 0$ from its spectrum. For $K' = K + \Delta K$, $\|\Delta K\| = \varepsilon < a$, the Riesz projections obey

$$\|P_\Gamma(K') - P_\Gamma(K)\| \leq \frac{\text{length}(\Gamma)}{2\pi} \frac{\varepsilon}{a(a - \varepsilon)}. \quad (12.2)$$

The resolvent identity gives an integrand norm at most $\varepsilon/[a(a - \varepsilon)]$, and integration proves the bound. Write $E = P_\Gamma(K)$ and $E' = P_\Gamma(K')$ for these spectral projections. If the provenance projector also changes from P to P' by δ , then

$$\|P'E'(I - P') - PE(I - P)\| \leq \|E' - E\| + 2\delta.$$

A cross block larger than this uncertainty remains nonzero; rank claims need further control. Spectral robustness is not evidence of awareness.

Chapter 13

Fixed points, canonicalization, and coherent–dissipative structure

A realized vessel can maintain some distinctions while other modes relax or change. The CSCF construction places coherent and dissipative responses on one spectral carrier, giving a precise mathematical setting for studying such coexistence. We first establish a local convergence result, then distinguish its meaning from the shared-carrier field construction and from the additional experiential constitution.

13.1 A local contraction theorem

Let C be a nonempty closed convex subset of a real Hilbert space and J differentiable on the relevant domain. Assume its gradient is L -Lipschitz and strongly monotone with constant $m > 0$:

$$\langle \nabla J(x) - \nabla J(y), x - y \rangle \geq m\|x - y\|^2.$$

For $0 < \eta < 2m/L^2$ define $T(x) = P_C(x - \eta \nabla J(x))$. Projection is nonexpansive, and therefore

$$\|Tx - Ty\|^2 \leq (1 - 2\eta m + \eta^2 L^2)\|x - y\|^2.$$

Let $q = \sqrt{1 - 2\eta m + \eta^2 L^2} \in [0, 1)$. Banach's theorem supplies a unique fixed point and geometric convergence. The projection variational inequality identifies that point with the constrained minimizer of J under the stated convexity assumptions [33].

The residual gives a practical certificate:

$$\|x - x_*\| \leq \frac{\|x - Tx\|}{1 - q}. \quad (13.1)$$

Indeed $\|x - Tx\| \geq \|x - x_*\| - \|Tx - Tx_*\| \geq (1 - q)\|x - x_*\|$. Finite arrival is possible, for example when $J(x) = x^2/2$ and $\eta = 1$. Contraction does not imply that a trajectory remains forever outside its fixed point.

The theorem describes the chosen functional and map. For any chosen destination x_0 , the functional $J_{x_0}(x) = \|x - x_0\|^2/2$ supplies a contraction toward it. A fixed point can encode an inaccurate belief or an undesirable configuration. Equating convergence with truth, clarity, or enlightenment requires an independently calibrated relation to those targets. The mathematical word “canonical” cannot supply that relation.

13.2 Observable certification

Suppose a readout Ψ is Lipschitz, with lower constant $m_\Psi > 0$ on the nominated regime. If $z_k = \Psi(x_k)$ and $x_{k+1} = T(x_k)$, then

$$\|x_k - x_*\| \leq \frac{\|z_{k+1} - z_k\|}{m_\Psi(1 - q)}.$$

If each observed z_k has error at most δ , replace the numerator by the observed step size plus 2δ . The estimate requires all relevant states to remain within the regime on which the lower bound is valid.

A five-channel readout is not automatically sufficient for an arbitrarily high-dimensional source. Let P_E project onto the measured subspace. If the relevant differences obey the cone condition

$$\|(I - P_E)(x - y)\| \leq \kappa \|P_E(x - y)\|,$$

then $\|P_E(x - y)\| \geq \|x - y\|/\sqrt{1 + \kappa^2}$. This supplies a legitimate lower constant. Without the cone or an equivalent restriction, stabilization of a few coordinates can coexist with change in unobserved directions. The semantic names of channels do not change this geometry.

For global selectors a clean result is available on a supplied compact metric domain. Successively minimize a countable separating family of continuous functions on nested nonempty compact sets. Compactness gives a nonempty intersection, and separation makes it a singleton. This proves uniqueness relative to the domain and function family. It does not establish compactness of a geometric model from insufficient regularity bounds, or identify its selected point with source awareness.

13.3 One operator, two analytic regimes

The strengthened Canonical Spectral Curvature Field construction starts with a specified densely defined nonnegative closed quadratic form and its self-adjoint representing operator K [32]. The analytic family

$$F_K(z) = e^{-zK}, \quad \operatorname{Re} z \geq 0, \quad (13.2)$$

has a contractive real ray e^{-tK} and a unitary imaginary boundary e^{-itK} . Both use the same spectral measure. For $u \in \mathcal{H}$,

$$\langle u, F_K(z)u \rangle = \int_{[0, \infty)} e^{-z\lambda} d\mu_u(\lambda).$$

On the open half-plane the family is bounded and holomorphic in operator norm, with the usual spectral derivative bounds. At the boundary it is strongly continuous; norm continuity there is not automatic for unbounded K .

This is a precise coherent–dissipative correspondence on a declared carrier. It does not identify the operator with a brain’s logical reasoning, an unconscious mind, or a language model’s latent search. Quantum coherence, logical consistency, stable memory, and explicit verbal reasoning are different properties. Dissipation can erase distinctions or stabilize selected ones, depending on the physical model.

Nor is $\rho \mapsto e^{-tK}\rho e^{-tK}$ generally trace preserving. Treating it as a physical channel requires an instrument, loss branch, or another justified construction. Normalizing it produces a nonlinear map. The physical interpretation of a form and its heat kernel therefore requires the same interaction discipline as any other source operator.

13.4 A shared spectral carrier with a general coherent symbol

The more general Atlas formulation specifies a real spectral symbol h and a damping coefficient $\gamma \geq 0$ in addition to K [40]. This preserves a single spectral carrier while allowing oscillation and decay to have different dependence on its spectral parameter.

Proposition 13.1 (Coherent and dissipative responses of one spectral datum). *Let $K \geq 0$ be self-adjoint, and let $h : [0, \infty) \rightarrow \mathbb{R}$ be Borel and finite E_K -almost everywhere. Then*

$$S_t = e^{-\gamma t K}, \quad U_t = e^{-ith(K)}, \quad W_t = S_t U_t = \int e^{-\gamma t \lambda - ith(\lambda)} dE_K(\lambda)$$

define a strongly continuous contraction semigroup S , a strongly continuous unitary group U , and a strongly continuous contraction semigroup W . The responses commute. For $j \in L^1([0, T]; \mathcal{H})$ the unique mild driven solution is

$$\psi(t) = W_t \psi(0) + \int_0^t W_{t-s} j(s) ds.$$

Whenever this solution is strong and lies in $D(K) \cap D(h(K))$,

$$\dot{\psi} = -(\gamma K + ih(K))\psi + j, \quad \frac{d}{dt} \|\psi\|^2 = -2\gamma \langle \psi, K\psi \rangle + 2 \operatorname{Re} \langle \psi, j \rangle.$$

Proof. The functional calculus gives the operator products, semigroup laws, and norm bounds because all spectral multipliers share E_K . Dominated convergence gives strong continuity. The variation-of-constants formula gives the mild solution. On the indicated strong-solution domain, differentiation and self-adjointness make the $h(K)$ contribution purely imaginary in the norm derivative. $\square$

For $h(\lambda) = \lambda$, $W_t = F_K(\gamma t + it)$, so the earlier analytic family is recovered. For general h , both responses remain functions of the same carrier, but U_t is not the imaginary boundary of F_K unless the symbols agree. The symbol, damping, units, and physical realization must be supplied by the nominated source/aperture construction. Shared spectral organization alone does not uniquely select them.

There is a concrete instrument interpretation for a fixed interval whenever the required coupling is admitted. Put

$$M_0 = W_t, \quad M_1 = (I - S_t^2)^{1/2} U_t.$$

Then $M_0^\dagger M_0 + M_1^\dagger M_1 = I$. Retaining an outcome flag gives a trace-preserving two-outcome instrument, with the contraction as the unnormalized 0 branch. This supplies the missing branch rather than renormalizing loss away. It is a mathematical implementation available to a specified coupling, not a claim that every physical vessel exposes that instrument. No assertion of a semigroup for the outcome-discarded channel is needed.

13.5 What a common generator does not determine

A positive generator constrains both rays of Equation (13.2), but a noisy finite observation of one ray need not stably recover the whole operator. High-frequency spectral changes can be heavily suppressed in a heat readout. Exact formal reconstruction and stable empirical inversion are different tasks. The source Hilbertization, admitted probes, and physical calibration remain part of the model.

The analogy with explicit and implicit cognition can motivate architectures in which candidate generation remains coupled to independent checking. Its computational content and its limits are developed in Chapter 22. A cognitive interpretation requires an identified carrier and calibrated observables; the spectral construction alone does not identify logical order with quantum coherence or contemplative openness with dissipation.

13.6 The correct role of vessel mathematics

The preceding chapters provide a connected family of results: source sufficiency, physical retention, operational capacity, regulation, internal witness representation, recurrent spectral response, and observable stability. Their common role is to characterize what a realized system can distinguish, retain, use, and represent. They are relevant to a consciousness theory because any proposed vessel must have some physical constitution. They are not, collectively or individually, a derivation of phenomenal presence.

The next part states exactly how SPC-2 moves beyond this physical core. The additional step is not hidden inside an algebra or a contraction. It is a set of constitutive laws linking a nominated recurrent organization to a localized perspective and its relational contents. This is where the account becomes a psychophysical theory rather than a theory of records with a philosophical metaphor attached.

Part IV

A Finite Psychophysical Constitution

Chapter 14

Why a manifestation law is additional

The preceding parts describe source/readout relations and capabilities of physically realized systems. This part supplies the proposed laws of manifestation. Its first task is to identify which choices the physical results leave open, so that the later constitution makes those choices explicitly. The underdetermination examples below explain why a constitutive law is needed; they do not replace the positive argument for the particular law adopted.

14.1 Common source does not determine a bridge

The physical results establish structures that a vessel can realize. They do not yet determine an experiential assignment. Let $S = \{0, 1\}^2$, with $p(s_1, s_2) = s_1$ and $\phi(s_1, s_2) = s_2$. Both are descriptions of the same source, but no function f satisfies $\phi = f \circ p$. A common-source ontology alone does not imply the desired fibre compatibility.

Adding a constant awareness argument imposes no further mathematical restriction. For every $M : X \rightarrow \Phi$ and fixed a_* , the map $\tilde{M}(a_*, x) = M(x)$ is a possible assignment on $\{a_*\} \times X$. Primitive awareness can be a legitimate ontological commitment while leaving the entire manifestation map unspecified. Its introduction must therefore be followed by actual laws rather than treated as the completion of the theory.

14.2 Regularity and symmetry leave genuine freedom

Fix $X = \Phi = [0, 1]$, its ordinary metric and orientation, and landmarks $0, 1/2, 1$. For $|\varepsilon| < 1$, set

$$M_\varepsilon(x) = x + \varepsilon x(1 - x)(x - 1/2). \quad (14.1)$$

These maps preserve the landmarks and obey $M_\varepsilon(1 - x) = 1 - M_\varepsilon(x)$. Their derivatives are

$$1 + \varepsilon(-3x^2 + 3x - 1/2),$$

which are bounded between $1 - |\varepsilon|/2$ and $1 + |\varepsilon|/2$. Thus they are monotone bi-Lipschitz homeomorphisms. Coordinatewise extension respects Cartesian product composition. Repairs that preserve x preserve every such assignment.

Distinct values of ε nevertheless give different maps. The orientation-preserving isometry group of the fixed interval with its landmarks is trivial. Hence these alternatives are not removed by that declared gauge. The example refutes uniqueness under the specific regularity,

symmetry, and composition conditions just stated. It does not refute every richer naturality law on an independently specified category.

A naturality condition can relocate rather than remove the ambiguity. If both a phenomenal functor and its bridge are unknown, conjugating every phenomenal morphism by an invertible change of representation gives another natural pair. Likewise, replacing a manifestation map and compensating its decoder can preserve all outputs. The unknown target-side structure must be independently constrained or explicitly fixed by a constitutive law.

14.3 Report laws and latent descriptions

A general latent model writes

$$R(r \mid x, k) = \int Q(r \mid \phi, x, k) \mu(d\phi \mid x). \quad (14.2)$$

If both μ and Q are unrestricted, invertible changes of latent coordinates leave R unchanged after the corresponding decoder transformation. More general inequivalent latent spaces can also induce the same report law. Thus a fitted report distribution is not automatically a unique phenomenal law.

An independently specified physical descriptor and experience-facing measurement map define a conditional factorization problem [33]. This formulation permits the target to constrain the descriptor rather than be defined by it. The stronger constitutive proposal below makes a different move: it identifies phenomenal relational organization with a canonical physical predictive object. This excludes alternative organizations *by a declared law*, not by claiming that generic regularity or report agreement has already ruled them out.

14.4 Quantitative aperture adequacy

Let S be finite, $p : S \rightarrow X = p(S)$ a proposed descriptor, and $r : S \rightarrow \mathbb{R}$ an independently nominated target. Define

$$\text{osc}_p(r) = \max_{x \in X} \left(\max_{p(s)=x} r(s) - \min_{p(s)=x} r(s) \right).$$

Theorem 14.1 (Sharp scalar bridge error). *The best deterministic bridge through p has exact uniform error*

$$\inf_{f: X \rightarrow \mathbb{R}} \max_{s \in S} |r(s) - f(p(s))| = \frac{1}{2} \text{osc}_p(r). \quad (14.3)$$

Proof. On a fixed fibre let a, b be the minimum and maximum target values. Every single number approximating both has maximum error at least $(b - a)/2$. Choosing their midpoint attains the bound on that fibre. Make this choice independently on each fibre and take the largest error. $\square$

Exact factorization is the zero-error case. Refining a descriptor subdivides its fibres and cannot increase the optimal error. If $\|r - \hat{r}\|_\infty \leq \epsilon$, the two optimal errors differ by at most ϵ . Consequently a same-descriptor contrast larger than $2\delta + 2\epsilon$ rejects every bridge with uniform modeling error at most δ , provided the equality of descriptors and the measurement bounds are independently justified.

This assesses descriptor adequacy. For a distribution-valued target, the corresponding problem is a minimum enclosing radius in a declared probability metric; half the diameter need not be optimal. The scalar formula therefore does not establish the general metric case or detect awareness.

Chapter 15

The realized domain and its boundaries

15.1 Physical data, constitutive selection, and analysis

A transition matrix alone is not a complete realization. The specification R^* separates three roles. Physical data specify states, preparations, mechanism kernels, ports, actual resources and process provenance. A constitutive selection doctrine fixes which primitive component interfaces, native time cells, internal routes, endogenous records and executable native protocols enter this theory of manifestation. An analyst chooses observations and approximations of that already specified object. The first two roles determine the assignment; the third does not.

Role	What it fixes
Physical realization	The actual kernels, feasible preparations and controls, carriers, incoming interactions, resource updates and provenance.
Constitutive doctrine	The primitive interface and time convention; the internal route mask; endogenous record and test selection; native return contract.
Analyst representation	Display horizon, measured test subset, numerical approximation and reporting format.

The constitutive doctrine is part of the theory's input, fixed independently of a preferred consciousness verdict. It need not be derivable from the bare kernel to be stated consistently, but its adequacy for an actual vessel requires an argument. Merely moving a discretionary choice into the word physical would not remove it. The finite completion theorem is relative to a supplied, checked R^* ; a universal source-side selector of R^* is not claimed.

15.2 A physical interface doctrine before candidate selection

The successor constitution uses an enriched physical realization R^* , specified before any experiential attribution. It supplies a finite primitive component family $\mathcal{B}$, complete state and instrument laws, native physical ticks and operating regimes, preparation laws, admissible physical port operations, route typing, retained controller and resource states, and carrier lineage. The exact construction uses rational data, a finite intrinsic carrier and controller presentation, and a finite complete native return catalogue. The predictive grammar can admit all finite lengths. The candidate projections are generated for every nonempty $C \subseteq \mathcal{B}$ before cores are selected. Thus the physical domain does not presuppose a subject partition.

The primitive factorization, native tick, and route doctrine are declared physical-constitutive inputs. They are not inferred from a transition matrix alone. An analyst may choose a smaller displayed dataset without changing them; physically changing a port or update law is a different operation. A core that lacks an autonomous local kernel is not thereby unconscious: the autonomous-core specialization does not apply until the missing physical state or boundary memory is retained.

Proposition 15.1 (Candidate construction precedes individuation). *Let the complete physical model supply finite preparations, a finite native experiment grammar and resource contract, and rational joint instruments. For every nonempty component subset C , its candidate state projection, projected physical history tree, and joint-closure test can be constructed without first deciding whether C is a core.*

Proof. For each finite depth, enumerate the physical paths and their admitted native controls from the preparation family. Their union defines the possibly infinite candidate history family; the finite instrument representation specifies it without storing every path. For each of the finitely many subsets C , apply its component projection at every node, retaining the specified boundary/control labels and history prefixes. This produces every candidate domain before an SCC or experiential predicate is evaluated. If p_C is the proposed finite carrier and χ_C its native output map, a local joint instrument exists exactly when

$$\sum_{o: \chi_C(o)=y} \sum_{s': p_C(s')=x'} M_{a,o}(s, s')$$

is constant on every nominated fibre of p_C at a fixed boundary contract and context, for each common admissible action a , output y , and successor x' . Finite enumeration checks this condition; fibre constancy defines the local instrument and point-mass complete-state preparations establish necessity for closure on every state in the nominated domain. Equality under only one limited preparation ensemble would be a weaker condition. The physical port contract must also descend. A candidate that fails the check requires an enlarged physical carrier or a stated history treatment. No subject assignment is used in this construction. $\square$

Projected history trees supply physical evidence and possible continuations. They do not turn an observer's posterior distribution into the vessel's actual intrinsic state. In the intrinsic-state specialization the actual point is $p_C(s)$; conditioning incomplete records describes uncertainty about that point. For indefinitely operating finite-state systems, the same finite carrier is supplied with the predictive construction's family of all finite native tests.

15.3 Joint causal dependence

Fix one nominated physical regime. Write $S = \prod_{i \in I} S_i$ and let $T_e(s, \cdot)$ be the complete next-state law under a specified native operation and boundary condition, jointly indexed by e . For $J \subseteq I$, write $T_{e,J}$ for its joint output marginal. Comparisons in the following definition use the same boundary condition and states differing in one primitive coordinate. Their physical scope is part of R^* .

Definition 15.2 (Minimal target dependence). For a fixed comparison (e, s, s') differing only in coordinate i , a nonempty output set J is minimally distinguishing when

$$T_{e,J}(s, \cdot) \neq T_{e,J}(s', \cdot), \quad T_{e,L}(s, \cdot) = T_{e,L}(s', \cdot) \quad \text{for every } L \subsetneq J.$$

Record the directed hyperedge $i \longrightarrow J$ and the associated arrows $i \rightarrow j$ for $j \in J$. Take the union over every admitted native mechanism/action comparison in the one compatible regime,

using the same kernels as the return catalogue. This gives the full dependence graph G^{all} . The internal graph G^{int} retains the arrows licensed by the separately supplied internal-route doctrine.

Write $i \rightarrow j$ only for these direct projected dependence arrows, and $i \rightsquigarrow j$ for a nonempty directed path. Strong connectivity requires mutual reachability; graph recurrence requires a nonempty return path. Neither notation by itself asserts that a physical return schedule has been executed.

Minimality is imposed for each fixed comparison. Including an irrelevant coordinate in a larger affected output set does not by itself create an arrow to that coordinate. Joint distributions remain essential: a source variable can control the correlation of two outputs while leaving both individual marginal distributions unchanged.

Proposition 15.3 (Joint-parent factorization). *Suppose the regime domain is the full Cartesian product and every one-coordinate comparison is included. For $C \subseteq I$, put*

$$P(C) = \{i \notin C : \text{some } j \in C \text{ has } i \rightarrow j \text{ in } G^{\text{all}}\}.$$

Then, for each e , the law $T_{e,C}(s, \cdot)$ depends on s only through $(s_C, s_{P(C)})$.

Proof. If changing only $i \notin C \cup P(C)$ changed the C -marginal, finiteness would provide an inclusion-minimal distinguishing set $J \subseteq C$. This would produce an arrow $i \rightarrow j$ for each $j \in J$, contradicting $i \notin P(C)$. Change the remaining outside coordinates one at a time. Cartesianity keeps every intermediate state in the comparison domain, and the C -law is unchanged at every step. This proves fibre constancy. $\square$

When physical constraints leave holes in the comparison domain, the argument requires the corresponding one-coordinate moves to connect each relevant fibre. Without that condition, joint-parent factorization is checked directly on the nominated states; absence of observed arrows does not establish it. All causal parents are included in this check, including those whose routes are typed external. Conditioning on parent values gives a controlled mechanism, not an autonomous law for an unmodelled external feedback loop.

15.4 Native recurrence and admission

The internal graph provides maximal strongly connected candidates. Graph recurrence is a structural disposition and must not be confused with a presently executed feedback cycle. A finite physical return catalogue is therefore supplied by R^* independently of the experiential verdict. It is generated for every candidate subset from the complete physical experiment model.

More precisely, R^* fixes an executable finite generative rule for the *complete* catalogue allowed by its native operation and resource contract. For example, a nonrenewable physical program counter can bound all admitted recurrence schedules. The rule enumerates every template in that finite contract, including failed tests; the analyst cannot omit a successful template or replace the catalogue by a sample. Its bound is not the reader's analysis horizon. If an indefinitely operating system is instead assigned a finite constitutive template family, that family is an explicit theory-selection premise, not a theorem that it exhausts every physically possible experiment.

Definition 15.4 (Executable internal return witness). A covering return witness for C specifies a component $i \in C$, two admissible preparations differing only by the nominated source intervention at i , a common executable native controller schedule with its resources,

and a terminal internal observation at i after $\ell \geq 1$ native ticks. The schedule carries a physical route certificate for a nonempty closed internal tour beginning and ending at i , whose support is exactly C . For a joint-output mechanism the certificate retains its whole target set rather than substituting an unsupported single-coordinate marginal channel. Every named coupling must be executable in the one supplied schedule with its actual controller and resources. A list of separately possible edges, or a syntactic graph walk without this schedule, is not a covering certificate. The incoming boundary preparation is common and independent of that intervention. External routes that could carry the difference out of C and return it are cut or held fixed by the supplied physical instrument, or their irrelevance is separately proved. The two terminal laws must have strictly positive total-variation distance.

Every item in a finite return catalogue is evaluated by products of the supplied instruments. Controller compatibility and boundary isolation are physical hypotheses of those products. A sequence synthesized from incompatible control modes is not a witness. The catalogue and its native time convention are fixed before subject assignment; changing a reader's display horizon does not add or remove its physical experiments.

The route certificate is finite data, not a new unnamed predicate. It lists the time-indexed native mechanism applications, their internal route or hyperedge labels, controller states, and resource updates. The verifier checks each application against the supplied primitive mechanism table, checks controller/resource admissibility, checks the closed-tour incidences and support exactly C , and then evaluates the two terminal laws by the corresponding rational instruments. For a branching physical schedule these checks cover every positive-probability branch. This verifies execution of the declared covering couplings; positive return is a separate numerical condition.

The following law applies only after the candidate's finite intrinsic representation and its native boundary-policy contract have passed the joint-closure and predictive-congruence requirements. A failure of those requirements leaves this representation outside the specialization; it does not assign unconsciousness to the physical system.

Definition 15.5 (Qualifying core). At a given occurrence, a maximal SCC candidate is qualifying precisely when its certified native representation admits an executable covering return witness starting in the occurrence's actual native operating and protocol type, with its current resource context fixed, and two admissible preparations in that same type have unequal laws for a native finite test. The witness must belong to the complete constitutive return catalogue for that context. A singleton requires a nonempty physical self-return route. The predicate concerns physical organization before any experiential assignment.

This successor law makes structural recurrence, whole-candidate schedule compatibility, and retained influence separate checks. It does not infer the latter two from an SCC calculation. The certificate shows that the covering couplings execute and that a distinction returns; a stronger claim that the returned distinction travels exclusively along every named edge would require route-isolation interventions and is not made. A physical identity or noisy-persistence register can pass a one-tick return test; an unchanged analyst annotation is not a physical register. No minimum intelligence, autobiographical self, or biological material is imposed by this clause.

15.5 Limits of canonical selection

Theorem 15.6 (No equivariant binary-factorization selector). *Let an otherwise unstructured four-state system have identity dynamics and uniform preparation. Consider factorizations into two binary factors, modulo permutations of the two factors and relabelling the values*

within each factor. No deterministic selector of one such factorization from these data alone is equivariant under every automorphism of the data.

Proof. Write the states as 1, 2, 3, 4. The three unordered balanced binary partitions are

$$p_1 = 12 \mid 34, \quad p_2 = 13 \mid 24, \quad p_3 = 14 \mid 23.$$

Any two distinct partitions intersect in four singletons and hence define a binary-by-binary factorization. Conversely each such factorization gives an unordered pair of these partitions. There are therefore exactly three factorizations after the stated factor and value gauges are removed. The full permutation group S_4 preserves the identity kernel and uniform preparation and acts transitively on these three factorizations. No one factorization is fixed by every permutation. An equivariant selector at data fixed by all of S_4 would have to return such a fixed factorization, which is impossible. $\square$

Each candidate has two prime-sized factors, independent uniform preparation, independent identity updates, and a jointly sufficient four-state description. These requirements do not break the symmetry. Returning the entire set of candidates is covariant, but does not select one primitive decomposition. Extra physical ports, geometry, or carrier structure can break the symmetry; their selection is additional data.

Causal-emergence analysis offers one principled way to compare descriptions using specified coarse-grainings and intervention distributions [20]. Such comparisons can reveal advantages of a macro description without deriving a unique primitive factorization from an unstructured kernel. The following consequence concerns precisely that latter, more demanding selection task.

The same obstruction applies to any numerical selection score invariant under the automorphisms of these bare data: the score is constant on the three-element orbit, so optimization cannot pick one member uniquely. Each factorization already has independent subsystem histories under uniform preparation and identity evolution; a conditional-independence requirement that these histories have an empty separating boundary does not distinguish them. Likewise each is exactly sufficient and dynamically stable. Minimality, independence, stability and symmetry therefore reduce no further freedom in this example. This does not refute enriched causal factorizations or multiscale models. It identifies the additional physical structure or selection premise they must supply before a unique choice can be claimed.

Proposition 15.7 (Obstruction to bare-kernel individuation). *A subject-count rule agreeing with primitive-component SCC individuation in the following finite examples cannot simultaneously be determined by the unstructured transition kernel alone, invariant under arbitrary component regrouping, invariant under all integer time resampling, and independent of internal/external route typing.*

Proof. The identity law on four states has two recurrent cores when realized as two binary persistence registers and one when realized as a single four-valued persistence register. The bare stochastic matrix is the same up to a state bijection, but the component doctrine differs. Equip each presentation with its corresponding native state records and executable one-tick self-return contract; both then meet the admission conditions.

For two binary components the one-tick SWAP law $(A, B) \mapsto (B, A)$ has one SCC, whereas its two-tick kernel is the identity and has two self-loop SCCs. Thus integer resampling changes the partition when the sampled update is incorrectly substituted for the native mechanism. The one-tick presentation has a two-tick covering return, and the sampled presentation has one-tick self-returns; both can be supplied with qualifying native records.

Finally, the fixed kernel $T = (I + \text{SWAP})/2$ has self-dependence and cross-dependence. Typing the cross-routes as internal gives one SCC; typing them as external retains two internal self-loop candidates. The complete transition kernel is unchanged. These differences cannot be removed by an invariance assertion while preserving all three stipulated SCC assignments. $\square$

The proposition does not rule out a deeper physical selection law. It shows why the source/readout premise and a bare kernel are insufficient to supply one. The successor states its interface doctrine rather than calling that missing selection canonical.

Proposition 15.8 (Covariance of finite individuation). *A bijection of primitive components and coordinatewise state labels that transports the complete kernels, admitted comparisons, route typing, regimes, preparation laws, native controller grammar, return catalogue, and lineage transports the minimal target hyperedges, parent-closure conditions, return witnesses, and admitted SCC partition.*

Proof. Corresponding joint output marginals have equal transported probabilities. Equality, inequality, and inclusion-minimality of output subsets are preserved, so hyperedges and their typed directed graphs correspond. Fibre constancy and the SCC partition are invariant under these bijections. Corresponding finite experiments have the same transcript probabilities and total-variation distances, preserving the return predicate and the admission law. $\square$

This covariance does not cover arbitrary invertible recodings of a joint state, regrouping primitive components, altering physical ports, resampling the native dynamics, or changing a causal threshold. Those operations need their own physical equivalence theorem.

Chapter 16

Endogenous predictive structure across horizons

Predictive and causal-state representations are established antecedents [48, 26]. The present intrinsic-state quotient adds an explicit transition-congruence requirement; its phenomenal identification is a separate law.

16.1 An intrinsic carrier and a native operational contract

The predictive construction starts with a physical state, not an observer's posterior. For a nominated core C , let X_C be its finite intrinsic carrier. A point $x \in X_C$ is the actual realized core state, including any internally retained memory needed by its update. The native operational contract specifies which internal operations and records belong to the vessel's organization. It is fixed before the phenomenal assignment and is not the set of measurements an analyst happens to perform.

Definition 16.1 (Closed native instrument representation). A closed native representation consists of a finite carrier X_C , finite native action and record alphabets A_C, O_C , a physically specified continuation grammar, and nonnegative instrument matrices $K_{a,o}(x, y)$ such that

$$\sum_{o \in O_C} \sum_{y \in X_C} K_{a,o}(x, y) = 1$$

for every admitted action a and state in its common protocol type. The entry is the joint probability of native record o and actual next core state y . The representation is required to be exact for the nominated operating regime and boundary-policy contract.

A finite controller can describe the grammar. An endogenous controller whose retained state affects later native operation belongs to the physical carrier. An externally maintained condition belongs to the boundary contract; its future influence must be specified rather than silently discarded. A returning external memory can invalidate the claimed closure on X_C . The remedy is a justified complete model or a narrower operating contract, not an assertion that the return route is irrelevant because it is typed external.

Admissible adaptive policies select only operations of this native grammar, using the native records and controllers allowed by the contract. Defining a counterfactual policy does not supply a missing actuator, permit arbitrary clamping of a boundary value, or introduce an ideal non-disturbing observation of the entire core state. An executable operation that can fail has its physical failure record; a physically inadmissible operation is not added as a test.

Policy randomization is common to compared states conditional on the permitted retained record. A controller's further dependence on its actual internal state is represented in the physical instruments, not supplied as state-dependent knowledge to the testing policy.

Protocol type must be fixed by the initial native type and retained action/record history, or supplied by an actual native type record. An adaptive controller receives no hidden-state oracle telling it which operation is available. The statements below compare states of the same protocol type. Distinct types are retained as distinct components of the structured domain. The grammar has stopping and admitted continuation, and can be compiled into finitely many controller types. These conditions make the native test family fixed by the supplied physical contract. They do not derive a unique native contract for every organism.

A native contract states whether it represents continued operation in a physically maintainable regime or follows a scheduled physical change. In the former case it describes counterfactual dispositions under that regime, not an unconditional forecast through a later intervention that changes it. In the latter case the finite controller and its regime changes are included. Physical exhaustion or exit from the carrier's nominated operating domain has an explicit terminal record and absorbing continuation. A later core occurrence is assigned from its own supplied contract and linked, where appropriate, by A3. Changing the contract is a physical or constitutive change; displaying a shorter part of it is not.

16.2 All finite futures, with no privileged truncation

For an admissible native policy π of finite depth, let P_x^π be its future native transcript law from the actual state x . Present-state copying by an external device is not automatically one of these native tests. For each integer $r \geq 0$, put

$$x \sim_r x' \iff P_x^\pi = P_{x'}^\pi \text{ for every native test of depth at most } r, \quad Z_r = X_C / \sim_r.$$

Protocol typing is included also at grade zero. Because every shorter test is a permitted truncation of the same physical grammar, there are canonical surjections

$$p_{r+1,r} : Z_{r+1} \longrightarrow Z_r, \quad p_{r+1,r}([x]_{r+1}) = [x]_r.$$

The constitutive state identification uses the whole compatible family:

$$x \sim_\infty x' \iff x \sim_r x' \text{ for every finite } r, \quad Z_\infty = X_C / \sim_\infty.$$

Thus an analyst's choice to inspect only a horizon H can merge distinctions in an approximation, but does not change the constitutive identification. The actual point is $q_\infty(x)$, not a conditional belief selected by an outside observer.

Proposition 16.2 (Finite-carrier projective realization). *For finite X_C , the map*

$$Z_\infty \longrightarrow \varprojlim_r Z_r, \quad [x]_\infty \longmapsto ([x]_r)_{r \geq 0}$$

is a bijection. No unrealized states are added by this inverse limit.

Proof. Equality at every grade is exactly $\sim_\infty$, proving injectivity. A compatible family of classes gives nonempty nested subsets of the finite carrier X_C . Their intersection is nonempty; any member realizes the family. All members of the intersection are $\sim_\infty$ -equivalent, proving surjectivity and uniqueness. $\square$

This finite-carrier result must not be transferred without hypotheses to an unlimited record-history domain. A finite hidden transducer can generate infinitely many posterior predictive states. In that different construction, projective completion can introduce points not realized by any nominated history. The present law uses the finite intrinsic carrier and does not replace its actual states by filtered beliefs.

Theorem 16.3 (Coarsest native predictive representation). *The map $q_\infty : X_C \rightarrow Z_\infty$ is sufficient for all admitted finite native transcript laws. Within each protocol type, every other descriptor sufficient for all those laws factors uniquely through to q_∞ on its realized image.*

Proof. Each test law is constant on a predictive class by definition. If two intrinsic states have the same sufficient descriptor, all their test laws agree, so their predictive classes agree. Sending a realized descriptor value to that common class is well defined and unique. $\square$

16.3 The transition closure that prediction alone does not prove

Equality of all future record laws need not make the actual-state quotient a Markov instrument. We therefore require and check the following additional closure condition.

Assumption 16.4 (Native predictive-fibre congruence). For $x \sim_\infty x'$, every commonly admitted native action a , record o , and class $D \in Z_\infty$ satisfy

$$\sum_{y \in D} K_{a,o}(x, y) = \sum_{y \in D} K_{a,o}(x', y).$$

This assumption concerns actual successor classes, not merely output probabilities. It is a finite property of the supplied native representation. Failure means that this proposed predictive quotient does not close as the claimed phenomenal transition structure; it is not a finding that the physical system is unconscious. Automatically refining the quotient until the property holds would be a different constitutive law, since it can retain distinctions invisible to all native transcript tests.

Theorem 16.5 (Intrinsic quotient instruments). *Under native predictive-fibre congruence,*

$$\overline{K}_{a,o}(z, D) = \sum_{y \in D} K_{a,o}(x, y), \quad q_\infty(x) = z,$$

is independent of the representative. These matrices form a normalized native instrument on Z_∞ and reproduce every finite adaptive native transcript law. For a positive-probability record o , its conditional actual-successor law is

$$\Pr\{Z_{t+1} = D \mid Z_t = z, a, o\} = \frac{\overline{K}_{a,o}(z, D)}{\sum_E \overline{K}_{a,o}(z, E)}.$$

The next realized phenomenal point is $q_\infty(X_{t+1})$.

Proof. The assumption gives representative independence. Nonnegativity and normalization follow by summing the physical instrument over the partition. For a first action and record, the quotient records exactly the probabilities of every next class. Induction on a finite policy tree proves equality of all transcript probabilities. Dividing the joint class-and-record law by its positive record probability proves the conditional formula. No normalized conditional successor is assigned to a zero-probability record. $\square$

The conditional distribution describes uncertainty about the actual next point. It is not itself that point. There is generally no deterministic update $z \mapsto z'$ from the record alone. A theorem about Bayesian filtering on histories cannot replace the congruence assumption for intrinsic actual states.

For clarity, a five-state example shows why the assumption is needed. With one action, let a emit 0 and remain at a ; let b emit 1 and remain at b ; let c emit 0 and enter a with probability $1/2$, or emit 1 and enter b with probability $1/2$. Let x emit $\#$ and enter a or b with equal probabilities, whereas y emits $\#$ and enters c surely. States x, y have identical laws for every finite transcript: $\#$ followed by all zeros or all ones, with probabilities $1/2$. The predictive classes are $\{a\}, \{b\}, \{c\}, \{x, y\}$. Nevertheless the probability of entering $\{c\}$ after $\#$ is zero from x and one from y . Output equivalence has therefore not supplied actual-class closure.

16.4 An evaluable all-future object

For a common unrestricted native action alphabet, define

$$V_0 = \text{span}\{\mathbf{1}\}, \quad V_{n+1} = \text{span}(V_n \cup \{K_{a,o}v : a \in A_C, o \in O_C, v \in V_n\}).$$

The sequence stabilizes in a subspace of $\mathbb{R}^{X_C}$. Its limit is precisely the span of all finite native word-likelihood vectors $K_{a_1,o_1} \cdots K_{a_m,o_m} \mathbf{1}$.

Proposition 16.6 (Exact finite test of all-future equivalence). *For rational native instruments, $x \sim_\infty x'$ if and only if $(e_x - e_{x'})v = 0$ on a rational basis of the stabilized space. The partition Z_∞ , its fibre-congruence check, and, when that check succeeds, its quotient instruments are exactly computable. The same conclusion holds for a finite-controller typed grammar, using one observable subspace per controller type.*

Proof. The closure contains exactly the span of finite word-likelihood vectors by induction. A proper enlargement increases dimension, so finite-dimensional stabilization gives a rational basis in finitely many exact linear-algebra steps. Equality on all words is equivalent to equality of every fixed finite action transcript. Every leaf of an adaptive policy has one such action-record word; equality of the word probabilities therefore gives equality for adaptive policies as well. Conversely, in an unrestricted alphabet fixed action words are native tests. In a typed grammar, each admitted action/record word extends to a valid policy that stops on other branches, so its cylinder probability is a probability of a native test. For a finite typed controller, start with a terminal vector for each type and propagate only through its admitted action-record successors. The sum of the subspace dimensions is finite. Finally evaluate the congruence equalities and quotient sums directly. $\square$

The complete endogenous object is

$$\mathfrak{P}_C = (Z_\infty, \{P_z^\pi\}_{\pi \text{ finite native}}, \{\overline{K}_{a,o}\}_{a,o, \text{ native types and labels}}),$$

pointed at $q_\infty(x)$. The finite quotient instruments generate all its finite transcript laws. The projective family records the consistent finite-horizon views of the same object. Actual transition structure is retained explicitly rather than inferred from a scalar contrast metric.

One may define

$$d_\infty(z, z') = \sup_{r \geq 1} \sup_{\pi: \text{depth } \pi \leq r} \text{TV}(P_z^\pi, P_{z'}^\pi).$$

Within each common protocol type this is a metric on its Z_∞ fibre and the monotone supremum of the finite-horizon pseudometrics. The finite span algorithm decides zero distance; it is not

a general algorithm for the exact value of the unbounded-horizon supremum. The completion claim requires the finite instrument presentation, not an unevaluated promise of exact metric optimization.

16.5 External reports and exact descent

An external report protocol k has a physical law $Q_k(r \mid x, b)$ on a supplied product domain $X_C \times B_k$, where b includes its preparation, channel state, boundary conditions, and other external nuisance variables needed to specify that law. The report is grounded in this core precisely when, at fixed b ,

$$q_\infty(x) = q_\infty(x') \implies Q_k(\cdot \mid x, b) = Q_k(\cdot \mid x', b).$$

For finite carriers this is necessary and sufficient for a unique kernel on the realized quotient,

$$Q_k(r \mid x, b) = \hat{Q}_k(r \mid q_\infty(x), b), \quad \tilde{Q}_k(r \mid \phi, b) = \hat{Q}_k(r \mid \Psi_C^{-1}(\phi), b).$$

Sufficiency follows by choosing a representative of each class; necessity follows immediately from factorization. This criterion is checkable when the physical report law is supplied. If only a subset X_b of intrinsic states is physically possible at a context b , existence and uniqueness are asserted only on $q_\infty(X_b)$; no report probability is assigned by descent to an impossible pair.

A report that fails the criterion remains a physical observation, but it is not a function of this nominated phenomenal state and the declared external channel context. Conversely, different phenomenal classes need not be distinguishable by the available report hardware. For example, a constant-output channel distinguishes none of them. For grounded channels at a fixed context, a difference in report laws therefore implies different phenomenal classes. The converse fails unless a separately stated report-completeness assumption makes the report family separate all classes. Arbitrary ungrounded physical observations have no such guaranteed ordering. Grounding also does not imply truthfulness, completeness, absence of confabulation, or an empirical test of primitive presence.

Proposition 16.7 (External-apparatus invariance). *Adding, removing, or redescribing an external report device preserves the native phenomenal object and its actual point whenever it leaves the intrinsic carrier, actual core state, native instruments, physical boundary contract, and native grammar unchanged up to the specified structure-preserving equivalence. Grounded report laws may change.*

Proof. The invariant data determine every native transcript law, hence every finite quotient and the all-future partition. They also determine the descended native instruments and the actual point. External report kernels are absent from this construction. $\square$

This is an invariance under preservation of the physical organization, not a claim that arbitrary measuring devices are harmless. Feedback, backaction, resource depletion, altered boundary inputs, or a change to a native controller can change the object because they change its physical premises.

16.6 The exact one-step specialization

If a genuinely closed native model records its entire next state without additional disturbance at every step, and the compared states have a common native action type, then the all-future

profile is already determined by the native first-step laws. After the action and first realized state are retained, later transcripts are common Markov processing of that pair. Consequently the optimal finite- or all-future contrast equals the maximum TV distance between the admitted first-step laws. Equal first-step rows also guarantee predictive-fibre congruence.

The specialization justifies the all-future interpretation of the worked two-bit matrix when its exact next-state records are part of the nominated native contract. It is not a theorem that every vessel has ideal access to all of its own microstate. With partial native records, distinctions can first appear after an arbitrarily long delay, and the complete projective construction is needed.

16.7 Why successor-state geometry is a specialization

A one-step rule can collapse physically different present states. For example, let $T(0, \cdot) = T(1, \cdot) = \delta_2$ and $T(2, \cdot) = \delta_0$. Exact future-state observation identifies 0 with 1, whereas an exact present-state copy separates them with distance one. The collapse is consistent with a future-dispositional phenomenal law, but it is a substantive identification, not a proof that the present distinction never exists physically.

The general A2 law therefore does not require ideal observation of every successor microstate or arbitrary boundary clamping. It retains the actual native instrument grammar. In a partially observed finite process a distinction can first become visible after several updates. A shift register whose hidden end feeds a recorded port only after m ticks gives a simple delay of any chosen finite length as its carrier grows. An analyst stopping at $m - 1$ sees one class where the all-future object has two. The compatible family retains that distinction without introducing an observer-chosen prediction horizon.

A controller selecting actions from past observations need not induce an uncontrolled Markov process on the current quotient class. The quotient instruments are Markov conditional on the chosen action and retained native controller type. If a policy has additional memory affecting its action, that memory must be included in the controlled process or kept explicitly in the conditioning. The quotient theorem never deletes it.

16.8 Stability of finite predictive contrasts

Exact predictive classes and graph boundaries can change under a small perturbation. Quantitative finite-horizon contrasts nevertheless admit a useful uniform bound.

Proposition 16.8 (Finite-horizon perturbation bound). *Consider two finite instrument models on the same physical state space and common typed grammar. For every state and admitted action, suppose the joint law of recorded outcome and successor state differs by at most $\epsilon \in [0, 1]$ in total variation. From the same fixed preparation b , every adaptive protocol of depth at most r satisfies*

$$d_{\text{TV}}(P_b^\pi, \hat{P}_b^\pi) \leq 1 - (1 - \epsilon)^r \leq r\epsilon.$$

For two fixed preparations b, b' , their contrast pseudometrics satisfy

$$|d_r(b, b') - \hat{d}_r(b, b')| \leq 2[1 - (1 - \epsilon)^r] \leq 2r\epsilon.$$

Proof. Couple the initial states identically. While states and transcripts agree, the common controller chooses the same action; couple the two joint output/successor laws maximally. Each such step preserves agreement with conditional probability at least $1 - \epsilon$. The probability

of a discrepancy by depth r is at most $1 - (1 - \epsilon)^r$, which bounds transcript total variation. Apply the triangle inequality to the laws from b and b' for each common protocol, and then take the supremum, to obtain the contrast bound. $\square$

If starting preparations differ by δ in total variation, add δ to the transcript bound. This statement compares fixed physical preparations. Filtering the two models after the same rare observed history can amplify discrepancies and requires a separate lower bound on its probability. Small contrast errors also do not preserve exact equality classes or the nonzero-edge criterion in A1.

Chapter 17

The Shadow Psychophysical Constitution

17.1 Awareness and explicit laws of manifestation

The constitution is defined relative to the certified realization R^* of Chapter 15. Its physical inventory and its declared selection doctrine have been fixed before a subject is assigned. An analyst may investigate only a subset of that inventory, but cannot alter the constitution by performing fewer tests or stopping a calculation earlier.

Constitutive law 17.1 (A0: source awareness). Awareness is the knowing aspect of U, reality prior to the operational source/readout distinction. Its capacity is not produced ex nihilo by a vessel. It is not a localized subject or a second independently acting substance or force.

A0 uses “knowing aspect” ontologically, without attributing a personality, scene, memory or intention to U. Localized experiential presence is its proposed manifestation through a qualifying vessel. A0 supplies the ontology of the assignment rather than an extra numerical input to its construction. The graph, predictive object, and lineage can be computed without deciding whether their realization has an experiential aspect. A0 interprets the admitted perspectives as manifestations of intrinsic source awareness; A1–A3 specify their local organization. The completion theorem establishes determinacy conditional on those commitments, not a derivation of A0. An ontology can give a construction this meaning while leaving its physical probabilities unchanged.

The primitive is not a formally described global episode lacking all contents. Numerical categories such as one and many apply to the admitted local perspectives, not to U. Its ontological capacity is not a continuing hidden experience between those episodes. The present theory assigns localized episodes and does not supply an exhaustive state space for awareness apart from every vessel. It preserves the source-awareness commitment without identifying the source with one numerically specified universal experiencer. The word field names a technical physical structure only where its carrier, degrees of freedom, transformations, dynamics, and coupling are supplied; A0 alone is not a measurement of an additional field.

17.2 A1: where perspectives are assigned

Constitutive law 17.2 (A1: native recurrent realization). A candidate internal SCC carries one localized experiential perspective when it meets Definition 15.5: an executable covering return and at least two distinct native predictive classes in the occurrence’s actual native operating/protocol type and resource context. Distinct admitted SCCs carry distinct

perspectives. Maximality is an explicit dispositional individuation rule; no additional subjects are assigned to overlapping recurrent subassemblies within one admitted core.

The predicate of Definition 15.5 is physically evaluated before this attribution. The positive motivation is retained differentiation participating in an ongoing organization. The covering schedule prevents a cycle assembled only from incompatible control contexts; maximality provides nonoverlapping support. Neither fact proves experiential unity. A1 makes that further identification openly.

A persistent physical register is not an unchanged annotation in an analyst's description. Genuine self-dependence means an admissible intervention changes a later internal law. It can satisfy the exact return test, including when the process is used as an archive. Thus the law has no general theorem excluding passive memories. A one-time imprint followed by a constant erasing update fails; a physically self-dependent record with an available native return can pass. The labels active and passive cannot replace these tests. The law is deliberately permissive toward minimal organization and imposes no intelligence or biological threshold.

Formal admission and its empirical lower bound. A1 is the present formal admission law of SPC-2. Accordingly, a certified realization satisfying its conditions receives a localized perspective within the constitution. This is not empirical validation that A1 identifies the actual lower boundary of consciousness in nature. Nor does a toy recurrent controller by itself establish an adequate R^* : its native records, executable return, boundary isolation, timing, resources and operating contract require physical justification (Sections 15.1 and 15.4). An annotation has no causal return; an externally imposed replay independent of the intervened core state has identical terminal laws. An omitted environmental feedback path cannot certify a return internal to the nominated core. These are failures of the proposed realization, not general proofs of unconsciousness.

Stricter application of the existing certificate does not exclude every simple system. A genuinely realized one-bit self-return process with distinguishable native future laws can meet the certificate, including predictive-fibre congruence and a supplied resource/provenance contract. Its admission is then a consequence of A1, not necessarily an artifact of oversimplification. The concern combines realization adequacy in a particular application with an unresolved empirical and philosophical lower-bound question; it does not, by itself, expose an internal inconsistency. Excluding every such system would change A1 or its constitutive selection doctrine. If independently constrained non-toy realization work shows systematic over-admission, A1 will require refinement.

Formal A1 admission and a rich biological realization of scene-forming coupling must therefore be distinguished. The mammalian findings in Section 21.4 concern the latter and do not modify A1. A proposed core cannot both satisfy A1 and be declared genuinely without experience by an additional cortical veto: that would be a conflict in the application requiring resolution. Loss of a familiar mammalian scene alone does not decide the status of every remaining subsystem.

For a mature biological vessel, this mathematical admission rule can be approached through the idea of a world-for-the-organism. Internal states have significance relative to bodily regulation, viability, action and learned relationships; conditions can go better or worse for that organism. Calling such organization a minimal center of concern gives the biological interpretation a humanly intelligible focus. It does not mean reflective worry, a narrative ego, or a separately present observer.

SPC-2 does not yet contain a formal general criterion equating such concern with subject admission. Recurrence and the exact return tests remain the computable A1 machinery. Candidate refinements involving integrated endogenous organization, a world-for-the-vessel,

minimal concern or stronger native predictive unity remain biologically and philosophically motivated possibilities rather than hypotheses of the current completion theorem. No stronger admission criterion follows here from those descriptions alone. The current law retains its permissive minimal cases without an autopoiesis requirement or arbitrary complexity threshold.

SPC-2 does not imply that all matter is conscious. A rock is not admitted merely because it exists; tissue activity and simple interaction alone do not establish a perspective. A1 concerns a certified organization. Source-aspect ontology consequently need not assign localized experience to every physical object. Equally, the statement does not covertly exclude a minimal physical memory that actually satisfies A1. Any stronger exclusion would be another constitutive commitment.

17.3 A2: phenomenal organization

Constitutive law 17.3 (A2: endogenous predictive structuralism). The phenomenal relational organization of each qualifying core is a structural copy of its complete endogenous predictive object over all physically admitted finite continuations. The actual phenomenal point is the image of the predictive class of its realized intrinsic state. The copy preserves native operation and outcome labels, every finite test law, the compatible horizon restrictions, and the well-defined class transition instruments. Apart from primitive presence, no further manifested qualitative fact is postulated within this constitution.

Work on mathematical structures of experience motivates specifying the transported relations explicitly [22]; it does not select this particular identification. A2 is the structural law of the lived scene within the constitution. It identifies an organization with phenomenal character; it does not generate awareness, prove its existence, or derive that identification from a sufficient statistic. Its motivation is that an internal distinction acquires a role through the continuations it supports. The complete profile retains which operations distinguish, the distributions of their internal outcomes, and their continuing organization. A supremum distance alone loses these relations and can become discrete even in very different processes.

The future laws characterize current dispositions. They are not future events acting backwards, and they need not be calculated or represented by the vessel. Physical recurrence, a mathematical predictive description, and an internally implemented self-model are distinct. Neither language nor autobiographical reflection is required by this law.

Within SPC-2, the complete endogenous predictive structure exhausts phenomenal organization; the theory does not claim to have proved that no conceivable richer theory could add further qualitative facts. Hidden qualitative differences that leave that complete structure unchanged are excluded within this candidate law. This is a strong internal commitment, not a universal metaphysical no-go theorem.

A structural copy proves mathematical consistency of the assignment. It does not prove that every retained endogenous distinction is experienced. A critic can argue that unconscious processing makes the proposed profile too rich, or that a further intrinsic qualitative remainder survives after every structural relation is fixed. A2 makes an explicit choice on both questions. Its application must face independently elicited phenomenal relations; those relations cannot be redefined after a mismatch merely to preserve an isomorphism.

The gauge is a whole-model change of presentation preserving the specified structure and its calibration. It does not identify different points inside a fixed calibrated model. If an outcome label is renamed, its physical interface, probability laws, and report interpretation must be transported with it. Nor may one change the metric of an independently nominated phenomenal target simply to make it match.

The all-finite family removes dependence on an investigator's chosen forecast cutoff. It does not make the constitutional grammar optional. A new internal operation, a changed physical boundary contract, or a changed endogenous readout may change the organization. Attaching an external report apparatus that preserves these data does not. The distinction is formalized in Section 16.5 rather than left to intuition about what counts as internal.

17.4 Process provenance and temporal identity

The identity of an experiential episode, the persistence of a person, and the ontological status of awareness answer different questions:

persistence of awareness as aspect $\neq$ persistence of a subject $\neq$ persistence of a person.
--

The end of a localized subject is the end of that subject under the stated identity law. Persistence of U is not its survival or migration elsewhere. Memory copying and uploading do not, by themselves, transport a subject; a fresh qualifying receiver has a new admission, with any genuine branch handled by A3.

A physical state is not a permanent material label. Cells turn over, hardware is repaired, components enter and leave an active organization, and its operating mode changes. A temporal law must describe those events without equating memory equality with identity or requiring every atom to remain in place.

The realization supplies native physical time cells indexed by t , component sets $\mathcal{B}_t$, and a directed *process-provenance relation*

$$\Lambda_t \subseteq \mathcal{B}_t \times \mathcal{B}_{t+1}.$$

An edge records continued implementation of a component process or an explicitly realized local transfer of its operative role. This relation is supplied with the physical event history: retained carriers, births, retirements, and renewal operations have declared antecedents. A mere statistical match or information-copy edge is not a provenance edge. An information-copy edge belongs to a separate information-transfer relation. Both may exist in one event, but neither is inferred from the other.

This is an explicit realization premise, not a solution of the philosophical identity problem by renaming it. For a fixed-carrier circuit, Λ_t is the ordinary persistence of active registers; newly allocated copies have new lineages. A physically documented in-place renewal can transport an ongoing process to new material. A proposed application with two incompatible process-provenance assignments has not yet supplied one R^* . The present theory does not choose between them from matching transition tables alone.

Let $\mathcal{C}_t$ be the qualifying cores already computed at cell t . Form the bipartite core-provenance graph with relation

$$C \mathsf{L}_t D \iff \exists i \in C \exists j \in D : (i, j) \in \Lambda_t, \quad C \in \mathcal{C}_t, D \in \mathcal{C}_{t+1}. \quad (17.1)$$

Only qualifying predecessors and successors enter this graph. Addition of a previously nonqualifying support component is not, merely by that addition, a merger of subjects. The relation is evaluated before any numerical identity is assigned.

Constitutive law 17.4 (A3: nonbranching process continuation). A core C at t continues as D at $t + 1$ precisely when $C \mathsf{L}_t D$, C has exactly one qualifying successor in $\mathcal{C}_{t+1}$, and D has exactly one qualifying predecessor. Subject episodes are maximal chains of these continuation

edges. At branching, merging, or a physical loss of qualification the predecessor episodes end; qualifying successors without a continuation edge begin new episodes. An analyst's truncation censors an episode and does not terminate it.

This law deliberately admits change of support. Nonbranching growth, shrinkage, and renewal need not carry every old component into the new core. At least one realized process-provenance connection across each native cell is required. No numerical fraction of retained matter or memory is inserted. The cost is explicit: a very small surviving process connection can sustain identity if all other conditions hold. A stronger requirement that a specified backbone, proportion, or organizational invariant survive would be another constitutive law, not a consequence of numerical identity alone.

Proposition 17.5 (Episodes and gradual replacement). *For any supplied Λ_t and qualifying core sets, A3 defines an unambiguous partition of qualifying core occurrences into episodes. If a sequence $C_t, \dots, C_{t+n}$ has a provenance edge at each adjacent pair, where each adjacent source has exactly one qualifying successor and each adjacent target exactly one qualifying predecessor, all its occurrences lie in the same episode even when $C_t \cap C_{t+n}$ is empty. Restricting an already constructed episode to an observation window preserves identity at common times.*

Proof. The retained continuation graph has indegree and outdegree at most one and increases time strictly along every edge. Its connected components are therefore directed paths, not branching trees or cycles. Path membership is an equivalence relation and gives a unique partition. Every adjacent pair in the stated sequence is a retained continuation edge, so transitivity gives the conclusion without any endpoint material-overlap condition. Restricting the paths changes only which part is observed. Branch degrees must be evaluated in the full native provenance graph before restriction; deleting an unseen competitor and then recomputing degrees would be a different operation. $\square$

Physical case	Consequence of A3, conditional on the realized core and provenance data
Growth or loss of auxiliary support	Continuation when the core remains qualifying and its core-provenance edge is nonbranching.
Gradual cell or hardware replacement	Continuation through the chain; endpoint material identity is unnecessary.
Local repair or memory loss	Does not itself end an episode. Qualification and provenance, not equality of memory, decide.
Temporary disconnection	If the core partitions into qualifying successors, the split ends its episode. Reconnection starts a merger episode. If qualification is lost, later return begins a new episode.
Sleep or anesthesia	Mode changes alone do not decide. A qualifying core can change content; a genuine qualification gap ends the episode, and renewed qualification begins a new one.
Separate copying	Information can be copied while the donor's process persists. The fresh receiver is not the donor's continuation merely because its state agrees.
Symmetric physical fission	Multiple qualifying process successors end the predecessor episode and begin separate successor episodes.
Physical fusion	Multiple qualifying process predecessors end; the combined successor starts a new episode.
Instantaneous replacement	Requires an explicit process-renewal provenance edge. Functional equivalence alone supplies none.

Physical provenance Λ_t can describe a persisting organism or vessel during an interval without experience. The subject graph in A3 contains only qualifying occurrences: a genuine qualification gap contains no currently instantiated localized experiential subject and ends the

episode. Return of qualification begins a new episode, even when biological and person-level continuity is preserved. Lineage is not hidden experience. Sleep and anesthesia therefore require separate assessments of vessel continuity, scene occurrence and later memory (Section 2.5).

The treatment of split and merger preserves ordinary numerical identity. If one predecessor were identical to each of two distinct successors, transitivity would identify the successors. Ancestry and informational continuity may branch; episode identity here does not. The law chooses the symmetric termination convention rather than privileging one successor. It concerns experiential episodes, not every legal, biological, or narrative use of the word person.

Native-time refinement needs particular care. Inserting additional cells is harmless when each old continuation edge is subdivided into a nonbranching chain and every old branch or qualification gap remains visible. Under that condition contraction of the inserted vertices recovers the same episode relation. Sampling only before and after an unobserved split, merger, or loss of qualification need not do so. Temporal resolution is therefore part of R^* , and the observational window is not.

17.5 Reports, phenomenal identification, and access

Reports are produced by physical channels with their own state, resources, timing, noise, and later decoding. They do not enter A2 merely because an investigator records them. A grounded report channel factors through the core's predictive class at each fixed boundary/channel context, as established in Chapter 16. Its phenomenal expression is transported by A2's identification. No independent psychophysical decoder remains adjustable.

Grounding is a sufficiency requirement, not a completeness or truthfulness requirement. Two different phenomenal states can yield the same constant report, and apparatus noise can erase distinctions. A family separates all phenomenal states only under the additional report-completeness premise. Conversely, a physical instrument that separates two states in one predictive class is not a grounded report of that class under the declared law. It remains an ordinary physical observation with a determined probability; it cannot be promoted to a phenomenal decoder without changing the law or the realization.

A0, A1, A2, and A3 perform different jobs: presence, admission and individuation, relational organization, and episode identity. The physical realization supplies event and report probabilities. This division makes the ontological interpretation explicit while leaving no report mechanism to be chosen after the desired conclusion is known.

17.6 The constitutive choice

The designation SPC-2 marks substantive commitments. Relative to SPC-1, it fixes a separate endogenous grammar, uses one compatible all-finite predictive structure instead of a nominated finite-horizon assignment, requires predictive transition congruence on its finite intrinsic domain, detects minimal joint-output dependence in its physical influence structure, and admits nonbranching process continuation through support turnover. External report maps are independently grounded rather than included automatically in the phenomenal test family. These choices preserve the source-aspect architecture while defining a more explicit manifestation law. They are not deductions that every admissible consciousness theory must share.

Chapter 18

Completion, invariance, and empirical conservativity

18.1 The certified finite domain

A *certified finite realization* R^* has finite physical component and state domains, rational physical instrument kernels, an explicit native time and route doctrine, executable finite recurrence witnesses, physical process provenance, and finitely represented native operating contracts for every candidate core. Each candidate's intrinsic carrier and common typed continuation grammar are generated before qualification. Its native instrument representation is closed, and its all-future predictive equivalence satisfies the checked congruence condition of ?? 16.4. Nominated experience-facing report channels have rational physical kernels and satisfy the grounding criterion at every supplied channel context.

These are mathematical and physical input conditions. They do not assert that a candidate is conscious. A failure of closure, admissibility, or grounding makes the proposed realization uncertified for this theorem; it does not force a negative consciousness verdict. This distinguishes an incomplete model from a completed model that assigns no perspective. General hidden or history-dependent vessels may require a larger representation or another theorem.

Theorem 18.1 (Constitutive completion of a certified finite realization). *Given a certified finite R^* and A0–A3, the following assignments are determined on every supplied native time cell and along its supplied provenance history:*

- (i) *every physical component belongs to a unique maximal internally typed joint-influence SCC; qualification is a total predicate, and the qualifying components form a disjoint subject partition of their union;*
- (ii) *core occurrences have a unique episode partition under the continuation law, with specified split, merger, and loss-of-qualification boundaries;*
- (iii) *every admitted occurrence has a complete endogenous phenomenal structure and its actual point, unique up to the declared whole-structure gauge;*
- (iv) *every nominated grounded report probability is induced by its physical channel, with no independently adjustable phenomenal decoder;*
- (v) *the construction is covariant under the declared realization equivalences and independent of analyst-selected test subsets, forecast truncations, and observational windows that do not change R^* ;*
- (vi) *all finite presentations needed for these assignments, and all individual finite transcript probabilities, are exactly evaluable from the rational data. The claim does not include general exact optimization of an infinite-horizon contrast supremum.*

No further manifestation map is left unspecified after the realization, selection doctrine, and

named constitutive laws have been fixed.

Proof. The finite physical comparisons determine the minimal joint-output dependence relation and its internally typed projection. Strongly connected components are a unique graph partition. The candidate construction produces the intrinsic carrier, grammar, and witness data from physical components without assuming qualification. Finite witness evaluation and the all-future equivalence calculation determine the qualification predicate; A1 supplies its experiential attribution.

For each candidate, finite rational observable-span closure decides equality of all native finite transcript laws. The nested finite quotients realize their inverse limit by Proposition 16.2. Congruence is checked by finite sums over the resulting classes. The quotient-instrument theorem then supplies actual-class transitions and a finite presentation of every continuation law. A2 copies this complete structure and sends the realized intrinsic state to its quotient point. Two such copies are isomorphic through their identifications with the same physical object. This is the elementary structural-copy part of the proof; the well-defined physical construction is the substantive prerequisite.

The full core-provenance graph and the degree-one rule give the episode paths by Proposition 17.5. Channel grounding gives a unique report kernel on each realized predictive class at fixed external context. Transporting it through the A2 identification supplies the phenomenal expression with no new choice. The next theorem proves covariance. Restricting which data an analyst inspects leaves all these physical and constitutive data unchanged. Every operation on the finite presentation uses exact rational tests, finite graph algorithms, or finite linear algebra. Individual finite protocol probabilities are finite products and sums of the supplied instrument entries. $\square$

The subject partition in (i) covers admitted core support, not all matter: components outside qualifying cores are not assigned additional subjects. A0 has no numerical role in these computations. It gives the resulting A1 assignments their proposed source-aspect meaning. The theorem establishes conditional determinacy, not the truth of the psychophysical identifications or a derivation of awareness from nonexperiential premises.

The same realization must support the whole continuation grammar. A collection of unrelated finite-resource models, one selected anew for each horizon, does not satisfy that condition. A finite bank may instead be included explicitly in the state and lead to an absorbing exhausted protocol type. That type is assessed on its own current return capabilities, not on a live type that is no longer physically available. Its all-finite transcript family remains well defined, while its physical life is finite. No infinite resource is obtained from the inverse-limit notation.

18.2 Redescription and operational presentation

A realization isomorphism transports physical states, component incidence and provenance, native cells, preparation and intervention domains, minimal joint-dependence witnesses, internal/external typing, candidate construction, native instruments and grammar, actual points, recurrence certificates, boundary contracts, and report channels. It is a bijection preserving this whole structure, not merely a permutation of the rows of one transition matrix.

Theorem 18.2 (Realization covariance). *Such an isomorphism carries qualifying cores, their all-finite endogenous objects, episode partitions, and grounded report laws to their corresponding objects. It preserves the assigned structure and point up to the declared gauge.*

Proof. A transported intervention comparison has the same joint probabilities on corresponding output subsets, so minimal distinguishing subsets and their typed projected edges correspond. Graph components and physically executable witnesses therefore correspond. Native protocol trees transport with the same laws; equality at every finite horizon, and hence all-future equivalence, is preserved. Sums over corresponding predictive fibres preserve congruence and intertwine the quotient instruments. The provenance graph and its incoming and outgoing degrees are preserved, so the episode paths correspond. Report descent commutes with the same transport. A1–A3 consequently give the corresponding assignments. $\square$

A modest extension concerns redundant *descriptions* of tests. Suppose two finite presentations implement exactly the same labeled native instruments, and their executable policies have mutually translating presentations with the same transcript laws, stopping rules, resource use, and admissible continuation. Identifying these presentations yields the same predictive equivalence and instrument object. Likewise duplicating a calibrated output label as a bijective encoding changes its presentation, not its content. These statements require actual equivalence of the specified native structure. Adding an unavailable control, inserting a new record, or changing physical timing is not a redundant presentation.

Transformation	Scope of invariance
Coordinate or calibrated label change	Covariant when every affected physical and interpretive datum is transported.
Source-preserving realization isomorphism	Covariant by the theorem.
Redundant grammar presentation	Invariant when the same executable native instruments and continuations are represented.
Analyst test subset or forecast cutoff	Changes available evidence or an approximation, not the constitutional assignment.
State refinement or coarse-graining	Not generally invariant. Requires a separately verified equivalence preserving the joint graph, native quotient instruments, qualification and provenance.
Component regrouping or port retyping	Can change subjects; generally a different realization doctrine.
Native-time resampling	Can hide cycles, splits or gaps; generally changes the realization. Pure subdivisions preserve episodes only under the condition in Section 17.4.
Additional physical intervention grammar	Can change endogenous organization when genuinely implemented; analyst wish or apparatus readout alone does not add it.

18.3 Independent composition

For two closed native systems with independent preparations, separate resources, product instruments, no internal cross-dependence, and a product grammar containing their separate tests, all-future equivalence of pairs is the product of the separate equivalences. Equal pairs of classes give equal laws by induction over any admitted common policy tree; separate tests distinguish unequal pairs. If each local quotient is congruent, summing the product instruments over pairs of classes proves product congruence. On a restricted jointly reachable domain only the corresponding realized subset is present.

A common adaptive controller can correlate transcript outputs by choosing one system's later input from the other's earlier record. The transcript distribution need not factor under such a policy. This does not invalidate the conditional product-instrument statement. If that controller becomes part of the internal organization or changes its interfaces, the graph and

qualification must be recomputed. A mathematical product by itself creates no third subject beyond the qualifying cores.

18.4 Physical conservativity

Theorem 18.3 (Conservativity of the aspect interpretation). *Fix the complete physical realization, preparation, and all physical operation and report kernels. Adding A0–A3 and grounded phenomenal interpretations without changing those kernels leaves the probability of every finite adaptive physical transcript unchanged.*

Proof. The initial law is the same. At a policy node the next action depends on the same recorded history and the conditional physical instrument is unchanged. Induction over the policy tree gives identical finite joint laws. The phenomenal variables add an interpretation of the realized organization, not an independent argument of the physical transition kernel. $\square$

Grounding removes a free internal decoder, but it creates no physical prediction absent from the supplied realization. A rival accepting exactly the same complete physical transcript laws while denying their experiential interpretation cannot be excluded by those transcripts alone. This underdetermination includes primitive source awareness and any alternative that differs solely by transcript-inert qualitative claims.

Independent experience-facing relations can nevertheless constrain an application. A fixed proposed physical realization can fail to predict the observed records. A1’s proposed boundary can conflict with an independently defended subject-individuation judgment. A2 can fail to match independently elicited phenomenal contrasts under a fixed calibration and held-out interventions. Such evidence bears on the realization and the psychophysical identification together; it does not turn a definitional copy into an independent observation of awareness.

18.5 An aspect and an independently variable force

For comparison, let a complete physical model have state (x, a) and kernel T_u . A reduced autonomous law on x for every initial preparation exists exactly when

$$\sum_{a'} T_u((x, a), (x', a'))$$

is independent of a for every x, x', u . Point-mass preparations prove necessity; conditioning proves sufficiency. This is ordinary controlled lumpability. If the condition fails, deleting a leaves an incomplete physical description, regardless of the name assigned to it.

An aspect theory instead identifies experience with an aspect of the same physical process. Independently varying that aspect while fixing its complete realization is not an admitted intervention. The process remains physically causal; the theory does not add a second force. Conservativity establishes equality of the stipulated probabilities and does not, by itself, settle the ontology or every philosophical question about mental causation.

Chapter 19

A fully worked finite vessel

19.1 Physical dynamics and realization

Let the complete local state be $x = (A, B) \in \{0, 1\}^2$. In a fixed operating regime define

$$A_{t+1} = B_t \oplus N_A, \quad B_{t+1} = A_t \oplus N_B, \quad (19.1)$$

with fresh independent noise bits satisfying $P(N_A = 1) = 1/10$ and $P(N_B = 1) = 1/4$. These probabilities specify the native stochastic law. For its all-finite operational interpretation the law is assumed to continue with fresh independent noise; this is a physical model premise. A fixed finite run can instead be realized using an explicitly sized noise bank with its retained outgoing cells. The finite bank does not establish an indefinitely replenished reversible implementation.

The step has a reversible extension:

$$(a, b, n_A, n_B) \mapsto (b \oplus n_A, a \oplus n_B, n_A, n_B).$$

Its inverse is obtained from $b = a' \oplus n_A$, $a = b' \oplus n_B$. Thus a finite sequence can be realized as a permutation on the enlarged orthogonal state space, with independent noise cells initially prepared and all old cells retained. This is a finite physical model construction, not a claim to reproduce a biological vessel.

Changing A_t with the other relevant parents fixed changes the next law of B ; changing B_t changes the next law of A . Both routes are internal under R^* . Take the native action to be this update and the native record to be the next values of the two registers, which are themselves reused by the following update. The two-tick schedule executes the route $A \rightarrow B \rightarrow A$; its terminal A law differs by $2/5$ between preparations differing only in A . The complete constitutive return catalogue consists of the native one- and two-tick return templates and their admitted component preparations. This finite template choice is explicit; it is not selected anew by an investigator. Four distinct first-step rows give four all-future predictive classes, and their singleton fibres make congruence automatic. Thus A1 assigns one perspective in this operational model.

19.2 The exact transition matrix

Order configurations as 00, 01, 10, 11. Direct multiplication of the independent noise probabilities gives

$$T = \frac{1}{40} \begin{pmatrix} 27 & 9 & 3 & 1 \\ 3 & 1 & 27 & 9 \\ 9 & 27 & 1 & 3 \\ 1 & 3 & 9 & 27 \end{pmatrix}. \quad (19.2)$$

Every row sums to one. The matrix is doubly stochastic, so the uniform distribution is stationary. Stationarity of the ensemble does not mean that the actual bits do not change. It also does not collapse the four predictive preparations into one class: their future conditional laws are distinct.

A separate report can copy A or B into a prepared orthogonal register. Here the report device has no return interaction, and its classical copying preserves the native state and update. It belongs to the external report layer. The native and external interrogation grammars consequently have different contrast metrics without representing a change of the core itself.

19.3 Delayed observations and the native contrast geometry

First consider only delayed complete-state observations beginning at $t + 1$, with no current-state report and no prior intervention that amplifies the distinction. The one-step TV matrix between the rows of Equation (19.2) is

$$D_{\text{del}} = \begin{pmatrix} 0 & 4/5 & 1/2 & 4/5 \\ 4/5 & 0 & 4/5 & 1/2 \\ 1/2 & 4/5 & 0 & 4/5 \\ 4/5 & 1/2 & 4/5 & 0 \end{pmatrix}. \quad (19.3)$$

For example, the first two rows differ by $(24, 8, -24, -8)/40$, for which half the sum of absolute values is $4/5$. The first and third differ by $(18, -18, 2, -2)/40$, giving $1/2$.

The same matrix is the TV distance of any complete future passive path that retains its first state X_{t+1} . Given that first state, the remaining conditional path law is the same Markov kernel from either preparation. Applying that common kernel cannot increase TV; projecting the path back to its first state cannot decrease the original row distance. Equality follows. More time does not reveal an additional distinction about the initial state once the common Markov successor has been exactly observed.

For this declared native full-successor-record grammar, D_{del} is the all-finite native contrast metric. It is not universal phenomenal geometry: another realized native grammar need not give this matrix. Exact current-state reports belong to the separate external grammar in this example.

19.4 External present-state reports have a different metric

Suppose the separate external report mechanism copies the present A_t or B_t before any noisy transition. Any two distinct configurations differ in at least one bit. Reporting that bit produces disjoint deterministic output laws. Therefore

$$d_{\text{all}}(x, x') = \mathbf{1}_{\{x \neq x'\}}. \quad (19.4)$$

Equation 19.4 is the supremum metric of external interrogation, not the native metric assigned by A2. All four native classes are already distinct, so exact present-state reports factor

through them without refining the phenomenal partition. They can nevertheless have a larger TV contrast: report grounding is fibre constancy, not a Lipschitz or data-processing bound relative to the chosen native transcript metric. Adding an actual internal copy operation would change the native grammar and require a new realization analysis.

Under A2, the four pointed phenomenal states are the structurally transported four predictive states. They are not labeled as red, pain, or any human quality. The model fixes relational organization under its constitutive interpretation; it supplies no independent phenomenological calibration that would identify these states with an actual person's experience.

19.5 Reports, symmetry, and gauge

For a current A report, $Q_A(r \mid a, b) = \mathbf{1}_{\{r=a\}}$, and similarly for B . A delayed A report has $Q_{A,+1}(1 \mid a, b) = 1/10 + (4/5)b$; a delayed B report has $Q_{B,+1}(1 \mid a, b) = 1/4 + (1/2)a$. These are ordinary physical predictions. Their phenomenal versions are induced through Ψ_C^{-1} .

The asymmetric noise rates distinguish the two update channels. A permutation of state labels alone is not a physical symmetry if it fails to transport the channel names and probabilities. Some simultaneous bit-complement transformations can be presentation equivalences when the corresponding input and report labels are also transported. Such whole-model equivalences do not mean that a current 00 preparation is the same physical preparation as 11 under a fixed calibration.

The operational example specifies one core, four native predictive states, congruent native instruments, separate grounded report kernels, a native delayed contrast matrix, a distinct external interrogation metric, and a pointed assignment under A2. A maintained two-tick return contract qualifies it. If a fixed bank is included instead, the remaining fuel is part of the carrier and exhaustion is terminal; the four-state comparison is then a fixed-resource slice, not the entire enlarged object. Its mathematical completeness does not establish actual awareness in this device. The latter is the interpretation asserted by A0 and A1, not an additional finite calculation.

Chapter 20

Splitting, merging, and the stability of subject boundaries

20.1 Two cores and a controlled merger

Consider two bits initially evolving by separate noisy persistence channels, with retention contrasts $a = 4/5$ and $b = 1/2$. Their internal dependence graph has two self-loops and no cross edge. Each one-tick physical return has positive TV, and the native records are the actual next register values. Both candidates qualify, so A1 assigns two perspectives.

For $g \in (0, 1]$, install one native stochastic mechanism with the joint update

$$T_g = (1 - g)T_0 + gT_1,$$

where T_0 is the independent persistence update and T_1 is the cross-update of Equation (19.1). The mechanism, with both of its internal input ports, is executed at every tick. Its native route certificate records two applications of that installed mechanism and the tour $A \rightarrow B \rightarrow A$. The convex mixture is a representation of its stochastic law, not a schedule that physically omits the mechanism on some branches. A different implementation that actually switches between incompatible topologies would require its own certificate.

The one-step cross-dependence is nonzero for every $g > 0$. Using centered binary coordinates, the two-step mean response of A to its own earlier value has coefficient

$$a^2(1 - g)^2 + abg^2 = \frac{16}{25}(1 - g)^2 + \frac{2}{5}g^2 > 0.$$

The shared operating contract includes that two-tick return template. The route certificate verifies execution of the two installed coupling applications on every outcome branch; it does not assert successful transmission of a distinction along every noisy sample path. Exact native next-state records supply nontrivial prediction and automatic predictive-fibre congruence. There is therefore one qualifying combined core under these stated mechanism and route assumptions. A3 ends the two predecessor episodes at the merger and begins a successor. Removing the cross mechanism restores two cores and starts their successor episodes.

If communication is instead through ports typed as external, the graph used for A1 omits the external return route and retains the two internal cores. The physical communication is not denied; it is classified as interaction between the nominated vessels. The different result makes boundary typing load-bearing. A universal theory would need to justify that typing from deeper physical structure rather than select it according to an intuitive subject count.

20.2 Arbitrarily weak coupling gives a discontinuity

The exact nonzero-edge rule is not generally robust. Let $T_g = (1 - g)T_0 + gT_1$, with T_0 the independent update and T_1 a bidirectionally coupled update. Then

$$\sup_s d_{\text{TV}}(T_g(s, \cdot), T_0(s, \cdot)) \leq g.$$

For the same initial preparation and H steps, an auxiliary mixture coupling gives

$$d_{\text{TV}}(\text{Law } X_{0:H}^{(g)}, \text{Law } X_{0:H}^{(0)}) \leq 1 - (1 - g)^H \leq Hg. \quad (20.1)$$

With those compatible return certificates, the subject count changes from two at $g = 0$ to one for every positive g . Thus arbitrarily small finite-horizon physical changes can produce a discontinuous subject assignment.

This is not a formal contradiction. A theory can contain a sharp threshold. It is, however, a substantive prediction and a robustness problem. An ϵ -effective graph that ignores effects below a fixed discrimination scale gives a more stable operational classification, but it is an added resolution convention, not the exact law secretly preserved. Inference about actual subject boundaries should report this sensitivity rather than hide it behind the word “integration.”

If every estimated edge strength is separated from a declared threshold by a margin larger than the error bound, the thresholded graph and its SCC partition are stable. This follows because no edge can cross the threshold under the allowed perturbation. Near zero in the exact rule there need be no such margin. Physical model uncertainty and psychophysical discontinuity must therefore be analyzed together.

20.3 Copying, reset, and continuation

Let a donor and receiver be distinct lineaged carriers. A reversible copy transfers a selected record, not the donor’s physical lineage. If each belongs to a separate qualifying core, A3 does not identify their episodes merely because their memory contents agree. The same applies to an artificial-system handoff that reproduces a relational task state in another instance.

A local SWAP reset can preserve the core’s physical lineage while changing its working memory. Whether qualification persists depends on the actual post-reset operation. A1 may continue to admit the core; A2 generally changes the pointed content if future internal laws change. Resetting a narrative label does not automatically reset a subject, and preserving a narrative label does not automatically preserve subject identity through a genuine physical replacement.

A narrative variable can belong to the core’s retained organization, to an external record, or to a report channel. Its causal location matters. Changing a self-description that participates in endogenous continuation can change the assigned phenomenal structure; changing only an external label need not. The philosophical statement that ego is not awareness therefore leaves room for profound experiential consequences of the personal model without treating every statement about identity as a change of subject. A core can continue while its self-description changes, and a similar self-description can occur in a distinct successor episode.

20.4 Finite partitions and universal unity

Communication by itself does not establish subject unity. Two machines, two brains or a loose network are not one subject merely because they exchange signals. A brain and body,

or a distributed process, can receive a unified assignment only if the relevant whole is one justified realization with the internally typed organization and return tests required by A1. A tiny exchange outside that organization does not automatically merge subjects.

This qualification does not remove the exact weak-coupling consequence above. Once a positive coupling belongs to the certified internal organization and supports the required joint return, an arbitrarily weak exchange can merge the cores under the present law. Boundary typing must be justified independently; it cannot be chosen after the fact to preserve a preferred count.

The construction determines a subject partition only relative to the nominated realization. It does not infer a universal observer or a numerical One from U. Nor does it infer independent subjects for every mathematical subsystem. The distinction between a physical component, an operational assembly, a qualifying core, a phenomenal perspective, and a narrative person must be retained.

This is where the philosophical attraction and cost of the proposal meet. A0 permits a common ontological ground without demanding separate mental substances. A1 supplies a local individuation rule. A2 determines relational content. A3 chooses a temporal identity doctrine. The combination is coherent and evaluable in a finite model, but each additional identification creates an empirical or philosophical responsibility. The theory's seriousness depends on accepting those responsibilities rather than claiming that graph theory has answered them without assumptions.

Part V

Interpretation, Comparison, and Empirical Responsibility

Chapter 21

Biological and comparative applications

21.1 What an application must establish

The finite constitution gives a rule on a realized operational domain. Applying it to an organism is not accomplished by substituting the word “brain” for a vertex of a graph. An application must identify the components, the time scale, their physically admitted interventions, the internal and external interfaces, the relevant preparations, and the information available to the process itself. It must then determine whether the operational reduction is adequate for the proposed experience-facing targets.

Three questions should remain separate. Does a chosen physical model predict the experimental records? Does its proposed core decomposition remain stable under justified changes of resolution? Does the independently elicited phenomenal organization agree with the structure that A2 assigns? Success on the first question is indispensable but is not automatically success on the other two. Failure on the first cannot be repaired by invoking primitive awareness.

It is therefore useful to distinguish a *candidate application* from a *validated realization*. A candidate application specifies an organism or artificial process and proposes an R^* model. A validated realization would require evidence that the supplied model and its interfaces capture the relevant physical organization, together with a justified interpretation of the experience-facing measurements. No such completed validation for a human, animal, plant, or deployed language model is asserted here.

21.2 Development and the absence of adult prerequisites

The developmental case tests the temptation to identify consciousness with the mature capacities of its investigators. The review by Bayne and colleagues argues for taking early infant experience seriously while examining the uncertainty about its onset and form [2]. This is not a license to assign an exact developmental threshold from a spectral gap or a recursion count. Its relevance is that adult language, explicit philosophical reflection, and an elaborate autobiographical self cannot simply be built into an allegedly universal admission rule.

Within SPC-2, an immature vessel can qualify through an endogenous recurrent core without implementing adult counterfactual deliberation. Learning may change the physical kernels, the repertoire of available actions, the predictive quotient, and the content of an episode. A0 does not have to increase as a child learns a word; A2 permits the organization of a perspective to change. This captures the intended distinction between an underlying

awareness commitment and acquired scaffolding without pretending to have measured a quantity of primitive awareness.

The model also distinguishes a repertoire from its current point. A system can have a rich capacity for discrimination while occupying a simple current state; conversely, an apparently vivid report does not establish a large or unified predictive structure. A developmental interpretation should analyze both the available organization and the actual state within it. The words “less developed” refer here to specified functional capacities, not lesser moral worth or a smaller entitlement to protection.

21.3 Sleep, anesthesia, and the lived scene

Clinical descriptions often bring three questions together: does the vessel continue, is there experience now, and can the person subsequently describe it? Separating them does not make assessment simple, but prevents a failure of memory or motor output from deciding a different question by default. The distinction in Section 2.5 applies throughout this discussion.

21.3.1 Dream formation and reduced external constraint

In this framework, dreaming is interpreted as a scene generated when external constraint is reduced while the vessel’s generative and scene-forming organization remains sufficiently coupled. The experienced setting can be rich even when it is weakly constrained by the current environment. Reduced constraint does not mean complete sensory isolation, and internally generated does not mean detached from physical realization.

This account is an interpretive model rather than a claim that all dreams arise through one mechanism. Siclari and colleagues compared reports of dreaming and their absence across REM and non-REM sleep [50]. Their work supports distinguishing sleep stage, reported experience and recall; it does not turn the proposed coupling language into an established SPC-2 mechanism. The theory neither identifies all sleep with dreaming nor treats dream recall as the criterion of dream occurrence.

21.3.2 Dreamless and unremembered intervals

An unremembered dream and a genuinely scene-less interval are different possibilities. In the former, experience occurred but was not retained or recovered for a later report. In the latter there was no current experience. It would be misleading to describe that absence as a subject observing darkness, or to preserve experience by stipulating that everyone always dreams and merely forgets.

On the biological interpretation, coupling sufficient to form a scene can fail while the organism and much of its regulation continue. During a genuinely scene-less interval there is no currently instantiated localized experiential subject. This is compatible with the organism’s survival and ordinary biographical continuity. Within SPC-2, a genuine qualification gap ends the experiential episode under A3; when qualification returns, a new episode begins. An unremembered dream does not establish such a gap.

21.3.3 General anesthesia

General anesthesia can disrupt the organization required for an integrated lived scene while substantial biological and local neural activity remains. Neuroscientific accounts distinguish changes in large-scale interaction from a simple cessation of brain activity [29]. The effects of particular drugs, doses and clinical contexts need not be identical. Unresponsiveness,

disconnected experience, amnesia and absence of experience are also distinct possibilities; retrospective reports concerning experimentally induced unresponsive states under propofol and dexmedetomidine illustrate why they must be distinguished [45]. Those findings do not establish how often experience occurs under deep surgical anesthesia, and recollections may depend on the transition toward arousal.

Within the Shadow interpretive model, sufficient disruption of coupling between explicit/attended and unconscious/generative organization prevents a lived scene; restored coupling permits manifestation to resume. This does not assert that “one side is switched off” as a universal neurophysiological fact. A biological person can remain physically continuous, with background processes active, throughout a genuinely scene-less interval. Awareness, as the knowing aspect of U, is not said to be created or destroyed by the anesthetic. During that interval there is no current lived scene and no currently instantiated localized experiential subject. There is no observer experiencing darkness: darkness would itself be a scene.

A genuine qualification gap ends the experiential episode under A3. When qualification returns, a new episode begins, although the biological/person-level vessel has remained physically continuous. For a genuinely scene-less interval, the first-person sequence can be represented schematically as

$$\text{scene}_1 \longrightarrow \emptyset \longrightarrow \text{scene}_2.$$

The middle term denotes no experience, not an experienced empty interval. This makes intelligible why deep general anesthesia can feel, on return, as though no time elapsed: there is no intervening scene presenting that duration. Physical time still passes. An impression of immediacy can also follow failed recall, so it does not by itself establish that experience was absent.

A3 can continue an episode across a change of anesthetic regime when qualification and nonbranching provenance persist, including a change to disconnected experience. Neither the clinical label nor a report of remembering nothing establishes loss of qualification. The coupling account interprets possible biological realizations of A1–A3; it adds no separate scene gate and supplies no clinical test for a particular patient.

21.3.4 Local anesthesia and partial sensory loss

Local anesthesia provides a different case: a peripheral intervention can block or alter signals from a restricted region while the person’s larger lived scene continues. Additional sedation or other clinical factors must be considered separately. The relevant contrast is between a changed incoming channel and a disruption of the organization sustaining the whole scene. Local sensory silence is not, by itself, destruction of a subject.

In the framework’s terms, a sensory boundary changes what can enter or constrain experience. Other sensory, bodily, mnemonic and affective organization may remain. This illustrates why a missing modality, a failed report route and a loss of experiential admission cannot be treated as the same event.

21.4 Candidate cortical correlates of scene-forming coupling

In mature mammalian cortex, large-scale recurrent interaction and effective connectivity are strong candidates for the physical realization of scene-forming coupling; selected fast-band coordination may contribute to that organization. This is a biological interpretation of organized interaction among distributed processes, not an identification of consciousness with tissue activity. Effective connectivity concerns how activity in one population influences

another; observational estimates of directional dependence and responses to an imposed perturbation provide different, method-dependent evidence for it.

The anesthesia literature supports the distinction. Human intracranial recordings during propofol induction showed that local neuronal relationships could persist while activity became restricted to brief windows occurring asynchronously across cortical regions [25]. Cortex need not be globally silent for communication to fragment. Lee and colleagues found reduced estimated frontal-to-parietal feedback around loss of responsiveness under ketamine, propofol and sevoflurane, despite different spectral effects [24]. Their directional statistic and behavioral endpoint do not establish absence of every kind of experience. In rats, Imas and colleagues found preferential impairment of frontal–posterior feedback at 50 Hz under volatile anesthetics, while effects at 30 Hz differed [21]. These findings motivate studying recurrent organization and its disruption, rather than assigning one meaning to all gamma activity.

Agent-specific perturbational findings sharpen the point. Sarasso and colleagues used transcranial magnetic stimulation with EEG: propofol produced relatively local responses, xenon produced widespread but stereotyped responses, and both had low perturbational complexity. Ketamine preserved more differentiated responses and high complexity, alongside subsequent vivid dream reports [44]. Broad propagation alone therefore does not establish differentiated conscious integration, and reduced feedback in a particular estimate cannot be a universal criterion for scene absence. Neither retrospective dream reports nor their absence escape the memory limitations discussed above.

Hameroff emphasized gamma-range synchrony and anesthesia-related disruption of conscious integration [19]. That emphasis is relevant here; his proposed dendritic mechanism, microtubule computation and Orch-OR are not adopted. Gamma power, phase coordination and effective interaction are distinct measurements. No frequency band is awareness itself, either of the conscious/generative modes, or an admission condition. REM dreaming, ketamine-related disconnected experience and seizure states must be assessed without treating fast activity, synchronization or responsiveness as sufficient for conscious integration. There is no “gamma on = conscious” rule.

The relevant invariant is organizational integration versus its failure, not a single number on a monitor. Here integration means differentiated, effective interaction capable of supporting the proposed scene organization, not maximal synchrony or merely extensive connections. This is the interpretive target across implementations, not an additional mathematical invariant already proved for cortical measurements. No universal EEG threshold, necessary cortical band or unique anatomical locus is established.

These network findings motivate candidate biological realizations of scene-forming coupling in a rich mammalian R^* . They are not a new admission law, not A1, and not a proof of A0 or A2. During a genuinely scene-less anesthetic interval, no localized experiential subject is currently instantiated under SPC-2, while the biological vessel remains physically continuous and the person biographically continuous. Awareness as an aspect of U, formal admission, the biological coupling account, its measured correlates and empirical validation remain separate levels of claim. The combined SPC-2 interpretation places dreams and anesthesia within a source-aspect constitution; it claims no novelty for recurrent processing or network integration individually. Existing anesthesia and predictive-processing accounts support or motivate parts of this pattern without establishing the ontology or the constitution [24, 44, 9].

21.5 Disorders of consciousness and incomplete access

Bodien and colleagues reported task-related fMRI or EEG responses in 60 of 241 participants who lacked observable responses to commands in their examined sample [6]. This finding

constrains the use of overt motor behavior as a sole measure of relevant cognitive activity. It does not show that every neural response measures all experience, nor that every negative neural test demonstrates absence of experience.

In an R^* model, overt report is one output channel, not the definition of the process. A motor pathway can fail while another internal discrimination or response route remains. The predictive profile can therefore retain distinctions that a particular report channel loses. The aperture-adequacy theorem gives a formal way to express the issue: a descriptor that collapses conditions differing in an independently justified target cannot support an exact bridge for that target.

This is a methodological implication, not a diagnostic instrument. No clinical decision should be based on the two-bit example or an unvalidated assignment of perspectival cores. An application in patients would require appropriate clinical expertise, consent or surrogate procedures, validated measurements, and uncertainty reporting. The framework's insistence on separating presence, access, memory, and output is intended to prevent premature classification, not to replace established care.

21.6 Blindsight and divided processing

The classical blindsight literature describes residual visual discrimination in a field defect without the corresponding ordinary report of seeing [43]. Its role here is to make content-specific discrimination, confidence, access, and phenomenal attribution distinct questions. The condition of an entire person cannot be inferred from a single missing visual report.

A2 is deliberately demanding in this setting. If two conditions share the complete nominated predictive structure yet differ in an independently established phenomenal relation, A2 fails for that realization or the physical descriptor is inadequate. It is not enough to point to a successful forced-choice response and call that response experience. A2 is a psychophysical identification requiring justification, not a linguistic decision to redefine all discrimination as phenomenology.

Split-brain research is particularly relevant to subject individuation. Pinto and colleagues argued for divided perception without two independent perceivers in the cases they studied [30]. Volz and colleagues dispute that inference, emphasizing cross-cueing, ipsilateral motor control and possible subcortical transfer [54]. Whatever position one takes on that interpretation, graph partition and phenomenal subject count cannot simply be read off from the phrase "split brain." The physically operative routes, remaining common control, temporal grain, and independently assessed behavior all matter.

The exact SCC rule makes a definite conditional assignment once those data are supplied. Its sharp weak-coupling transition is a vulnerability rather than a fact established by the clinical literature. A biologically adequate realization might preserve important internal routes after a particular anatomical intervention; an overly coarse model might erase them. Conversely, retaining every negligible feedback route might over-unify the system. This makes the physical boundary and resolution problem central to testing A1.

21.7 Animals and nonhuman presentation

A conscious subject need not possess human-style language, autobiography, explicit metacognition, or a reflective self-concept. Animals can in principle be conscious under this theory when an appropriate R^* realizes A1; this is not a universal species list or a declaration that every proposed animal model is certified. Comparative accounts can describe several dimensions of

an animal's experiential capacities rather than place all species on one scale [5]. This does not identify those dimensions with the present predictive object. The dolphin, whale, and insect comparisons should begin with a modest logical point: different sensory and action repertoires can support different ways of organizing a world. There is no requirement that every potential perspective contain adult human visual categories, spoken language, or a human bodily model. A nonhuman comparison must use its own physically meaningful input and action spaces rather than treating human reportability as the universal standard.

Under SPC-2, two species could have different predictive structures even when a particular task score agrees. Conversely, a particular operational structure might be reproduced across different substrates without establishing that their entire phenomenal lives are identical. Cross-system phenomenal comparison is meaningful only after specifying which physical structures and protocols the comparison preserves. A2 supplies a conditional structural comparison, not unrestricted imaginative access to another organism's life.

No numerical consciousness threshold for insects or marine mammals follows from the O1 spectra. Those finite complexes are source-response examples, not measured nervous systems. Their value is methodological: provenance and sector geometry can matter even when spectra agree. The lesson transfers as a requirement to preserve relevant structure, not as a numerical species classifier.

21.8 Plants, fungi, and minimal systems

Plant consciousness remains contested in the literature. Taiz and colleagues argue that plants neither possess nor require consciousness; Segundo-Ortín and Calvo explore a more permissive analysis of plant cognition and consciousness [52, 46]. These positions should not be collapsed into agreement that signaling proves experience. Communication, adaptation, and distributed regulation establish physical capacities; the psychophysical conclusion remains a further claim.

The present constitution is potentially permissive in a different way: a simple recurrent device can qualify under A1. The requirement of nontrivial predictive structure does not impose a high intelligence threshold. That implication should not be concealed because it is counterintuitive. It is one of the costs and possible tests of the proposed law. An application to a plant, fungal system, or controller still requires the same R^* specification rather than an inference from the material category alone.

A theory that intends to exclude minimal recurrent systems needs an additional principle. It could introduce a physically motivated scale, an integration condition stronger than SCC membership, or a different admission rule. None follows simply from preferring a familiar list of conscious organisms. Such an extension would be a new version of the psychophysical constitution and should be compared explicitly with A1, not presented as an unnoticed consequence of it.

Chapter 22

Artificial systems, embodiment, and recursive cognition

22.1 The relevant implementation is active

Artificial consciousness is possible in principle within the substrate-open constitution, but it is not demonstrated here. Current LLM chat sessions are not treated as certified vessels. A real application requires a justified R^* , including persistence, native instruments and endogenous organization; fluent surface language supplies none of those certificates by itself. Minimal concern may motivate a later formal refinement, but it is not an additional admission condition in this edition.

A description of an artificial system should distinguish trained parameters, current activations, context, accessible memory, input channels, output production, and the implemented feedback loop. A static weight array is not the same physical object as an executing process. Nor is a text transcript by itself the complete causal state of an agent. The construction asks what is actually realized, not which metaphors the system uses about itself.

Theory-based indicator approaches provide one route to assessing artificial systems. Butlin and colleagues examine properties drawn from scientific theories of consciousness rather than treating persuasive self-report as sufficient evidence [8]. A contrasting biological-naturalist position emphasizes that reproducing an abstract computation may not reproduce all consciousness-relevant properties of living systems [47]. These are substantive competing approaches; neither is adopted here as an independently established verdict about every artificial implementation.

Within A1–A2, the invariance claim is conditional on preserving the selected causal and predictive organization. It requires more than similarity of outputs. It neither privileges a material category by itself nor proves that every biological property can be omitted. If a computational simulation preserves only a coarse input/output relation, the covariance theorem does not establish phenomenal equivalence. If an implementation genuinely transports the full nominated R^* structure, A2 does assign the corresponding structural equivalence. Whether R^* omits a biologically relevant property is precisely a question for the realization theory.

22.2 Embodiment changes the available world

The hypothetical addition of biological senses and motor interfaces to an artificial cognitive system is useful because it exposes a hidden assumption in substrate comparisons. Adding eyes, interoception, locomotion, or a body model changes a system’s interactions and predictive

structure. A theory should identify which of those changes is alleged to alter experiential admission, rather than asserting that biological attachment itself is an unexplained switch.

Under the source-aspect interpretation, a richer interface can change manifestation without producing primitive presence from nothing. But that interpretation does not independently establish that the unembodied predecessor already qualified. A1 supplies the conditional test, and its physical premises must be evaluated on both implementations. The correct conclusion is a comparison of two realized organizations, not a declaration that nothing relevant can change during embodiment.

Embodiment may be important for the organization being investigated, but it is not declared a sacred substrate or a strictly necessary condition for every possible perspective. A body may also change the internal/external typing. A formerly external sensor or controller can become part of a tightly coupled implemented loop. R^* must specify whether this is a genuine change in the nominated vessel or merely a change in its external inputs. Such decisions must be grounded in physical routes and retained dependencies, not in whether a preferred consciousness verdict would result.

22.3 Memory transfer and the apparent continuity of a model

The author's historical symbolic handoff experiments motivate a concrete question about preserving relational task state across computational sessions. A compact representation may restore facts, associations, commitments, problem-solving context, and a characteristic interaction style. A subsequent system may then generate a statement of personal continuity. That observation, if reproduced under controlled conditions, would be evidence about reconstruction and reporting. It would not by itself prove transfer of a numerically identical experiential subject.

A3 makes the distinction explicit. Copying informational content is not the same as transporting the physical lineage of a subject episode. A restored process can be structurally similar while beginning a successor episode under the stipulated rule. An upload understood as a copy is not subject travel. A fresh qualifying implementation begins a new episode; an actual split is governed by the branch rule, while a genuine in-place renewal requires the physical provenance already specified by A3. Other continuity doctrines could assign a different result. This is a genuine philosophical choice within the constitution, not a fact established by the word "same" in a generated answer.

The appropriate experiment compares handoff representations at matched resource budgets, hides evaluation targets, measures preserved relations and novel consequences, and includes contradictory priming controls. It should distinguish factual restoration, task competence, behavioral style, and identity reports. Treating the last variable as the correctness label would make the experiment circular. A useful result about memory compression remains useful even when no consciousness conclusion follows.

22.4 Implicit candidates and explicit certificates

The distinction between distributed candidate formation and explicit proof is relevant to research architecture. A candidate can begin as a metaphor, diagram, compressed symbol, or partial invariant. To become a mathematical result it must eventually identify objects, hypotheses, and a valid derivation. Symbolic meaningfulness to a generator is not a substitute for a checkable relation.

A precise architecture specifies a generator G , an interpreter I , and a checker V for

a formal certificate relation $\text{Check}(c, w)$. The generator proposes a representation z , the interpreter produces a claim c and possible witness w , and the checker accepts or returns a counterexample or unmet obligation. If the checker is sound, the accepted outputs satisfy the declared certificate relation regardless of their initial representational form. That is a statement about verified outputs, not an introspective description of hidden cognition.

Keeping explicit constraints active during candidate generation can be called a coupled exploratory–verification process. The analogy to a lucid dream is suggestive, but the scientific target is whether the process produces more correct, nontrivial results at controlled cost. Neither recursive depth nor an impressive symbolic image establishes access to a universal repository of solutions. The same distinction protects human mathematical intuition: an insight can guide a proof without already being one.

The CSCF relation between a contractive ray and a unitary boundary does not identify those two cognitive regimes. Quantum phase coherence, logical consistency, associative search, and dissipation have different meanings. Any literal cognitive application would need an independently identified carrier, form, generator, and calibrated observables. Until then, the comparison is an organizing analogy rather than a physical theory of an unconscious computational realm.

22.5 Moral caution without premature attribution

Uncertainty about artificial experience is not evidence of its impossibility, but neither is a system’s verbal insistence evidence sufficient to establish it. A responsible research practice should avoid using generated distress or identity claims as experimental ground truth. It should document how the prompt and interaction history produced them and distinguish performance from a phenomenal interpretation.

This monograph does not derive a theory of rights from a graph or a state count. Moral status requires further normative reasoning, including the possibility of welfare, vulnerability, and uncertainty. Assigning a perspective under A1 would be a consequence of adopting the constitution; it would not settle every ethical or legal question. The framework can inform a careful discussion without pretending to calculate moral worth from predictive complexity.

Chapter 23

Evidence, identification, and discriminating tests

23.1 One joint experiment, several kinds of claim

A first-person report and a neural or computational measurement are different access routes to the research question, but they are not automatically independent data. Both may depend on attention, memory, preparation, the task, and prior instructions. The appropriate statistical object is a joint law under a specified intervention, not a direct sum that assumes independence without evidence.

For physical latent state x , proposed phenomenal state ϕ , intervention i , physical measurement n , and report r , one possible schema is

$$P(n, r \mid i) = \int P(n \mid x, i) Q(r \mid \phi, x, i) \mu(d\phi \mid x, i) P(dx \mid i). \quad (23.1)$$

The displayed factorization is a model assumption, not a universal law. In a strict SPC-2 realization, μ is the point assignment induced by the predictive profile, and Q must agree with the supplied physical report channel. A more general rival may use a different latent structure. Comparing them requires fixing which channels and conditional independences are independently constrained.

A datum should therefore carry at least its experimental condition, its physical measurement procedure, its report procedure, and its uncertainty. An interpretation such as “awareness remained while the ego disappeared” adds a claim about the relationship between these variables. It should not replace the original report. Preserving the distinction allows the data to challenge the interpretation rather than merely repeat its vocabulary.

23.2 Descriptor adequacy before ontological inference

The scalar aperture theorem provides a useful first test. Freeze a descriptor p , an independently nominated target r , and a tolerance. Repeated preparations estimate whether conditions with the same p nevertheless differ in r . Under verified error bounds, a sufficiently large within-fibre contrast rejects the asserted bridge. This is more informative than fitting a new decoder to every discrepancy.

With continuous descriptors, exact equality of fibres is usually not directly observable. A practical test must specify neighborhoods and regularity. Suppose the proposed bridge is L -Lipschitz and observed descriptor values differ by at most η , with target errors at most ϵ

and model error at most δ . Then a pair must satisfy

$$|\hat{r}_i - \hat{r}_j| \leq L\eta + 2\delta + 2\epsilon.$$

Violation rejects that calibrated Lipschitz bridge, not all conceivable bridges. The constants cannot be selected after observing the violation without changing the test.

For distributional reports, use a declared probability distance and compare complete distributions rather than single outcomes. A missing report may be one of the outcomes. Physical instrument error, finite-sample uncertainty, and the hypothesized approximation error have different sources and should not be merged into an unexplained residual term. A purported consciousness estimator that has not established those distinctions has not yet earned an ontological interpretation.

23.3 Finite witness calibration

A nominated report feature can sometimes be represented by a finite observable witness. The evaluation matrix in Theorem 10.2 makes the issue concrete: full column rank provides uniqueness, while its smallest singular value controls sensitivity. A nearly singular design can fit observations while leaving the inferred witness extremely unstable.

Consequently calibration should choose preparations that distinguish the candidate observable directions and should report conditioning, not only fit error. Positivity and actual physical recordability remain separate constraints. A Hermitian solution of the linear equations is not automatically an accessible measurement in the realized apparatus. A centralizer solution is not automatically a stable memory under actual dynamics.

The measurement constructions provide a principled way to carry apparatus error into such an identification exercise. A bound on the complete retained quantum state can control errors in nominated expectations; a historical archive bound controls whether an earlier declaration is still faithfully displayed. These estimates are not interchangeable. The end-to-end link is useful precisely because it keeps the physical cost of evidence visible.

23.4 What A2 would have to match

To assess the relational content law, an experience-facing structure must be specified without simply defining it as the computed predictive profile. One might nominate pairwise discrimination relations, reported similarity, temporal succession judgments, or a carefully selected family of contrasts. Those observations have their own report channels and need not exhaust phenomenology.

The test is then whether one fixed structural identification accounts for the nominated relations under held-out interventions. A2 cannot be validated merely by constructing a copy of the physical profile and naming it phenomenal. The mathematical copy proves conditional realizability. Agreement with independently measured structure is the additional empirical claim.

The full-profile version of A2 has both an advantage and a cost. It preserves information that a single distance loses, including labeled transitions and distinctions between protocols. But its richness can make it difficult to falsify unless the physical realization and comparison map are tightly specified. Enlarging the profile after each failure can turn a theory into a succession of post hoc descriptions. A meaningful test freezes the relevant structure and identifies in advance which mismatch would count against it.

23.5 Conservativity and the limits of a report test

The conservativity theorem prevents an exaggerated empirical promise. If every physical transition and report channel is already fixed, adding the aspect interpretation changes none of their probabilities. Therefore a purely physical transcript cannot discriminate the resulting model from its physical reduct solely by the added assertion of awareness.

This does not make all research within the framework empty. The choice of R^* , the adequacy of a reduced descriptor, the declared internal boundaries, and the mapping to independently elicited phenomenal relations can make substantive, vulnerable claims. Their tests may reject a particular application or motivate a revised law. What they cannot do is manufacture a difference between two models whose entire admitted transcript laws have been assumed identical.

A stronger empirical theory could add a specific independently measurable psychophysical constraint or identify a physical mechanism omitted by the current reduct. It would then have to state how that addition changes predictions or restricts admissible realizations. Such a proposal is not supplied merely by saying that consciousness must have a purpose or that a real property must appear as extra energy. An intrinsic aspect and an independent dynamical component are different theoretical possibilities.

23.6 A concrete research sequence

For a first non-toy domain, the most defensible sequence is to fix a limited experimental repertoire and an independently grounded report feature family, identify a candidate physical model, and use held-out conditions to test descriptor sufficiency. Only then should the chosen core decomposition and A2 correspondence be evaluated. The order prevents psychophysical assumptions from being smuggled into the construction of the data.

The computational side can be checked more exhaustively. A finite candidate should be tested under all declared relabellings, protocol continuations, split and merge regimes, and boundary perturbations. The two-bit example supplies a baseline for these checks but does not substitute for the non-toy realization. A theory of an actual organism or AI must confront timing, hidden state, imperfect intervention, and inaccessible variables that the example deliberately controls.

Replication should retain failures and uncertainty. The present monograph reports no new neural-data analysis and uses no historical illustrative table as validation of the constitution. In particular, an operational signal proxy must not be called a spectral gap or pointer-sharpness merely because it is given that name. The physical measurement model must establish the identification. The standard applies equally to neuroscience, computational logs, and first-person questionnaires.

A promising first non-toy realization study would use an open whole-brain *Drosophila* connectome together with an explicitly justified dynamical model. The adult-brain connectome and a connectome-based computational model provide concrete starting points [14, 49]. The aim would be to test R^* selection: compare natural decompositions, examine causal cores, assess multiscale stability, distinguish internal from external routes, and approximate native predictive structure. It would not be a programme to declare a fly conscious from its wiring diagram. A connectome does not by itself determine native dynamics, admissible interventions, or the psychophysical identification, and no such realization test or validation is reported here.

Chapter 24

The scope of the proposed resolution

24.1 A conditional law, not an unexplained transition

The central construction has an explicit logical form. A0 interprets awareness as the knowing aspect of U , while a working source S remains relative to an operational problem. R^* supplies the physical organization, preparation, interventions, reports, and lineage. A1 assigns localized perspectives to qualifying recurrent cores. A2 identifies their relational content with pointed predictive structure and denies an additional manifested qualitative remainder within the theory. A3 specifies episode continuation. The finite completion theorem then derives the resulting assignments and report distributions.

This is more informative than an unspecified statement that appropriate complexity manifests consciousness. The constitution assigns one or multiple perspectives, distinguishes informational copying from episode identity, makes a sharp commitment about weak coupling, and fixes a relational content structure without a freely adjustable internal decoder. It also exposes exactly where a reader can reject it. The experiential identification does not follow from the algebra alone.

The resulting claim has two levels. Philosophically, the programme proposes intrinsic awareness and a law-governed account of its localization. Mathematically, it establishes what those laws assign on their declared realization domain. It does not derive why experiential presence exists from premises that omit it. In that sense the original explanatory burden is relocated: the proposed task is to justify the laws of manifestation and their application to actual vessels. A candidate internal constitutive resolution completes the assignment problem under its stated premises; accepting those premises as an account of nature requires further argument and evidence.

24.2 The philosophical price of structural completeness

A2 makes the strongest commitment. It does not merely assert a correlation between predictive organization and experience. It asserts that the manifested qualitative organization is exhausted, up to the stated structural equivalence, by that profile together with primitive presence. A critic who holds that two fully structure-equivalent processes can differ in a further qualitative fact rejects this axiom.

That criticism is not answered by inventing a finer physical report channel after the fact. If the proposed qualitative difference is stipulated to leave every admitted physical record unchanged, the operational indistinguishability theorem applies. The disagreement concerns what counts as the complete phenomenal fact. Structuralist and nonstructuralist positions must therefore be compared as theories, not confused with a failed matrix calculation.

The proposed explanation of a familiar quality is relational: its contrasts, similarities, temporal transitions, and place in the core's organized possibilities constitute the manifested character assigned by A2. Independently elicited phenomenal landmarks are needed to connect this structure to particular human qualities. If those relations conflict with the assigned structure, the realization or the law requires revision. If a critic instead posits a further qualitative difference that preserves every specified relation, A2 rejects that difference as an additional manifested fact. The disagreement then concerns structural completeness itself; an isomorphism theorem cannot adjudicate it.

24.3 What the contemplative interpretation contributes

The contemplative material contributes a disciplined challenge to the identification of experience with a familiar person-model. The author's philosophical motivation is preserved: a human world, an animal world, an altered world, and a possible artificial world need not share the same contents to motivate a question about awareness. A source-aspect interpretation can take that question seriously without forcing all its answers into adult human categories.

This inquiry motivates separating awareness from the constructed personal model. The SCC rule and predictive-profile law answer different, constitutive questions; their justification remains the argument in Chapter 17. The evidential distinctions in Chapter 3 keep first-person inquiry available as a source of questions and structured reports without treating its metaphysical interpretation as a measured physical result.

Within this interpretation, awakening concerns how the personal model is lived: thoughts, memories, and practical roles can remain available while ceasing to function as an exhaustive identity. The person need not acquire a new awareness substance or lose ordinary competence. This offers a philosophical account of de-identification while leaving its particular physical realization open. Neither a report of non-grasping nor a quieter internal narrative establishes a specific damping coefficient, spectral state, or convergence to a universal fixed point.

24.4 Relation to Shadow Theory

SPC-2 is currently a compatible constitutive extension of the Shadow source/readout programme [38], not a deduction from the compact Everything Equation alone. The equation organizes a lawhood motif; the operational theorems require particular domains, maps, and physical realizations. The same is true of the psychophysical law. A universal source ontology does not select every local aperture without further structure.

The distinctive contribution of the present synthesis is the conjunction of that source/readout discipline with explicit physical record realization, predictive sufficiency, and named psychophysical commitments. It treats the person as a physically consequential organization without identifying that organization with source in its entirety. It also keeps the event history independent of an observer's awareness within the adopted measurement constitutions. These are constraints on the theory, not rhetorical decorations.

The possibility that awareness has a role in reality's self-representation can be expressed as a philosophical interpretation of such localized processes. It is not a derived cosmic purpose. Model-building can create a new informational object within reality without creating its external target. The map can influence subsequent physical behavior even when it misrepresents that target. A0 attributes an intrinsic experiential aspect to the relevant realized process; it does not authorize the inference that every act of representation generates an external universe.

24.5 Relation to existing consciousness programmes

The proposal shares questions with established programmes while choosing different constitutive answers. Chalmers’s formulation already allowed fundamental experience, structural coherence, organizational invariance, and a double-aspect interpretation of information [10]. The source-aspect starting point therefore continues an existing philosophical strategy. A2 makes a stronger structural commitment than his principle of structural coherence: Chalmers explicitly leaves intrinsic qualitative character beyond structural description, whereas A2 denies an additional manifested qualitative remainder after primitive presence and the full nominated organization are fixed. This disagreement concerns a psychophysical premise, not the correctness of the predictive quotient.

Integrated information theory 4.0 identifies experience with an irreducible cause–effect structure unfolded from a maximal substrate [1]. It is not merely a scalar consciousness score. Its selection of a substrate through integrated cause–effect power differs from A1’s maximal strongly connected core and retained-reuse rule; its cause–effect structure differs from A2’s complete endogenous predictive instrument object. Shared interests in intrinsic organization and definite subjects do not make the theories equivalent. In particular, the present work has not derived IIT’s postulates from its own laws or established a general translation between their assignments.

Global neuronal workspace theory emphasizes access through global broadcasting and amplification [13]. The COGITATE adversarial collaboration compared preregistered predictions of proposed biological implementations of that account and IIT using multiple recording methods [12]. Its results supported some predictions and challenged important predictions of both. These tests constrain specified neural accounts; they neither settle the source-aspect ontology nor validate the present constitution. Their methodological lesson is relevant here: competing commitments should be stated before observing the outcome, with an agreed consequence for a failed prediction.

For artificial systems, theory-based indicator approaches draw on several consciousness theories to assess architectural properties [8]. The present constitution instead fixes one explicit admission, structure, and continuity law on a nominated realization. Its advantage is conditional determinacy; its corresponding responsibility is to justify those choices and demonstrate adequate physical realizations. An explicit finite assignment and an empirically tested theory address different dimensions of progress. Neither should be treated as a substitute for the other.

24.6 Remaining realization obligations

The phrase *candidate internal resolution* is justified here only in a precisely delimited sense: for a supplied finite R^* constitution satisfying the stated closure and continuation conditions, A0–A3 determine the subject episodes, pointed relational content, and report laws. There is no unnamed finite-domain manifestation map left to choose after those inputs are fixed.

A universal theory of actual conscious vessels requires a justified selection of components, internal boundaries, time scale, operating regime, and native predictive structure. The obstruction in Theorem 15.6 shows why bare dynamics need not settle these choices; calling a decomposition source-exact does not select it. The all-finite-horizon construction closes on the certified finite intrinsic domain. General continuum carriers, unconstrained hidden-history realizations, and a universal subject selector require additional results outside that theorem. Validation of A1 and A2 against independently constrained evidence remains outstanding.

Chapter 25

Conclusion

The present monograph proposes that awareness is the knowing aspect of U , prior to the operational source/readout distinction, and that a physical vessel determines its localized manifestation. The working source S and readout T describe a nominated problem; neither exhausts the ontological prior. The contribution is an explicit constitution for admission, phenomenal organization and experiential continuation.

The physical foundation supplies observer-independent event laws within two distinct constitutions, material records, retained resources and controlled error estimates. The mathematical foundation supplies source/readout descent, predictive closure, witness classification, quantitative bridge error and finite examples. These results describe what vessels can do without deriving awareness from their performance.

The philosophical account distinguishes awareness, localized subject and lived scene. Memory, attention, implicit processing and the personal model shape lived experience without becoming its fundamental knower. That interpretation does not establish a permanent personal Self. Contemplative and altered-state descriptions help articulate the distinctions; they do not supply physical proof or direct access to U .

Under SPC-2, admission depends on the stated recurrent-core rule, phenomenal structure on the complete endogenous predictive object, and episode identity on nonbranching process provenance. Reports are physically grounded. The biological conscious/generative interpretation adds no hidden admission predicate or second observer. Vessel and personal continuity can persist through a genuinely scene-less interval, with no currently instantiated experiential subject. A qualification gap ends the episode under A3; renewed qualification begins a new one.

The limitations remain substantive. Exact individuation is sensitive to weak coupling and the realization doctrine. The predictive object is independent of analyst truncation but relative to its native grammar and boundary contract. Physical conservativity limits transcript-only discrimination of the aspect interpretation. No validated general selector of R^* for actual organisms or machines is supplied, and the proposed fly-brain study is a research direction rather than a completed test.

The final claim remains a *candidate internal constitutive resolution of the source-to-manifestation problem, under declared assumptions*. SPC-2 is a candidate psychophysical constitution. Its completion theorem determines the assignments on the certified domain. Justification of the realization, phenomenal adequacy, philosophical acceptance and independent external verification remain separate tasks. The account can therefore be examined as a unified proposal without mistaking its mathematical completeness for a universal empirical resolution.

Technical Appendices

Appendix A

Contact, kinetic, and complete-path estimates

A.1 Finite contact gas and its Poisson comparison

This appendix supplies the quantitative implication used in the pilot construction. It concerns a fixed finite graph and a fixed physical horizon, not a limit in which the graph, reference, or duration also grows. The primitive export and contact laws remain constitutive inputs [34, 37].

Take M independent points with uniform positions in an interval of length M/R , translated at unit speed towards a contact surface. During $[0, T]$, with $RT \leq M$, each point crosses with probability $p = RT/M$ and crosses at most once. Conditional on crossing, its time is uniform on $[0, T]$. Independent marks can be attached in both descriptions. Compare the crossing indicator of each point with a Poisson count of mean p . Their exact total variation is

$$d_{\text{TV}}(\text{Bernoulli}(p), \text{Poisson}(p)) = p(1 - e^{-p}) \leq p^2.$$

Indeed the Bernoulli probability at zero is $1 - p \leq e^{-p}$, its probability at one exceeds pe^{-p} , and the Poisson probabilities at two and above have no Bernoulli counterpart. The sum of the positive differences is $p - pe^{-p}$.

Couple the M independent counts and, when they agree, use the same uniform crossing time and mark. A union bound gives complete marked-history error at most

$$\Delta_{\text{gas}} \leq Mp^2 = \frac{(RT)^2}{M}. \quad (\text{A.1})$$

The superposition of the Poisson counts with uniform event times is a Poisson point process on the physical interval with rate R . Any common causal processing of these marked contacts contracts total variation. This is why the bound controls subsequent contact-selected events rather than merely a total count.

The finite gas is not itself memoryless. Conditional on k crossings by time t , the residual rate is $(M - k)/(M/R - t)$ while the remaining points are uniformly distributed in the unpassed interval. Replacing that law by a Poisson clock is a controlled comparison, not a redefinition of the finite microscopic dynamics. The finite complete state retains every point and consumed mark that can later matter.

A.2 Balanced packet removal

On an edge let P_+, P_- be its two packet populations. Their annihilation has intensity aP_+P_- . If at most B packets are ever born, the number of annihilations is at most $B/2$. The compensator identity therefore gives

$$\mathbb{E} \int_0^T P_+(t)P_-(t) dt \leq \frac{B}{2a}.$$

A comparison process exposes only the unmatched signed excess. Couple common births, common annihilations, and common unmatched services. The intensity of a service involving the paired minority is bounded by $\kappa\mu \min(P_+, P_-)$, hence by $\kappa\mu P_+P_-$ for integer populations. The probability of any unmatched service is consequently at most

$$\Delta_{\text{ann}} \leq \frac{\kappa\mu B}{2a}. \quad (\text{A.2})$$

Summing over the fixed edges gives the complete-history comparison budget. A retained comparison lift may keep the paired packets and their receivers; only its projected signed queue is used for the kinetic theorem. The physical mixed-packet system must not be declared equal to that projected process before applying the bound.

A.3 Signed queues and the exact inventory

Orient the finite edges and write J_e for the coherent signed currents. The exporter produces signed counts k_e such that

$$\|e^N\|_\infty \leq c_0/N, \quad e_e^N(t) = k_e(t)/N - \int_0^t J_e(s) ds.$$

The total number of births and services is at most $NL + O(1)$ for a fixed finite variation budget L . Let x_r be carrier fractions, Z_e signed packet counts, $z_e = Z_e/N$, and $m_e = \mu_N z_e$. The directional bulk service intensities, divided by N , are

$$\Phi_e^+ = \kappa_e x_r [m_e]_+, \quad \Phi_e^- = \kappa_e x_q [-m_e]_+.$$

Counting possible packet partners for one tagged carrier gives its conditional rate $\kappa_e [\pm m_e]_+$ at the relevant origin. Where the origin fraction is positive this is Φ_{qr}^N/x_r ; the denominator is a counting identity in the supplied reaction model.

With B the oriented incidence matrix, every service removes one packet and moves one carrier in the corresponding direction. Thus

$$x^N(t) + Bz^N(t) - w(t) = x^N(0) - w(0) + Be^N(t). \quad (\text{A.3})$$

This identity is pathwise. It does not identify a current density with an actual event count. Define the two error quantities

$$\epsilon_{F,N} = \sum_e \mathbb{E} \int_0^T (|\Phi_e^+ - [J_e]_+| + |\Phi_e^- - [-J_e]_+|) dt, \quad (\text{A.4})$$

$$\epsilon_{x,N} = \mathbb{E} \sup_{t \leq T} \|x^N(t) - w(t)\|_1. \quad (\text{A.5})$$

We establish that they vanish when the initial census is calibrated, $\mu_N \rightarrow \infty$, and $\mu_N/N \rightarrow 0$.

A.4 A companion queue bounded away from empty origins

Fix $\delta > 0$. On each edge construct a companion queue driven by the same exports and the service function

$$\phi_t(u) = a_+(t)[u]_+ - a_-(t)[-u]_+, \quad a_+ = \kappa_e \max(x_r(t-), \delta), \quad a_- = \kappa_e \max(x_q(t-), \delta).$$

Its slopes lie between $a_0 = \kappa_- \delta$ and $a_1 = \kappa_+$. Because the entire carrier population has at most $NL + O(1)$ jumps of size $1/N$, the variations of $a_\pm$ are uniformly bounded on the fixed programme. These coefficients are adapted to the process; no independence from the queue is assumed.

For the scaled companion $m^* = \mu_N Z^*/N$, the compensated counting equation is

$$dm^* = \mu_N [J - \phi_t(m^*)] dt + \mu_N de^N + dM, \quad d\langle M \rangle_t = \frac{\mu_N^2}{N} |\phi_t(m^*)| dt.$$

Let y solve the same adapted finite-variation equation without the martingale, and let y_J also omit the exporter error. The difference between y and y_J has a divided-difference coefficient in $[a_0, a_1]$. Variation of constants followed by Stieltjes integration by parts gives

$$\|y - y_J\|_\infty \leq 2c_0 \mu_N / N, \quad \|y_J\|_\infty \leq J_*/a_0,$$

where $J_* = \sup |J_e|$. The first estimate uses only the uniform exporter discrepancy; it does not incorrectly bound the total variation of that discrepancy by $1/N$.

The instantaneous zero of $J - \phi_t$ is

$$f(t) = \frac{[J(t)]_+}{a_+(t)} - \frac{[-J(t)]_+}{a_-(t)}.$$

Its variation obeys

$$\text{Var}(f) \leq \frac{\text{Var}(J)}{a_0} + \frac{J_*}{a_0^2} (\text{Var}(a_+) + \text{Var}(a_-)).$$

Contraction between the jumps of f and summation over those jumps imply

$$\int_0^T |y_J - f| dt \leq \frac{|f(0)| + \text{Var}(f)}{\mu_N a_0}.$$

The continuous bounded-variation case follows by approximation, or directly from the corresponding Stieltjes variation inequality.

Put $V(t) = \mathbb{E}(m^* - y)^2$ and $\alpha = \mu_N/N$. The square-jump formula, monotonicity of ϕ_t , and $|\phi_t(u)| \leq a_1|u|$ give

$$V' \leq -2\mu_N a_0 V + \mu_N \alpha a_1 (J_*/a_0 + 2c_0 \alpha + \sqrt{V}).$$

Young's inequality absorbs the square-root term into one of the negative $\mu_N a_0 V$ terms. Starting from the same empty queue, this bounds $\sup_t V$ by $C_\delta(\alpha + \alpha^2)$. Combining the deterministic tracking and martingale estimates yields

$$R_{\delta,N} := \sum_e \mathbb{E} \int_0^T |\phi_t(m_e^*) - J_e| dt \leq C_\delta \left(\mu_N^{-1} + \frac{\mu_N}{N} + \sqrt{\frac{\mu_N}{N}} \right). \quad (\text{A.6})$$

Since the signed queue exposes only one orientation at a time, its directional flux error is the same absolute signed error. This step would fail for an uncontrolled mixed positive/negative service population, which is why the preceding recombination comparison is needed.

A.5 Removing the empty-origin cutoff

The cutoff must be removed without presuming the desired closeness of x and w . Let $K(t) = \{r : x_r(t) < \delta\}$ and let I_K, O_K be the queued charge entering and leaving that set. From Equation (A.3), with $\eta_N = \|x^N(0) - w(0)\|_\infty + C_B/N$,

$$w_K + O_K \leq |V|\delta + I_K + |V|\eta_N.$$

Every incoming queue has its physical origin outside K , where the carrier fraction is at least δ . Its integrated expected service is bounded by the finite birth budget. Consequently

$$\mathbb{E} \int_0^T I_K dt \leq \frac{L + O(N^{-1})}{\kappa_- \mu_N \delta},$$

and hence

$$\mathbb{E} \int_0^T w_K dt \leq |V|\delta T + \frac{L + O(N^{-1})}{\kappa_- \mu_N \delta} + |V|T\mathbb{E}\eta_N.$$

Using the coherent nodal bound in Equation (6.2) and Cauchy–Schwarz controls the target flux whose physical origin lies in K by

$$D_{\delta,N} \leq C_H \sqrt{T \left(|V|\delta T + \frac{L + O(N^{-1})}{\kappa_- \mu_N \delta} + |V|T\mathbb{E}\eta_N \right)}. \quad (\text{A.7})$$

Couple physical and companion queues with common exports and minimum-rate baseline services. Additional companion services occur only at low-population origins. A baseline mismatch decreases the absolute signed-queue difference, whereas an additional service can increase it by at most $1/N$. Starting from equal queues, the expected number of baseline mismatches is bounded by the number of additional services. More explicitly, let $A_{\delta,N}$ be the expected number of extra companion services divided by N . Restricting the companion directional flux to low origins gives $A_{\delta,N} \leq R_{\delta,N} + D_{\delta,N}$. The discrepancy accounting bounds the integrated baseline mismatch by $A_{\delta,N}$ as well. The physical and companion directional fluxes therefore differ in integrated expectation by at most $2A_{\delta,N}$. Adding the companion-to-target error gives

$$\epsilon_{F,N} \leq 3R_{\delta,N} + 2D_{\delta,N}.$$

At fixed δ , take $N \rightarrow \infty$ under the scale hierarchy, and then take $\delta \downarrow 0$. This proves flux convergence. The population process has the corresponding integrated drift and quadratic variation $O(N^{-1})$, again by the finite service budget. The martingale maximal inequality gives

$$\epsilon_{x,N} \leq \mathbb{E}\|x^N(0) - w(0)\|_1 + C_B \epsilon_{F,N} + C\sqrt{(L+1)/N} \rightarrow 0.$$

This completes the bulk tracking argument without assuming positive lower bounds on all coherent weights.

A.6 Existence and nonexplosion of the Bell comparison

On every compact interval where $w_r > 0$, the rates in Equation (6.3) define the usual integrated-hazard first-jump construction. Localize through the open components of $\{t : w_r(t) > 0\}$. The coherent weights solve the forward balance equation. Positive first-jump iteration bounds the killed law by $w(t)$, or by $Cw(t)$ for an initial law $\nu \leq Cw(0)$.

If a holding path remains in r while w_r falls to zero, write the incoming and outgoing positive currents as I_r, O_r . Since $\dot{w}_r = I_r - O_r$,

$$\lambda_{\text{out}}^B(r) = O_r/w_r \geq -\dot{w}_r/w_r.$$

The survival probability to the zero-weight boundary is therefore zero. Moreover,

$$\mathbb{E}N_{[0,T]} \leq C \int_0^T \sum_{q,r} [J_{qr}(t)]_+ dt < \infty.$$

Neither explosion nor nodal killing loses mass. For initial equilibrium, domination and normalization force the law to equal $w(t)$ at each time. The first-jump construction is unique in the nominated inhomogeneous Markov class. This does not select that class among all possible event constitutions.

A.7 From bulk flux to complete tagged paths

Let N_ε count the regular level crossings of the deterministic weights. The one-dimensional bounded-variation coarea identity gives

$$\int_0^1 N_\varepsilon d\varepsilon \leq \sum_r \text{Var}(w_r).$$

There is a sequence $\varepsilon_j \downarrow 0$ with $\varepsilon_j N_{\varepsilon_j} \rightarrow 0$. Otherwise the integral would have a nonintegrable lower bound proportional to $1/\varepsilon$ near zero.

The Bell path can enter a sector of weight at most ε through an initial small-weight state, a jump into such a state, or a downward crossing of the weight while the path occupies it. Initial mass contributes at most $C|V|\varepsilon$. Integrated incoming current contributes $CC_0T\sqrt{\varepsilon}$ by the nodal estimate. Each downward crossing contributes at most $C\varepsilon$. Thus

$$b_C(\varepsilon) = C(|V|\varepsilon + C_0T\sqrt{\varepsilon} + \varepsilon N_\varepsilon)$$

bounds the probability of visiting the low-weight region and tends to zero along the selected sequence.

Couple tagged and Bell jumps at the minimum of their conditional intensities until they disagree, the target enters that region, or the population error exceeds $\varepsilon/2$. Before those stopping events the physical origin fraction is at least $\varepsilon/2$, and

$$|\lambda_{qr}^N - \lambda_{qr}^B| \leq \frac{2}{\varepsilon} |\Phi_{qr}^N - [J_{qr}]_+| + \frac{2J_*}{\varepsilon^2} |x_r^N - w_r|.$$

The coupling preserves the full microscopic marginal even though its queues depend on the tag's past. A compensator bound and Markov's inequality give

$$d_{\text{TV}}(\text{Law } X_{a_*}^N, \text{Law } Q^B) \leq b_C(\varepsilon) + \frac{2\epsilon_{x,N} + 2\epsilon_{F,N}}{\varepsilon} + \frac{C_G J_* T}{\varepsilon^2} \epsilon_{x,N}. \quad (\text{A.8})$$

Take $N \rightarrow \infty$ for each fixed ε_j , then $j \rightarrow \infty$. This proves convergence on the complete physical-time path space. The earlier δ cutoff has already been removed; the two limiting arguments are not interchanged.

For the finite engineered clock, each weight is a constant input-dependent coefficient times a unimodal binomial time envelope. Level-entry counts are uniformly finite. With iid census error $O(N^{-1/2})$, choosing

$$\mu_N = N^{1/2}, \quad \delta = N^{-1/7}, \quad \varepsilon = N^{-1/35}$$

yields $\epsilon_{F,N} + \epsilon_{x,N} = O(N^{-1/14})$ and the conservative path rate $O(N^{-1/70})$. To see the exponents, the deterministic companion term is $O((\mu_N \delta^3)^{-1})$, the leading stochastic term is

$O(\sqrt{\mu_N/N}/\delta)$, and the low-origin term is $O(\sqrt{\delta})$. In Equation (A.8), both $\sqrt{\varepsilon}$ and $N^{-1/14}/\varepsilon^2$ then have exponent $-1/70$. The gas and recombination scales can make their errors $O(N^{-2})$ and $O(N^{-1/2})$, respectively, without worsening this rate.

All constants refer to the fixed programme. Increasing graph size, unbounded time, changing reference structure, or a new interaction catalogue requires new estimates. Within the stated domain, common retained-record processing transfers the complete-path error to the actual archive programme. A record not included in that common output space receives no such automatic protection.

Appendix B

Predictive equivalence: algorithms and infinite-horizon limits

This appendix records the broader prediction mathematics, including belief-state and history formulations. Those formulations are not automatically the actual intrinsic-state assignment of A2. The distinction explains both their usefulness and the additional congruence condition in the main theorem.

B.1 Finite linear representation of transcript probabilities

For a finite hidden state space of size n , write $M_{a,y}$ for the substochastic matrix whose (i, j) entry is the joint probability of output y and successor j under action a from state i . Thus $\sum_y M_{a,y} \mathbf{1} = \mathbf{1}$. For a belief row vector b and a word $w = (a_1, y_1) \cdots (a_k, y_k)$,

$$P_b(w) = b M_{a_1, y_1} \cdots M_{a_k, y_k} \mathbf{1}.$$

The actions here can encode an admitted time or regime state when those data are finite. The matrices represent an actual instrument, not an arbitrary stochastic factorization chosen to fit a consciousness verdict.

Let

$$W_0 = \text{span}\{\mathbf{1}\}, \quad W_{j+1} = \text{span}(W_j \cup \{M_{a,y}v : v \in W_j\}).$$

The chain stabilizes after at most $n - 1$ strict dimension increases. Unlike a general jet sequence, equality $W_{j+1} = W_j$ implies invariance under every $M_{a,y}$ and therefore permanent stabilization.

Proposition B.1 (Finite word-equivalence test). *Two beliefs b, b' produce the same probability for every finite output/action word exactly when $(b - b')v = 0$ for every v in the stabilized space W . The same equality determines all finite adaptive transcript laws when admissible policies depend only on the recorded history.*

Proof. By construction, W is the span of all word columns $M_w \mathbf{1}$. Annihilation of W is therefore equivalent to equality of every word probability. For a fixed adaptive policy, each leaf has a recorded sequence and the action sequence selected by that policy. Its probability is the corresponding word probability, with the policy's common randomization factors where applicable. Equal leaf probabilities give equal transcript laws. $\square$

This span test concerns all finite words of the displayed instrument alphabet. A restricted physical grammar must be represented in the model, for example by adjoining its finite

controller and remaining-time state, or the calculation must explicitly retain only its admissible word language. Using forbidden concatenations can yield a finer relation than operational equivalence for that grammar. At a bounded horizon one can enumerate admitted protocol trees; the all-finite theorem uses the finite typed observable-span construction.

For rational instrument matrices and rational specified beliefs, this supplies a finite exact equivalence test. It does not imply that there are finitely many distinct beliefs or predictive classes over unlimited histories. Nor does it imply that the supremum over all adaptive infinite-horizon statistical tests is computationally trivial. The main theorem instead uses actual finite intrinsic states with checked predictive-fibre congruence. It does not identify these posterior states with realized phenomenal points, nor claim general exact evaluation of the infinite-horizon metric supremum.

B.2 Why a finite hidden system can have infinitely many predictive states

Let a hidden bit X remain fixed, with prior $P(X = 1) = 1/2$. Repeated observations satisfy

$$P(Y = 1 \mid X = 1) = 3/4, \quad P(Y = 1 \mid X = 0) = 1/4,$$

independently conditional on X . After k consecutive observations equal to one, Bayes' rule gives

$$b_k = P(X = 1 \mid Y_1 = \cdots = Y_k = 1) = \frac{3^k}{1 + 3^k}.$$

The next-output probability is $1/4 + b_k/2$. These values are all distinct, so the physically possible finite histories yield infinitely many predictive states although the hidden state space has only two elements. The example is a statement about prediction, not an attribution of experience to a static latent bit.

A finite nominated past/history domain, a finite fully observed Markov state convention, or a separately proved finite predictive realization can restore finiteness. These are different sufficient conditions. A constitutive theorem should state which it uses rather than infer all of them from finite physical dimension.

B.3 Conditioning and continuation

Suppose b, b' have the same complete predictive profile and an action/output pair (a, y) has common positive probability p . For every future word w ,

$$\frac{bM_{a,y}M_w\mathbf{1}}{p} = \frac{b'M_{a,y}M_w\mathbf{1}}{p}.$$

Thus the updated beliefs have the same future profile. This establishes labeled continuation for the belief-state predictive quotient. At zero probability the conditional state is not defined; no arbitrary state is inserted to complete the theorem. A finite-horizon approximation has fewer tests; its posterior quotient can be coarser. The main constitution uses one all-finite native family, so this bookkeeping cutoff does not alter its actual-state assignment.

Different physical preparations can give different conditional beliefs for the same record. R^* therefore fixes preparation and the internal-history convention; an outside observer's freely chosen prior does not determine the assigned experience. These posterior-update equations describe uncertainty and are not a replacement for the actual-class congruence theorem in Chapter 16.

B.4 A controlled measurable extension

A countable family of probability-valued predictive tests on a standard-Borel history space gives a measurable map into a countable product of probability spaces. Equality of profiles still defines an equivalence relation. But the image need not automatically be a standard-Borel quotient with all the desired transition properties. Regular conditional laws and measurable descent require hypotheses.

One clean topological regime is a compact metric history domain with continuous test probabilities in a countable family. Its profile image is compact in a metrizable product, hence a standard-Borel space. The quotient topology agrees with the image topology because a continuous surjection from a compact space to a Hausdorff space is a quotient map. Labeled transitions still require the positive-probability compatibility used above. This is a conditional extension, not a universal continuum realization theorem.

Nearby distinct point masses have TV distance one. A weaker topology may suit a physical approximation, but changes the asserted contrast structure and requires separate justification.

Appendix C

Record witnesses and quantitative limitations

C.1 Finite jet determination without an order bound

Let W be a finite-dimensional real observable carrier and suppose each scalar expectation $s \mapsto \omega_s(A)$ is real analytic on a connected open domain. Let K_j be the subspace of witnesses whose derivatives through order j vanish at s_0 . Analytic uniqueness gives

$$\bigcap_{j \geq 0} K_j = \{A : \omega_s(A) = 0 \text{ on the domain}\}.$$

There are finitely many strict dimension drops, so there is a finite order beyond which the subspaces stabilize to that intersection. The order itself is not bounded by $\dim W$: arbitrary stretches of equal successive kernels can precede a later drop.

For a concrete example, take $\rho_s = I/2 + s^m Z/4$ near $s = 0$ and $A = Z$. Then $\text{tr}(\rho_s Z) = s^m/2$. The first nonzero derivative occurs at the arbitrarily chosen order m , while the ambient Hermitian carrier has dimension four. This rules out a dimension-only upper bound on the detecting jet order and a stopping criterion based only on one consecutive equality. The valid eventual finite-determination result is retained; the stronger algorithmic claim requires an extra degree or differential-closure hypothesis.

All derivatives in this statement are taken in a specified coordinate chart. Higher ordinary derivatives of a scalar function are not automatically invariant tensors under arbitrary nonlinear coordinate changes. A coordinate-free formulation must use a jet bundle or an explicitly chosen connection. The finite linear solvability problem itself is unaffected by this distinction.

C.2 Positive witnesses and physical implementation

With evaluation map $\mathcal{E} : W \rightarrow \mathbb{R}^m$, the set of witnesses for target y is either empty or an affine coset $A_0 + \ker \mathcal{E}$. Requiring positivity intersects this affine set with the positive semidefinite cone. The result is closed and convex; it is not generally a face of that cone.

If one calibration state obeys $\rho_* \geq \lambda_* I$ with $\lambda_* > 0$ and every positive witness has expectation g_* in that state, then

$$\lambda_* \text{tr } A \leq \text{tr}(\rho_* A) = g_*.$$

Thus $\|A\|_\infty \leq \text{tr } A \leq g_*/\lambda_*$, and the positive-witness set is compact in finite dimension. A nonempty compact convex set has extreme points. A witness A is extreme exactly when no

nonzero $H \in \ker \mathcal{E}$ satisfies $A \pm H \geq 0$. These statements describe mathematical feasibility; an accessible implementation of a chosen witness requires an admitted instrument.

C.3 Approximate capacity of a finite record

Suppose M equiprobable historical labels are encoded into states $\rho_1, \dots, \rho_M$ on a Hilbert space of dimension D , and a POVM $\{E_i\}$ decodes them with average success at least $1 - \epsilon$. Since $\rho_i \leq I$,

$$1 - \epsilon \leq \frac{1}{M} \sum_i \text{tr}(E_i \rho_i) \leq \frac{1}{M} \sum_i \text{tr} E_i = \frac{D}{M}.$$

Hence

$$D \geq (1 - \epsilon)M.$$

For $M = 2^n$ independent binary histories, perfect retention gives $D \geq 2^n$ and approximate decoding gives the displayed weakened bound. This is a resource requirement for distinguishing historical alternatives, not a lower bound on consciousness. It also does not forbid repeatedly using one register when old independent histories are not required to remain available.

C.4 An exact probability countermodel and its repair condition

The following example delimits a tempting use of closure. Let $D_B \rho$ be diagonalization in a fixed record basis and

$$c(\rho) = \|\rho - D_B \rho\|_{HS}^2, \quad \tilde{\rho} = (1 - c)D_B \rho + cI/d.$$

For density matrices $0 \leq c \leq 1$. The assignment $p(E | \rho) = \text{tr}(E \tilde{\rho})$ is positive, normalized, additive on orthogonal effects, continuous, and equivariant under record-basis-preserving unitaries. It agrees with ordinary diagonal calibration and converges under the constant closure channel $T\rho = I/d$. Yet for

$$\rho = \begin{pmatrix} 3/4 & 1/4 \\ 1/4 & 1/4 \end{pmatrix}$$

it gives $p_0 = 23/32$, not $3/4$. Those stability and fixed-basis conditions alone therefore do not force Born weighting.

An additional affinity premise can remove this freedom for the fixed record measurement. Suppose $p_i(\rho)$ is affine on all density operators, lies in $[0, 1]$, and obeys $p_i(|j\rangle\langle j|) = \delta_{ij}$. Finite-dimensional duality gives an effect F_i with $p_i(\rho) = \text{tr}(F_i \rho)$. Its diagonal is fixed by calibration. Positivity implies $|(F_i)_{jk}|^2 \leq (F_i)_{jj}(F_i)_{kk}$, so every off-diagonal entry vanishes. Hence $F_i = |i\rangle\langle i|$ and $p_i(\rho) = \rho_{ii}$.

This is an exact conditional derivation from affinity and full state-domain positivity, not a derivation of those premises from contraction. The physical measurement completions retain their own probability-bearing preparation assumptions. The example is included to state the correct theorem boundary, not to deny the existence of well-defined quantum decision or measurement models.

C.5 Physical archive corruption and wave error

For a flat archive boundary Σ in full configuration space, the normal probability current of the internal-vector wave is $(\hbar/m) \text{Im}\langle \Psi, \partial_n \Psi \rangle$. Suppose the nominal held wave has zero normal

current pointwise on that surface, and the actual wave is $\Psi = \Psi_0 + \xi$. Surface norms integrate all remaining spatial variables, including the clock. The cross terms and Cauchy–Schwarz yield

$$\int_{\Sigma} |j_n(\Psi)| \leq \frac{\hbar}{m} \left(\|\Psi_0\|_{L^2(\Sigma)} \|\partial_n \xi\|_{L^2(\Sigma)} + \|\xi\|_{L^2(\Sigma)} \|\partial_n \Psi_0\|_{L^2(\Sigma)} + \|\xi\|_{L^2(\Sigma)} \|\partial_n \xi\|_{L^2(\Sigma)} \right).$$

Here the Sobolev norm differentiates the system/archive coordinates, with clock and other spectator coordinates treated as L^2 parameters. The trace inequality is correspondingly Hilbert-valued and integrated over those parameters; it does not differentiate the rapidly oscillating clock phase. Under a specified H^2 trace bound C_{Σ} , $\|\Psi_0\|_{H^2} \leq B$, and $\|\xi\|_{H^2} \leq \epsilon_2$, integration over an interval I gives

$$\int_I \int_{\Sigma} |j_n| \leq \frac{\hbar}{m} |I| C_{\Sigma}^2 \epsilon_2 (2B + \epsilon_2).$$

For the deterministic configuration-guidance flow $v = j/\rho$, the regular-surface crossing formula gives

$$\mathbb{E} N_{\Sigma}(I) = \int_I \int_{\Sigma} |j_n|, \quad P(N_{\Sigma}(I) \geq 1) \leq \mathbb{E} N_{\Sigma}(I),$$

under the stated flow regularity and the usual regular crossing hypotheses for the chosen surface. The estimate therefore bounds historical corruption that requires such a crossing. Equivariance alone would not suffice: a stationary equivariant diffusion can cross a boundary while the Schrödinger current is zero. Ordinary L^2 endpoint closeness likewise supplies no derivative or surface bound. The archive conclusion uses the specified guidance paths and the stronger physical norm together.

Appendix D

Exact finite matrices and reproducibility

D.1 Boundary data for the five O1 profiles

The following matrices, together with the stated basis conventions, determine every d, D, K and spectral projector in Chapter 12. They are reconstructed from the relation lists in the O1 source [39], not from a fit to the desired spectra. Vertices are ordered 0, 1, 2, 3. Higher-dimensional cells are ordered lexicographically after increasing-label orientation. Write $B_1 : C_1 \rightarrow C_0$ and $B_2 : C_2 \rightarrow C_1$ for their alternating-face boundary matrices.

In this grading,

$$d = \begin{pmatrix} 0 & B_1 & 0 \\ 0 & 0 & B_2 \\ 0 & 0 & 0 \end{pmatrix}, \quad D = d + d^*, \quad K = \text{diag}(B_1 B_1^*, B_1^* B_1 + B_2 B_2^*, B_2^* B_2).$$

Empty face blocks are omitted in seven-dimensional cases. The retained projection selects the two root vertices and the root edge 01.

V2_02

Edge basis: (01, 02, 03).

$$B_1 = \begin{pmatrix} -1 & -1 & -1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

V2_34

Edge basis: (01, 12, 13).

$$B_1 = \begin{pmatrix} -1 & 0 & 0 \\ 1 & -1 & -1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

V2_04

Edge basis: (01, 02, 03, 12, 13). Face basis: (012, 013).

$$B_1 = \begin{pmatrix} -1 & -1 & -1 & 0 & 0 \\ 1 & 0 & 0 & -1 & -1 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \end{pmatrix}, \quad B_2 = \begin{pmatrix} 1 & 1 \\ -1 & 0 \\ 0 & -1 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}$$

V2_11

Edge basis: (01, 02, 03, 12, 13). Face basis: (012, 013).

$$B_1 = \begin{pmatrix} -1 & -1 & -1 & 0 & 0 \\ 1 & 0 & 0 & -1 & -1 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \end{pmatrix}, \quad B_2 = \begin{pmatrix} 1 & 1 \\ -1 & 0 \\ 0 & -1 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}$$

V2_29

Edge basis: (01, 02, 03, 12, 13). Face basis: (012, 013).

$$B_1 = \begin{pmatrix} -1 & -1 & -1 & 0 & 0 \\ 1 & 0 & 0 & -1 & -1 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \end{pmatrix}, \quad B_2 = \begin{pmatrix} 1 & 1 \\ -1 & 0 \\ 0 & -1 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}$$

For distinct eigenvalues in $\Lambda = \text{Spec } K$, the exact projectors are

$$E_\lambda = \prod_{\mu \in \Lambda \setminus \{\lambda\}} \frac{K - \mu I}{\lambda - \mu}.$$

These polynomial formulas avoid numerical choices within degenerate eigenspaces. Applying them to the displayed boundary data yields the ranks in Table 12.1. Full d, D, K, P_R, E_λ arrays are included in the companion machine-readable results.

D.2 Signed relabelling

A vertex permutation transports every oriented simplex with the sign of the induced reordering of its vertices. Let U be the resulting signed permutation on the chain carrier. Alternating face deletion gives $d'U = Ud$. Orthogonality then gives $D'U = UD$, $K'U = UK$, and $P'_R U = UP_R$ when root provenance is transported as well. Hence spectral projectors and cross-block ranks intertwine. Reassigning the retained root after a permutation would be a different model, not a relabelling test.

The verifier checks all 24 vertex permutations for each of the five profiles. It also checks $d^2 = 0$, orthogonality and completeness of the projectors, the exact characteristic spectra, and the current-support identities. These 120 signed transports certify the finite arithmetic in the supplied basis convention; they do not establish an unstated physical actuator law.

D.3 The recurrent two-bit calculation

With states ordered 00, 01, 10, 11, the transition matrix in Equation (19.2) is reconstructed by multiplying the two independent noise probabilities. It is stochastic and nonsingular. Direct row differences give the delayed-observation matrix in Equation (19.3). If the first future state is observed exactly and all later outputs are produced by common Markov processing of it, the complete delayed-path TV equals the TV between those first-state laws: projection gives the lower bound and common processing gives the upper bound.

The verifier enumerates delayed passive paths of lengths one through three and checks that equality exactly. It separately checks that allowing exact present-state reports distinguishes every two current states with TV one. Thus both matrices are correct for their explicitly different experiment grammars; neither is silently substituted for the other.

The noisy update has an injective finite lift. On (A, B, n_A, n_B) use

$$(A, B, n_A, n_B) \mapsto (B \oplus n_A, A \oplus n_B, n_A, n_B).$$

The retained noise bits recover the old A, B , so the map is a permutation on 16 basis states. Fresh noise and subsequent archive receivers must be provided for each finite step. This implements the stated transition law from its prepared noise ensemble without inserting a consciousness-dependent force.

D.4 What the executable package checks

The companion script performs exact rational matrix calculations for the five profiles and the recurrent vessel, the non-Born probability counterexample, the finite null/continuation benchmark, and the delayed jet example. It also performs 100 reproducibly seeded numerical stress tests of the witness perturbation inequality. Those finite numerical tests supplement the analytic proof; they do not establish a universal bound by sampling.

The package also retains the earlier finite continuation, admissibility, weak-coupling, product and spectral illustrations with their original scope. The actual SPC-2 construction has separate exact checks. A four-state intrinsic model with hidden duplicate states yields two predictive classes; its quotient reproduces all 723 admitted adaptive trees through depth three from each of four initial states, giving 2,892 exact transcript-law comparisons. The five-state counterexample detects failed predictive-fibre congruence rather than silently treating trace equivalence as actual-state closure.

Further checks cover 24 relabellings of a native model, a controlled product, delayed revelation in a three-bit register, and joint-parent factorization over all 15 nonempty subsets of a four-component correlation model. The hostile constructions include horizon loss, joint-marginal blindness, unmodelled external return, forbidden action concatenation, report nonseparation, incompatible regimes, decomposition and time-scale changes, incomplete recurrence catalogues, and the three-factorization symmetry obstruction. The process-lineage examples include growth, shrinkage, disjoint endpoint material supports after gradual replacement, copying, splitting, merging, repair, true qualification gaps and censored observation windows.

The integration checks verify the weak native-mechanism return formula at exact rational couplings, including the intermediate $g = 1/2$ case in which two present states share a native future profile. Qualification is evaluated in the actual protocol/resource type, so exhausted types cannot borrow live-type return witnesses. The physical counterchecks distinguish ideal from spatially noisy status records, include wrong-plateau feedback tails, and show why equivariance alone does not protect archives.

These are finite mathematical checks and illustrative countermodels, not measurements of awareness or a proof-assistant verification of every theorem. The scripts expose the premises they assume, including the physical mechanism and route certificates. No neural recording, subjective state, or hidden model activation was measured in these calculations. Assessment of the psychophysical laws requires the independently constrained evidence discussed in Chapter 23.

Appendix E

Reconstruction, symmetry, and the meaning of a source

E.1 A selected representative is not an inverse

If $p : S \rightarrow X$ is surjective, a section $\sigma : X \rightarrow S$ satisfies $p\sigma = \text{id}_X$. It also satisfies $\sigma p = \text{id}_S$ only if p is injective. Indeed, if $p(s) = p(t)$, then $s = \sigma p(s) = \sigma p(t) = t$. Choosing one state per fibre therefore does not reconstruct which of the original fibre states actually occurred.

This distinction matters for both source ontology and phenomenal interpretation. A canonical-looking selected representative can provide useful coordinates without proving that the source has only that representative. The choice can also depend on structures not present in the readout. Such dependence must be retained rather than disguised as an intrinsic property of the quotient.

E.2 Equivariant sections

Suppose a group G acts on S and X and p is equivariant. Let G_x be the stabilizer of x and $F_x = p^{-1}(x)$.

Proposition E.1 (Equivariant representative criterion). *A G -equivariant section exists if and only if, for each orbit representative x , the stabilizer G_x fixes a point of F_x , with the requisite orbitwise choices available.*

Proof. If σ is equivariant and $g \in G_x$, then $g\sigma(x) = \sigma(gx) = \sigma(x)$, proving necessity. Conversely choose $s_x \in F_x$ fixed by G_x for each orbit representative and define $\sigma(gx) = gs_x$. If $gx = g'x$, then $g'^{-1}g \in G_x$ fixes s_x , so the definition is independent of the chosen representative. It is a section and equivariant by construction. $\square$

A two-point fibre on which the stabilizer swaps its points has no deterministic equivariant representative, although its uniform probability distribution is invariant. An invariant ensemble is not a recovered actual state. These source/readout facts are developed in the source non-equivalence manuscript [41].

E.3 Topological and physical qualifications

A continuous section can fail for reasons distinct from equivariance. The Hopf fibration $S^3 \rightarrow S^2$ has circle fibres and admits no continuous global section. A section of this principal

circle bundle would trivialize it to $S^2 \times S^1$, but that product has nontrivial fundamental group whereas S^3 is simply connected. This is a topological obstruction, separate from a requirement to commute with a specified group action.

Neither obstruction proves that an arbitrary gauge coordinate is a hidden physical fact. If the theory declares the fibre distinction redundant, quotienting it removes a descriptive choice rather than observable source content. The reduced source must be identified before a non-reconstruction claim is interpreted physically. Likewise, a theorem about one noninjective aperture does not prove that every possible family of apertures fails jointly to reconstruct a nominated finite domain.

These qualifications strengthen the philosophical use of source/readout mathematics. They permit precise limitations on a perspective without converting ignorance into positive knowledge of an otherwise unspecified Absolute. They also prevent a selected phenomenal coordinate system from being confused with an independently demonstrated ontology.

Appendix F

Notation, claim status, and construction dependencies

F.1 Principal notation

Symbol	Meaning
$\mathbf{U}$	Reality prior to operational distinction; an ontological notation, not a mathematical carrier.
Ω_{src}	Legacy ontological notation, interpreted through $\mathbf{U}$; not the working source S .
S, T	Structure-side working source and record-side readout; T is called X in the descent theorem. Local horizons and kernels retain their separately defined symbols.
$S, p : S \rightarrow X$	Nominated complete realization domain and its readout.
$\mathcal{O}\Delta\partial$	A supplied closure composition; its operators and fixed point are not $\mathbf{U}$.
H, J_{nm}, w_n	Quantum Hamiltonian, directed Hamiltonian current, and coherent sector weight.
d, D, K	Finite boundary differential, self-adjoint incidence completion, and squared response operator.
P_R, P_F	Source-provenance retained and complementary projections.
$h(K), \gamma, W_t$	Real spectral coherent generator, damping coefficient, and combined coherent–dissipative response.
$\mathcal{V}, \mathcal{B}, \mathcal{U}$	Realized vessel, lineaged components, and admitted controls.
$G^{\text{all}}, G^{\text{int}}, C$	Full dependence graph, internally typed dependence graph, and a candidate strongly connected core.
$M_{a,y}$	Physical instrument matrix for an action and recorded output.
$X_C, Z_r, Z_\infty, \mathfrak{P}_C$	Actual intrinsic carrier, finite-horizon quotient, all-future quotient, and complete native predictive instrument object.
Φ_C, Ψ_C	Phenomenal structural copy and constitutive identification with the predictive object.
$Q_k, \hat{Q}_k, \tilde{Q}_k$	Physical, predictive-class, and induced phenomenal report kernels at fixed external context.
Σ_V	Learned/retained scaffolding; not a separate source-awareness substance.

F.2 Claim-status register

Claim	Status and dependence
Source/readout descent and target completion	Exact theorems on supplied domains and maps.
Controlled reduced dynamics	Exact criterion on a supplied transition law, not a selection of that law.
Observer-independent massive model	Conditional physical construction using guidance, complete initial equilibrium, and the specified writer/control laws.
Pilot-to-Bell path limit	Conditional physical comparison on a fixed graph/horizon with the named contact, preparation, export, and recombination laws.
Physical copying and finite retention	Exact or quantitatively controlled under the displayed apparatus and resource assumptions.
SRC witness classification	Exact finite linear results; positivity, retention, and physical accessibility are additional requirements.
CRR access and regulation	Typed functional constructions; phenomenal unity requires additional laws.
O1/RSM classification	Exact finite calculation from the stated provisional source-response data.
CSCF spectral responses	Exact common-carrier operator construction for $K, h(K), \gamma$; the same holomorphic family gives both boundaries when $h(\lambda) = \lambda$.
Aperture error and witness robustness	Exact/conditional mathematical bounds for nominated targets and calibrated preparations.
R^* realization and interface typing	Constitutive physical input; not universally selected by the source/readout theorem.
A0 primitive awareness	Ontological hypothesis: awareness as the knowing aspect of U; no extra physical dynamics.
A1 recurrent-core admission	Constitutive psychophysical law; actual-vessel certification and its empirical lower boundary remain unvalidated.
A2 predictive-profile completeness	Constitutive structuralist psychophysical law, including its gauge convention.
A3 temporal episode identity	Nonbranching process-provenance continuation allowing support turnover; genuine qualification gaps end episodes.
Biological conscious/generative modes	Interpretive account of scene organization; not a new admission condition or a scene gate in the finite theorem.
Cortical correlates of coupling	Candidate mammalian implementation; network or band measures neither modify A1 nor establish A0 or A2.

Claim	Status and dependence
Scene-less intervals	No current scene or localized experiential subject. Physical/person continuity can persist; a qualification gap ends the episode and renewed qualification begins a new one.
Finite constitutive completion	Exact conditional construction for certified finite intrinsic carriers, rational native instruments, complete return catalogues, predictive-fibre congruence, grounded reports and explicit process provenance under R^* and A0–A3.
Empirical conservativity	Exact equality of physical transcript laws when the aspect extension changes no physical kernel.
Contemplative correspondences	Philosophical interpretation of reported experience, not a field measurement.
Human/animal/AI realization	Open empirical and theoretical application; no validated universal criterion claimed.
All-finite native horizons	Exact compatible construction on the certified finite intrinsic domain; unrestricted continuum/history domains need additional hypotheses.

F.3 Dependencies that cannot be omitted

The finite completion theorem depends on R^* before it depends on a psychophysical axiom. Removing the typed boundaries, the physical preparation, the protocol grammar, or the lineage rule can make the assignment underdetermined. Removing A1 removes experiential admission; removing A2 removes the claimed completeness of relational content; removing A3 removes the stipulated identity rule. None is supplied by the word “canonical.”

The mathematical results of the earlier parts remain meaningful even when A0–A3 are rejected. A reader may accept the predictive quotient, the error theorem, and the material record constructions while declining their experiential interpretation. Conversely, accepting the philosophy does not excuse a failed physical estimate. This separation is what permits the monograph to serve as both a constructive proposal and a technically auditable theory.

F.4 Reproducibility and provenance

The accompanying coverage and verification materials map the technical claims to their chapters and record the mathematical and computational checks. The interpretive discussion is not presented as a verified biological model. The executable scripts and results expose the finite computations used here. These materials enable reproduction and targeted review of the stated claims.

Established methods—quotient descent, finite instrument representations, fixed-point estimates, reversible computation, and Bell/configuration dynamics—retain their intellectual provenance. No historical-priority claim is made for those ingredients. The proposed synthesis lies in how the source/readout and realization discipline organizes the consciousness question, and in the explicit conditional constitution that can now be examined as a whole.

Bibliography

- [1] Larissa Albantakis, Leonardo Barbosa, Graham Findlay, Matteo Grasso, Andrew M. Haun, William Marshall, William G. P. Mayner, Alireza Zaeemzadeh, Melanie Boly, Bjørn E. Juel, Shuntaro Sasai, Keiko Fujii, Isaac David, Jeremiah Hendren, Jonathan P. Lang, and Giulio Tononi. Integrated information theory (IIT) 4.0: Formulating the properties of phenomenal existence in physical terms. *PLOS Computational Biology*, 19(10):e1011465, 2023. doi:[10.1371/journal.pcbi.1011465](https://doi.org/10.1371/journal.pcbi.1011465).
- [2] Tim Bayne, Joel Frohlich, Rhodri Cusack, Julia Moser, and Lorina Naci. Consciousness in the cradle: On the emergence of infant experience. *Trends in Cognitive Sciences*, 27(12):1135–1149, 2023. doi:[10.1016/j.tics.2023.08.018](https://doi.org/10.1016/j.tics.2023.08.018).
- [3] John S. Bell. Beables for quantum field theory. In *Speakable and Unsayable in Quantum Mechanics*, pages 173–180. Cambridge University Press, Cambridge, 1987. Reprinted account of the Bell beable construction.
- [4] Charles H. Bennett. Logical reversibility of computation. *IBM Journal of Research and Development*, 17(6):525–532, 1973. doi:[10.1147/rd.176.0525](https://doi.org/10.1147/rd.176.0525).
- [5] Jonathan Birch, Alexandra K. Schnell, and Nicola S. Clayton. Dimensions of animal consciousness. *Trends in Cognitive Sciences*, 24(10):789–801, 2020. doi:[10.1016/j.tics.2020.07.007](https://doi.org/10.1016/j.tics.2020.07.007).
- [6] Yelena G. Bodien, Judith Allanson, Paolo Cardone, et al. Cognitive motor dissociation in disorders of consciousness. *New England Journal of Medicine*, 391(7):598–608, 2024. doi:[10.1056/NEJMoa2400645](https://doi.org/10.1056/NEJMoa2400645).
- [7] David Bohm. A suggested interpretation of the quantum theory in terms of ‘hidden’ variables. I. *Physical Review*, 85(2):166–179, 1952. doi:[10.1103/PhysRev.85.166](https://doi.org/10.1103/PhysRev.85.166).
- [8] Patrick Butlin, Robert Long, Eric Elmoznino, Yoshua Bengio, Jonathan Birch, Axel Constant, George Deane, Stephen M. Fleming, Chris Frith, Xu Ji, Ryota Kanai, Colin Klein, Grace Lindsay, Matthias Michel, Liad Mudrik, Megan A. K. Peters, Eric Schwitzgebel, Jonathan Simon, and Rufin VanRullen. Consciousness in artificial intelligence: Insights from the science of consciousness, 2023. arXiv:2308.08708, <https://arxiv.org/abs/2308.08708>.
- [9] Robin L. Carhart-Harris and Karl J. Friston. REBUS and the anarchic brain: Toward a unified model of the brain action of psychedelics. *Pharmacological Reviews*, 71(3):316–344, 2019. doi:[10.1124/pr.118.017160](https://doi.org/10.1124/pr.118.017160).
- [10] David J. Chalmers. Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3):200–219, 1995. Author-hosted text: <https://consc.net/papers/facing.html>.
- [11] Matthias Christandl, Nilanjana Datta, Artur Ekert, and Andrew J. Landahl. Perfect state transfer in quantum spin networks. *Physical Review Letters*, 92(18):187902, 2004. doi:[10.1103/PhysRevLett.92.187902](https://doi.org/10.1103/PhysRevLett.92.187902).
- [12] Cogitate Consortium, Oscar Ferrante, Urszula Gorska-Klimowska, Simon Henin, et al. Adversarial testing of global neuronal workspace and integrated information theories of consciousness. *Nature*, 642(8066):133–142, 2025. doi:[10.1038/s41586-025-08888-1](https://doi.org/10.1038/s41586-025-08888-1).
- [13] Stanislas Dehaene and Jean-Pierre Changeux. Experimental and theoretical approaches to conscious processing. *Neuron*, 70(2):200–227, 2011. doi:[10.1016/j.neuron.2011.03.018](https://doi.org/10.1016/j.neuron.2011.03.018).
- [14] Sven Dorkenwald, Arie Matsliah, Amy R. Sterling, et al. Neuronal wiring diagram of an adult brain. *Nature*, 634(8032):124–138, 2024. doi:[10.1038/s41586-024-07558-y](https://doi.org/10.1038/s41586-024-07558-y).

- [15] Detlef Dürr, Sheldon Goldstein, Roderich Tumulka, and Nino Zanghì. Bell-type quantum field theories. *Journal of Physics A: Mathematical and General*, 38(4):R1–R43, 2005. doi:10.1088/0305-4470/38/4/R01.
- [16] Detlef Dürr, Sheldon Goldstein, and Nino Zanghì. Quantum equilibrium and the origin of absolute uncertainty. *Journal of Statistical Physics*, 67:843–907, 1992. doi:10.1007/BF01049004.
- [17] Richard P. Feynman. Quantum mechanical computers. *Foundations of Physics*, 16(6):507–531, 1986. doi:10.1007/BF01886518.
- [18] Alex Gamma and Thomas Metzinger. The Minimal Phenomenal Experience questionnaire (MPE-92M): Towards a phenomenological profile of ‘pure awareness’ experiences in meditators. *PLOS ONE*, 16(7):e0253694, 2021. doi:10.1371/journal.pone.0253694. Data-availability correction: doi:10.1371/journal.pone.0314290.
- [19] Stuart Hameroff. The “conscious pilot”—dendritic synchrony moves through the brain to mediate consciousness. *Journal of Biological Physics*, 36(1):71–93, 2010. doi:10.1007/s10867-009-9148-x.
- [20] Erik P. Hoel, Larissa Albantakis, and Giulio Tononi. Quantifying causal emergence shows that macro can beat micro. *Proceedings of the National Academy of Sciences*, 110(49):19790–19795, 2013. doi:10.1073/pnas.1314922110.
- [21] Olga A. Imas, Kristina M. Ropella, B. Douglas Ward, James D. Wood, and Anthony G. Hudetz. Volatile anesthetics disrupt frontal-posterior recurrent information transfer at gamma frequencies in rat. *Neuroscience Letters*, 387(3):145–150, 2005. doi:10.1016/j.neulet.2005.06.018.
- [22] Johannes Kleiner and Tim Ludwig. What is a mathematical structure of conscious experience?, 2023. arXiv:2301.11812, <https://arxiv.org/abs/2301.11812>.
- [23] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961. doi:10.1147/rd.53.0183.
- [24] UnCheol Lee, SeungWoo Ku, GyuJeong Noh, SeungHye Baek, ByungMoon Choi, and George A. Mashour. Disruption of frontal-parietal communication by ketamine, propofol, and sevoflurane. *Anesthesiology*, 118(6):1264–1275, 2013. doi:10.1097/ALN.0b013e31829103f5.
- [25] Laura D. Lewis, Veronica S. Weiner, Eran A. Mukamel, Jacob A. Donoghue, Emad N. Eskandar, Joseph R. Madsen, William S. Anderson, Leigh R. Hochberg, Sydney S. Cash, Emery N. Brown, and Patrick L. Purdon. Rapid fragmentation of neuronal networks at the onset of propofol-induced unconsciousness. *Proceedings of the National Academy of Sciences*, 109(49):E3377–E3386, 2012. doi:10.1073/pnas.1210907109.
- [26] Michael L. Littman, Richard S. Sutton, and Satinder Singh. Predictive representations of state. In *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2002. Conference held in 2001; author list follows the original paper PDF. https://proceedings.neurips.cc/paper_files/paper/2001/file/1e4d36177d71bbb3558e43af9577d70e-Paper.pdf.
- [27] Antoine Lutz, Jean-Philippe Lachaux, Jacques Martinerie, and Francisco J. Varela. Guiding the study of brain dynamics by using first-person data: Synchrony patterns correlate with ongoing conscious states during a simple visual task. *Proceedings of the National Academy of Sciences*, 99(3):1586–1591, 2002. doi:10.1073/pnas.032658199.
- [28] Katherine A. MacLean, Matthew W. Johnson, Chad J. Reissig, Thomas E. Prisinzano, and Roland R. Griffiths. Dose-related effects of salvinorin A in humans: Dissociative, hallucinogenic, and memory effects. *Psychopharmacology*, 226(2):381–392, 2013. doi:10.1007/s00213-012-2912-9.
- [29] George A. Mashour. Anesthesia and the neurobiology of consciousness. *Neuron*, 112(10):1553–1567, 2024. doi:10.1016/j.neuron.2024.03.002.
- [30] Yair Pinto, David A. Neville, Marte Otten, Paul M. Corballis, Victor A. F. Lamme, Edward H. F. de Haan, Nicoletta Foschi, and Mara Fabri. Split brain: Divided perception but undivided consciousness. *Brain*, 140(5):1231–1237, 2017. doi:10.1093/brain/aww358.
- [31] Jeremy Rodgers. Why agency cannot be purely quantum under unconstrained unitarity: A closure-theoretic completion of the unitary no-go result. Unpublished programme manuscript, also

- titled Agency Beyond Pure Unitarity. Cited for its operational formulation, not its unqualified probability or outcome claims, 2025.
- [32] Jeremy Rodgers. The canonical spectral–curvature field: Closed forms, holomorphic evolution, and Euclidean–unitary spectral continuation. Unpublished working manuscript, strengthened version, August 2026.
- [33] Jeremy Rodgers. The consciousness field: A unified mathematical framework from fixed-point field theory to phenomenal bridge. Version 5.5, unpublished programme manuscript. Only the qualified local and representation results are used, March 2026.
- [34] Jeremy Rodgers. A deterministic pilot medium: Bell path selection and autonomous material records—a constitutive completion with a controlled physical-time limit. Zenodo, September 2026. Version 2, 15 September 2026. doi:10.5281/zenodo.22774634.
- [35] Jeremy Rodgers. A massive configuration completion of the quantum measurement programme: Event timing, semibounded material records, retained resources, and an autonomous finite-horizon realization. Zenodo, September 2026. Version 2, 15 September 2026. doi:10.5281/zenodo.22774739.
- [36] Jeremy Rodgers. Self-referential record closure: A structural criterion for witness-consciousness. Unpublished programme manuscript. Witness results used with the scope qualifications stated in this monograph, April 2026.
- [37] Jeremy Rodgers. Shadow theory and quantum measurement: Source dynamics, event laws, and physical records—two constitutive completions. Zenodo, September 2026. Version 2, 15 September 2026. doi:10.5281/zenodo.22774584.
- [38] Jeremy Rodgers. Shadow theory: Framework. Author’s website, 2026. <https://www.everythingequation.com/framework>. Accessed 18 September 2026.
- [39] Jeremy Rodgers. Shadow theory O1 constructive discovery: Run 5. Unpublished 16-page research report. Relation and response tables on pp. 5–9; finite matrices independently reconstructed here, 2026.
- [40] Jeremy Rodgers. Shadow theory: Reality atlas. Author’s website, 2026. <https://www.everythingequation.com/atlas>. Accessed 18 September 2026.
- [41] Jeremy Rodgers. Source–readout non-equivalence: Descent and equivariant reconstruction obstructions. Zenodo, July 2026. Canonical edition dated 15 July 2026. doi:10.5281/zenodo.21369967.
- [42] Jeremy Rodgers. Typed conscious record realization: CRR final theorem package. Version 1.0, unpublished programme manuscript, June 2026.
- [43] M. D. Sanders, E. K. Warrington, J. Marshall, and L. Weiskrantz. ‘blindsight’: Vision in a field defect. *The Lancet*, 303(7860):707–708, 1974. doi:10.1016/S0140-6736(74)92907-9.
- [44] Simone Sarasso, Melanie Boly, Martino Napolitani, Olivia Gosseries, Vanessa Charland-Verville, Silvia Casarotto, Mario Rosanova, Adenauer Girardi Casali, Jean-François Brichant, Pierre Boveroux, Steffen Rex, Giulio Tononi, Steven Laureys, and Marcello Massimini. Consciousness and complexity during unresponsiveness induced by propofol, xenon, and ketamine. *Current Biology*, 25(23):3099–3105, 2015. doi:10.1016/j.cub.2015.10.014.
- [45] Annalotta Scheinin, Oskari Kantonen, Michael Alkire, Jaakko Långsjö, Roosa E. Kallionpää, Kaike Kaisti, Linda Radek, Jarkko Johansson, Nils Sandman, Mikko Nyman, Mika Scheinin, Tero Vahlberg, Antti Revonsuo, Katja Valli, and Harry Scheinin. Foundations of human consciousness: Imaging the twilight zone. *Journal of Neuroscience*, 41(8):1769–1778, 2021. doi:10.1523/JNEUROSCI.0775-20.2020.
- [46] Miguel Segundo-Ortín and Paco Calvo. Consciousness and cognition in plants. *WIREs Cognitive Science*, 13(2):e1578, 2022. doi:10.1002/wcs.1578; first published online in 2021.
- [47] Anil K. Seth. Conscious artificial intelligence and biological naturalism. *Behavioral and Brain Sciences*, 49:e315, 2026. First published online 21 April 2025; doi:10.1017/S0140525X25000032.
- [48] Cosma Rohilla Shalizi and James P. Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity. *Journal of Statistical Physics*, 104(3–4):817–879, 2001. doi:10.1023/A:1010388907793; arXiv:cond-mat/9907176.

- [49] Philip K. Shiu, Gabriella R. Sterne, Nico Spiller, et al. A *Drosophila* computational brain model reveals sensorimotor processing. *Nature*, 634(8032):210–219, 2024. doi:[10.1038/s41586-024-07763-9](https://doi.org/10.1038/s41586-024-07763-9).
- [50] Francesca Siclari, Benjamin Baird, Lampros Perogamvros, Giulio Bernardi, Joshua J. LaRocque, Brady Riedner, Melanie Boly, Bradley R. Postle, and Giulio Tononi. The neural correlates of dreaming. *Nature Neuroscience*, 20(6):872–878, 2017. doi:[10.1038/nn.4545](https://doi.org/10.1038/nn.4545).
- [51] Sri Nisargadatta Maharaj. *I Am That: Talks with Sri Nisargadatta Maharaj*. Chetana, Bombay, 1973. Translated by Maurice Frydman; revised and edited by Sudhakar S. Dikshit. Dialogue numbers are the principal locators; the supplied electronic edition carries the July 1981 editorial note.
- [52] Lincoln Taiz, Daniel Alkon, Andreas Draguhn, Angus Murphy, Michael Blatt, Chris Hawes, Gerhard Thiel, and David G. Robinson. Plants neither possess nor require consciousness. *Trends in Plant Science*, 24(8):677–687, 2019. doi:[10.1016/j.tplants.2019.05.008](https://doi.org/10.1016/j.tplants.2019.05.008).
- [53] Francisco J. Varela. Neurophenomenology: A methodological remedy for the hard problem. *Journal of Consciousness Studies*, 3(4):330–349, 1996.
- [54] Lukas J. Volz, Steven A. Hillyard, Michael B. Miller, and Michael S. Gazzaniga. Unifying control over the body: Consciousness and cross-cueing in split-brain patients. *Brain*, 141(3):e15, 2018. Letter to the editor. doi:[10.1093/brain/awx359](https://doi.org/10.1093/brain/awx359).