

FLOW SHAPING PRINCIPLE: CORRESPONDENCE-CALIBRATED PATHWAY-SPECIFIC UPDATING

KAI WANG

Independent Researcher, Australia

Correspondence: wangkai@connection-theory.org

Version 3.7.0 | August 2026

Adaptive systems can assign persistent change to configurations through which consequential episodes pass. This paper formulates the Flow Shaping Principle (FSP) as a falsifiable claim about veridical consequence-configuration correspondence. FSP concerns active updating only. Foundational meanings are imported from Connection Theory and exhaustively registered in a version-bound appendix. For a pre-specified boundary, resolution, property set, and primary partition Π , the system must first demonstrate Π -specific updating. $A_corr(\Pi)$ is measured from update-available substrate through a pre-registered readout capacity-matched to, but distinct from, the system's attribution computation. A separate diagnostic reads the system-generated attribution signal. $S_corr(\Pi)$ measures persistent change in the correct configuration aligned with consequence-warranted direction or content; localization alone is insufficient. FSP predicts that A_corr prospectively predicts held-out S_corr and that calibration transfers across two pre-specified correspondence-degradation classes. The zero-information boundary and information upper bound follow from data processing; they are analytic constraints, not claimed discoveries or empirical wagers. The primary canonical target is an instrumentable artificial adaptive system; neural and human organizations remain candidate extensions. Entry conditions, independent measurements, controls, and failure criteria are specified, but no complete canonical test has been reported.

Keywords: adaptive systems; complex systems; continual learning; credit assignment; pathway-specific updating

1. Introduction and Scope

1.1 The Explanatory Problem

An adaptive organization encounters each new episode through an organization already shaped by earlier episodes. Neural circuits route activity through existing pathways, continual-learning systems update parameters through existing representations and optimization structures, and institutions process events through established roles and procedures. Yet observing that an episode was followed by organizational change does not establish that the change was assigned to the configuration that carried the episode, nor that the direction of change corresponded to the consequence produced by that episode.

This distinction is central. An update may be global rather than pathway-specific. It may be localized to the correct pathway but carry the wrong content. It may also be driven by an unavailable, delayed, corrupted, or misassigned consequence. A general organizational claim must therefore distinguish three relations: which configuration carried the episode, which consequence belongs to that episode, and whether the persistent change in that configuration is aligned with the consequence that should guide it.

Mechanism-specific theories remain indispensable. Neuroscience studies synaptic and circuit plasticity [1]; artificial intelligence studies credit assignment, optimization, replay, and continual learning [9]; organizational research studies incentives, authority, routines, and review. FSP does not replace these mechanisms. It asks whether a narrow correspondence-calibration relation can be measured and defeated under controlled intervention.

1.2 Imported Concepts and Ownership Boundaries

Connection Theory v4.6 [12], archived under version DOI 10.5281/zenodo.21714702, owns the foundational meanings imported here. Appendix A is the sole and exhaustive register of those imported meanings, their bound content, and their role in the present paper; no other meaning is imported from that work. The phrase *operating adaptive organization* is compositional shorthand for an organization that is currently operating and is adaptive in the respective meanings registered in Appendix A, and introduces no additional theoretical primitive. The claims of the present paper remain indexed to those registered meanings. Dependency-governance consequences of later changes to Connection Theory are specified exclusively in Appendix A.5.

Three logical roles must be separated. First, the selected configuration must be update-capable; this is a methodological gate rather than an empirical discovery. Second, a consequence or retained proxy must be causally available to the organization; this is the feedback-availability boundary inherited from Connection Theory. Third, the update-available substrate must preserve veridical correspondence between the consequence and the configuration to which the consequence belongs. The zero-information version of the third condition is an attribution constraint, is not claimed as a novelty of FSP, and is not part of the empirical claim evaluated here. FSP's empirical burden begins only with calibrated, prospective, cross-manipulation prediction.

FSP addresses active updating only. FSP does not model retained capability, the persistence or degradation of capability-relevant persistence state during inactivity, Organizational Continuity across interruption, current applicability, or the mechanisms of inheritance, interpretation, and reconstruction across boundaries or carriers. Evolutionary Relativity v3.6.1 [13], archived under version DOI 10.5281/zenodo.21715199, examines accumulated effective organizational change across observational timescales. No definition, prediction, or failure criterion in the present paper is imported from or depends on that framework.

Connection Theory's update-non-commutativity prediction is referenced only to delimit ownership of causal order effects. It is neither imported into nor tested by the present FSP framework, as recorded in Appendix A.4.

1.3 Bounded Scope and Present Status

FSP applies only to consequence-sensitive adaptive organizations. It does not cover every physical case in which propagation deforms a pathway. The paper therefore

makes no evidential use of river channels, material wear, or other reciprocal physical couplings that lack an independently specified consequence-attribution process.

The current primary test domain is artificial adaptive systems, where internal states, replay indices, eligibility structures, and routed updates can be measured directly. Neural systems are a candidate extension. Human organizations are a candidate transfer domain pending a feasible independent readout of correspondence in the state actually available to the updating process. Mentioning a domain does not increase validated scope.

The proposition, entry conditions, measurement constraints, and prospective failure criteria have been specified, but no complete canonical evaluation comprising Manipulation Check 1a and Predictions 1b–1c has yet been reported. This is a pre-validation theoretical status, not empirical support. Domain scope and manipulation coverage are separate evidential dimensions: neither is established by argument, by success in a single domain, or by success on a single transfer pair.

Complex-systems perspective. FSP treats an adaptive organization as a history-dependent networked system whose configurations route consequential episodes and whose persistent changes depend on feedback availability and consequence-configuration correspondence. Its tools are a pre-specified organizational partition, controlled correspondence degradation, matched readouts at distinct stages and held-out cross-manipulation transfer. Its contribution to complex-systems research is a falsifiable relation between information available at an attribution stage and later pathway-specific organizational change, with failure criteria that distinguish global updating, misattribution and downstream interference.

2. The Flow Shaping Principle

2.1 Operational Frame and Entry Condition

Every confirmatory test is indexed to a frame $F = (B, R, P, \Pi)$. B specifies the system boundary; R specifies observational resolution; P specifies the property or update outcome under study; and Π is the single primary partition used to individuate flow-carrying configurations. Π must be fixed before primary outcomes are examined. Additional partitions may be exploratory or multiplicity-controlled, but they cannot rescue failure under the primary Π .

Before FSP is tested, an independent entry phase must establish that the system performs Π -specific updating in the selected frame. Under a high-correspondence condition, the pre-specified correspondence-correct update measure must exceed a no-correspondence baseline by a pre-registered margin. This requirement prevents a system that performs only global, undifferentiated, or provenance-insensitive updating from being treated as a failed instance of a pathway-specific principle.

High measured correspondence with negligible correspondence-correct change therefore does not, by itself, falsify FSP. It means that the entry condition for Π -specific updating was not established in that system and frame. This exclusion must be decided

from the independent entry phase, not after the main cross-manipulation result is known.

2.2 Veridical Correspondence and Correspondence-Correct Shaping

Let an episode e traverse a configuration cell C_i in Π and produce or become associated with consequence y_e . The ground-truth correspondence specifies that y_e belongs to C_i for the purpose of the pre-specified update. Pair permutation can preserve the identity of C_i , the identity of y_e , and both marginal distributions while destroying this veridical correspondence.

$A_{corr}(\Pi)$ denotes the amount of ground-truth consequence-configuration correspondence recoverable from the substrate state available to the update process by a pre-registered readout whose functional class and input access are capacity-matched to, but not identical with, the attribution computation performed by that process. It is not the information recoverable by an unconstrained external observer. A decoder that can invert a systematic reassignment unavailable to the system's update rule would measure information present in the substrate, not correspondence available to updating, and would therefore overestimate A_{corr} . Every readout is scored against the true pairing: a legible trace pointing to the wrong consequence-configuration pair has low correspondence fidelity even when configuration identity remains fully decodable. An update-identical readout of the attribution signal already produced by the system is a separate diagnostic quantity; it must not replace the capacity-matched primary A_{corr} because doing so moves the test boundary downstream from attribution to consolidation and interference.

$S_{corr}(\Pi)$ denotes persistent correspondence-correct shaping. The measure must require both correct location and correct update direction or content. A large update localized to C_i does not count as positive S_{corr} if it reflects the consequence belonging to C_j . Depending on the system, S_{corr} may be operationalized as a signed projection of the update in C_i onto a pre-specified consequence-consistent direction, a correct-content update contrast net of wrong-content and matched non-target updates, or another single pre-registered measure with the same correspondence requirement.

Primary A_{corr} must be read from the association-carrying substrate before or during attribution through the capacity-matched-but-not-identical decoder; the update-identical diagnostic is read from the attribution signal produced by the system; and S_{corr} is measured from persistent organizational change afterward. None may be defined from another. Both decoder classes, their permitted inputs, and the justification for their relation to the actual update process must be fixed independently of S_{corr} . This temporal, functional, and measurement separation is necessary for prospective prediction and for distinguishing attribution-stage loss from downstream consolidation or interference (Figure 1).




Figure 1. Temporal, functional, and measurement separation of the three quantities. Primary $A_{corr}(II)$ is read from the association-carrying substrate before or during attribution, through a decoder capacity-matched to but not identical with the attribution computation performed by the update rule. The update-identical diagnostic is read from the attribution signal the system has already produced. $S_{corr}(II)$ is read from persistent organizational change and requires both correct location and correct update direction or content. Both decoder classes, their permitted inputs, training splits, chance levels, and the justification for their relation to the actual update process are fixed independently of S_{corr} . The two brackets mark stages in which predictive loss may arise. Comparison of the paired curves can localize differential predictive loss before versus after the system-generated attribution signal, conditional on the registered readouts and manipulations. No additive or causal decomposition of total loss is assumed; a stronger estimate of relative stage contribution would require additional pre-specified identification assumptions or interventions. Imported meanings are registered in Appendix A.

2.3 Principle Statement

Flow Shaping Principle (single-quantity form). In an operating adaptive organization that has independently satisfied the Π -specific entry condition, an independently measured, update-capacity-matched $A_{corr}(II)$, representing veridical consequence–configuration correspondence available to the attribution stage of the update process, should prospectively predict held-out correspondence-correct persistent shaping $S_{corr}(II)$. A calibration obtained under one pre-specified class of correspondence degradation should transfer, within a pre-specified error criterion, to a distinct pre-specified class under the same primary frame, primary decoder definition, and outcome definition. Support is indexed to the tested ordered manipulation pair. An update-identical readout is reported as a diagnostic comparator and is not substituted for the primary A_{corr} test.

This is deliberately narrower than the claim that flow changes organization. It does not assert that every consequence produces change, that more available information always produces better performance, or that adaptation is globally beneficial. It asserts that once correspondence-sensitive updating has independently been established, the same measured correspondence quantity should organize loss of correspondence-correct shaping across different ways in which correspondence is damaged.

2.4 What Is and Is Not Empirical

If A_{corr} is zero, correspondence-correct discriminative updating cannot be generated from that substrate without another information source. Likewise, A_{corr} places an information-theoretic upper bound on S_{corr} . These statements follow from the structure of attribution and are not the empirical contribution of FSP.

The risky claim is that operating adaptive systems satisfying the entry condition exhibit a sufficiently stable relation between update-rule-relative A_{corr} and S_{corr} for calibration to predict held-out results and transfer across a distinct manipulation class. The gap between A_{corr} and S_{corr} reflects how efficiently the update rule reads and uses available correspondence. FSP therefore risks failure if that readout-and-use relation changes across manipulations, even when the attribution boundary remains true. Separate substrates may carry capacity and fidelity, or different corruption mechanisms may alter how correspondence is used.

A particularly important pre-commitment concerns cross-manipulation failure. If Prediction 1b succeeds separately within two manipulation classes but Prediction 1c fails between them, the single-quantity form of FSP is not supported. The failure may motivate a successor hypothesis distinguishing attribution capacity from attribution fidelity, but that two-axis successor is not validated by the data that defeated the single-axis claim. It requires a new specification and a new pre-registered canonical test.

3. Core Experimental Test

3.1 Manipulation Check 1a — Zero-Correspondence Verification

Hold episode content, consequence content, their marginal distributions, consequence availability, total exposure, update opportunity, and nonspecific disturbance constant. Randomly permute episode-consequence pairings so that the consequence remains available but the update-matched readout carries no usable information about the ground-truth correspondence under the primary II. Verify independently that A_{corr} has reached the pre-specified chance or equivalence region and that no pre-specified side channel restores the true pairing.

S_{corr} in this condition should be equivalent to the no-correspondence baseline within pre-specified bounds. Other organizational changes may still occur: global updates, homeostatic adjustments, wrong-content updates, or changes in non-target configurations are not required to vanish. This check establishes that the correspondence manipulation and A_{corr} readout jointly created the intended zero-correspondence condition.

If A_{corr} does not reach the equivalence region, a side channel remains, or veridical S_{corr} remains above baseline under a purported zero-correspondence condition, Manipulation Check 1a has failed and Predictions 1b–1c are not interpretable. Such a failure first indicts the manipulation or readout specification; it does not by itself falsify FSP.

3.2 Prediction 1b — Within-Manipulation Prospective Prediction

Choose a manipulation class M1 that degrades correspondence across graded levels. Examples include controlled trace delay or degradation, provided each level is accompanied by an independent substrate readout of A_corr. Manipulation intensity is not itself A_corr, and A_corr may not be inferred from S_corr.

Use an independent calibration set to estimate the relation between A_corr and S_corr and, where applicable, a retention window or transition region. The mapping class, fitting procedure, uncertainty model, held-out levels, and error criterion must be specified before the primary held-out outcomes are inspected. The fitted relation must then predict S_corr for held-out M1 levels, episodes, or initial states.

Prediction 1b is not supported if independently measured A_corr fails to predict held-out S_corr under the pre-specified criterion, or if a retention window estimated in calibration fails prospectively. A monotonic result inferred only after observing the full response curve is insufficient.

3.3 Prediction 1c — Cross-Manipulation Generalization

Use the mapping calibrated under an attribution-loss class M1 to predict S_corr under a distinct manipulation class while preserving the same B, R, P, primary Π , update-matched A_corr readout, and S_corr outcome. Manipulation classes must not be treated as interchangeable. Attribution loss includes trace delay, decay, or ablation that weakens the available correspondence. Content misattribution M2 preserves the target configuration but pairs C_i with consequence y_j through marginals-preserved permutation: the location is correct and the consequence content is wrong. Location misattribution M3 preserves consequence content y_e but reassigns provenance or routing to C_j: the content is correct and the update location is wrong (Figure 2).




Figure 2. M1, M2 and M3 are not members of one scheme. Panel A classifies misattribution by which relation is falsified: M2 preserves the target configuration and corrupts consequence content; M3 preserves consequence content and reassigns provenance or routing to C_j . The fourth cell is a severe control, not part of the core protocol. Panel B shows a distinct mechanism family—trace delay, decay, or ablation—displayed as a procedural sequence of candidate manipulation levels ordered by intended severity. The displayed order is not a measurement: each level requires an independent capacity-matched A_{corr} readout, and empirical ordering and monotonicity are outcomes rather than premises. Manipulation Check 1a is implemented by full marginals-preserved permutation and is therefore the zero-correspondence endpoint of the pair-permutation family used for M2, not a level of the Panel B axis and not a confirmatory transfer condition. Because the two panels manipulate different things, their control requirements are stated separately. Prediction 1b tests prospective calibration within Panel B; Prediction 1c tests transfer to a pre-specified cell of Panel A. Manipulation coverage is a reporting axis distinct from evidential readiness and cross-domain scope.

Prediction 1c is supported only if the calibration predicts held-out outcomes for the pre-specified ordered transfer pair within the error criterion. A first canonical test may examine M1–M2; a separate M2–M3 or M1–M3 transfer would expand manipulation coverage. Success on one pair does not establish manipulation-general A_{corr} . Cross-manipulation failure rejects the single-quantity claim for that transfer even if both within-class curves are individually orderly.

3.4 Confirmatory Commitments

The following items must be fixed before the primary test: the frame F; the single primary Π ; the entry margin; the capacity-matched-but-not-identical primary A_{corr} decoder class; the update-identical diagnostic readout; the input access and update-readout justification for both; the single primary S_{corr} measure; the unsaturated range; the separately named manipulation classes and ordered transfer pair; the

calibration and held-out sets; the mapping class; equivalence bounds; and the cross-manipulation error criterion.

The first canonical test must report A_{corr} and the resulting A_{corr} -to- S_{corr} calibration under both pre-registered readout classes. The capacity-matched-but-not-identical decoder is the primary test of correspondence accessible to the attribution process. The update-identical readout conditions on the attribution signal already produced by the system and is a diagnostic of downstream consolidation, competition, decay, saturation, and interference. The two curves must be reported together even if only the primary curve is used for the confirmatory 1b–1c decision.

The unsaturated range must be established in an independent calibration phase by showing that the selected outcome is not already at a ceiling or floor and remains responsive to graded exposure or update opportunity in the chosen interval. “Saturation” may not be invoked after a failed primary result unless the pre-specified diagnostic demonstrates that the calibration itself was invalid.

A complete canonical test should report Manipulation Check 1a together with the two theory-risk predictions, 1b and 1c. Binary intact-versus-scrambled differences alone test neither prospective calibration nor transfer. Within-manipulation prediction without cross-manipulation transfer supports, at most, a manipulation-specific model. A successful 1c result is indexed to the tested ordered manipulation pair; broader manipulation coverage requires additional pre-specified transfer pairs.

4. Operational and Measurement Constraints

4.1 The Primary Partition Π

Pathway specificity is always relative to a partition. Π may distinguish network modules, parameter groups, routed expert nodes, neural ensembles, synaptic populations, or another pre-specified configuration set. The partition must be justified independently of the primary outcome—for example, from architecture, intervention accessibility, or a training-independent functional definition.

Only one Π is primary in a confirmatory test. Multiple resolutions may be registered, but they require an explicit multiplicity procedure or exploratory status. Repartitioning configuration space after observing where an effect appears is prohibited because it can convert global or misplaced change into apparent specificity.

4.2 Independent Measurement of A_{corr}

Primary A_{corr} must be measured from the substrate that carries correspondence into the attribution stage of the update process, independently of persistent organizational change, by a decoder whose functional class and input access are capacity-matched to, but not identical with, the attribution computation performed by the actual update rule. The purpose of this match is to measure correspondence that a process with the system’s readout capacity could use without defining A_{corr} as the system’s own computed credit signal. An unconstrained decoder measures information present in the substrate rather than correspondence available to updating; an update-

identical decoder measures the attribution signal after that stage has already been completed. The latter is valuable as a diagnostic comparator but cannot serve as the primary A_{corr} readout for the attribution claim.

Delay, pair permutation, trace ablation, provenance reassignment, and similar procedures are manipulation types, not measurements of A_{corr} . Each manipulation level requires an independent capacity-matched primary readout and an update-identical diagnostic readout. Both must be evaluated against ground-truth pairings, not merely configuration identity. Their functional classes, feature access, training splits, chance levels, calibration procedures, and matching justifications must be fixed independently of S_{corr} . A more powerful exploratory decoder may be reported, but it cannot replace either pre-registered class or rescue a failed primary result.

4.3 Measurement of S_{corr}

S_{corr} must be persistent, pathway-level, and correspondence-correct. Global performance is insufficient because another pathway may compensate. Localization is also insufficient because a consequence can be assigned to the correct location with the wrong content. The outcome must therefore be signed or content-sensitive relative to a pre-specified consequence-consistent update.

One possible artificial-system measure is the signed projection of the persistent parameter or routing update in target cell C_i onto an update direction defined from the true consequence, minus the corresponding wrong-direction component and a matched non-target contrast. One possible neural measure is a pre-specified change in the target ensemble or tagged synaptic population that is aligned with the behavioral or physiological direction associated with the true episode-consequence pair. These examples illustrate the requirement; they are not universal formulas.

4.4 Marginals-Preserved Controls

The decisive misattribution control preserves the episode set, the consequence set, both marginal distributions, total consequence availability, total internal flow, update count, and relevant computational or metabolic expenditure, while randomly permuting the episode-consequence correspondence. Scrambling content is not an adequate substitute because it changes input statistics and effective update magnitude.

A no-consequence condition may be included to test the imported Connection Theory feedback-availability boundary, but it is not the decisive FSP control. FSP requires the consequence to remain available while veridical correspondence is selectively lost or corrupted.

5. Candidate Realizations

5.1 Artificial Adaptive Systems: Primary Canonical Target

Instrumentable artificial systems provide the cleanest initial attack. A routed continual-learning network can expose which expert, module, memory cell, or parameter group carried each episode; which consequence belongs to that episode; what

substrate state carries the pairing into the update; and where the persistent update occurs. Continual-learning settings are especially relevant because interference and replay make provenance explicit design problems [9].

A canonical experiment could first establish Π -specific updating under veridical high-A_corr conditions. It could then use trace delay or controlled retention degradation as M1, marginals-preserved pair permutation as the content-misattribution class M2, and systematic provenance reassignment as the location-misattribution class M3. Primary A_corr would be read through a separately implemented decoder with the same functional capacity and input access as the routing, eligibility, or gradient-allocation readout used by the actual update rule. The system’s own attribution signal would be recorded in parallel as the update-identical diagnostic. S_corr would be evaluated from persistent correct-location and correct-direction updates in the target module. Calibration on M1 could first be tested on M2; a distinct transfer involving M3 would extend, rather than be assumed by, that manipulation coverage.

Sleep-like replay in artificial networks is an application setting, not a separate principle. Offline replay has been reported to reduce catastrophic forgetting in artificial neural networks [11]. A specifically FSP-relevant experiment must go further: replay volume, episode marginals, consequence marginals, compute, and update count must be matched, while only veridical episode-consequence correspondence is permuted. Benefits from more computation or from replaying cleaner content would not test the core claim.

5.2 Neural Systems and Sleep: Candidate Extension

Neural plasticity supplies plausible association-carrying substrates, including activity-dependent traces, synaptic tags, ensemble reactivation, and circuit-specific consolidation [1]. The present obstacle is not merely decoding accuracy. FSP requires a principled justification that the chosen neural decoder has a functional class and input access matching the readout used by the biological update process. Because that readout format is not currently known, neural A_corr is not yet fully defined for a confirmatory test, even if an external decoder can recover cue or episode identity.

Sleep is relevant only as a candidate setting for internal reorganization. Whether sleep is classified as an operating mode or an interruption depends on the organization selected as O. Continuing neural-metabolic organization may remain active through sleep, whereas a narrower wake-task organization may cease active manifestation. FSP applies only where the selected organization remains active and update-capable. Where that organization is interrupted, its persistence and any later reconstruction fall outside the present framework; FSP resumes only when an active update-capable organization again exists.

Replay and reorganization during sleep are supported by observations of hippocampal and cortical reactivation [5,14], homeostatic synaptic regulation [2], maintained learning efficiency [3], and branch-specific structural change after learning

[15]. These findings motivate a test environment but do not establish the FSP calibration claim.

Targeted Memory Reactivation (TMR) offers a possible biological correspondence manipulation because learned cues can be presented during sleep and compared under veridical versus permuted cue–memory pairings while preserving the cue and memory marginals. Existing TMR research supports the more modest claim that sleep cues can influence later memory, with variable and generally limited effect sizes [4,8]. A valid FSP test would additionally require a justified update-matched neural readout of veridical cue–memory correspondence and a pre-specified correspondence-correct neural or behavioral update measure. External decodability alone is insufficient.

5.3 Human Organizations: Candidate Transfer Domain

Human organizations maintain audit trails, responsibility assignments, case histories, and review procedures that can preserve correspondence between decisions and consequences. These are candidate association mechanisms, not independent measurements of A_{corr} . At present, no substrate-level readout has been established that measures the correspondence actually available to the organizational update process independently of the later procedural or behavioral change.

The domain therefore remains a candidate transfer target rather than a current realization of the calibrated claim. Ethical limits, delayed outcomes, selective reporting, personnel turnover, and the organization’s control over its own evidence further complicate marginals-preserved intervention. Organizational examples may motivate measurement development but do not enlarge validated FSP scope.

6. Discussion and Limitations

6.1 The Main Contribution

The broad statement that organization constrains flow is partly constitutive of organization, and the zero-information attribution boundary is a logical constraint. Neither is the empirical contribution of FSP. The proposed contribution is the update-rule-relative transfer claim: after Π -specific updating is independently established, a readout matched to the actual update process should measure veridical consequence–configuration correspondence that prospectively predicts correspondence-correct shaping for held-out conditions and retains predictive meaning across a pre-specified manipulation pair.

The primary capacity-matched A_{corr} -to- S_{corr} relation records the effective efficiency with which correspondence accessible at the attribution stage is converted into persistent shaping. Prediction 1b tests prospective calibration within a manipulation class; Prediction 1c tests whether that relation is stable enough to transfer across a distinct class. Manipulation Check 1a is quality control for those tests, not a third independent empirical wager. The update-identical diagnostic curve begins after the system has produced its own attribution signal and therefore measures a different segment: downstream consolidation, competition, decay, saturation, and interference.

The single-quantity claim can be killed by orderly within-class prediction that fails to transfer across the pre-specified pair. Such a result would leave feedback availability and the data-processing boundary intact while rejecting the FSP transfer claim.

6.2 Design Implications and Their Status

Adaptive-system designers often have good reasons to preserve causal provenance, eligibility records, traceable decision paths, and protected replay. Human organizations may likewise benefit from audit trails, post-event review, and mechanisms that reduce silent reassignment of consequences. These implications follow from the attribution boundary, which FSP does not claim as novel; they do not depend on, and are not evidence for, FSP's calibrated correspondence-to-shaping claim.

The empirical FSP claim could therefore fail while these engineering practices remain useful. Keeping practical design guidance separate from experimental support avoids treating usefulness as validation.

6.3 Recorded Failure and Removed Branches

The entries below record two distinct kinds of defeat: withdrawn theoretical hypotheses and withdrawn methodological claims about what a registered comparison can identify. Unless expressly stated otherwise, neither type alters the core FSP proposition or Predictions 1b–1c.

An earlier working hypothesis proposed that protected sleep-like reorganization would emerge only after sufficiently integrated or centralized nervous organization developed. Comparison with sleep-like states in *Cassiopea*, *Hydra*, and *C. elegans* did not support that boundary [6,7,10]. The hypothesis is withdrawn in its previous form. What organizational features predict sleep-like states remains open, but that question is not required by the FSP core claim.

An earlier formulation, appearing in §6.4 of versions up to 3.6.3, stated that the relative contributions of attribution-stage and downstream loss must be estimated from the paired calibration curves. That formulation exceeded what the registered comparison identifies. Comparison of the paired curves can localize differential predictive loss before versus after the system-generated attribution signal, conditional on the registered readouts and manipulations, but it does not identify an additive or causal decomposition of total loss without additional pre-specified assumptions or interventions. The stronger identification claim is withdrawn.

The present version also removes physical flow analogies, civilization-level illustrations, unrestricted evolution equations, and extended claims about dynamic stability. Those branches were either non-evidential, imported from Connection Theory, or insufficiently constrained to contribute to the correspondence-calibration test.

6.4 Limitations and Next Test

The main limitation is empirical. No complete canonical pass consisting of Manipulation Check 1a plus Predictions 1b–1c has yet been reported. Artificial systems offer the most feasible first test, but even there the result depends on a capacity-

matched-but-not-identical primary A_corr readout, an update-identical diagnostic readout, a ground-truth correspondence definition, a direction-sensitive S_corr measure, and strict separation of entry, calibration, and held-out evaluation.

Neural transfer currently lacks a justified decoder class matched to the biological update readout, not merely sufficient data or decoding power. Human-organizational transfer likewise requires a measurement concept tied to the readout actually used by the organizational update process rather than an audit record observable only to an external analyst. Until those problems are solved, FSP scope remains unvalidated beyond the proposed artificial canonical target.

The difference between the capacity-matched and update-identical calibration curves is itself a required empirical result. It can localize differential predictive loss to stages before versus after the system-generated attribution signal, conditional on the registered readouts and manipulations. It does not by itself provide an additive or causal decomposition of total loss. Any stronger estimate of relative stage contribution requires additional identification assumptions or interventions specified in advance.

Manipulation coverage is a reporting axis separate from evidential readiness and cross-domain scope. Success on one ordered transfer pair, such as attribution loss to content misattribution, supports only that pair. Establishing manipulation-general validity requires additional pre-specified transfers, including tests that distinguish content misattribution from location misattribution. Coverage cannot be inferred from a single success.

The prevalence of systems satisfying the Π -specific entry condition is also an open empirical question. A conditional principle restricted to entry-passing systems risks being true but narrow, or nearly empty, if most real update rules are global or provenance-insensitive. Determining how common entry-passing systems are, and which organizational features enable them, is a necessary companion research program rather than an assumption of FSP.

Order effects among task-facing and reorganization phases remain outside the core FSP test; the referenced ownership boundary is recorded in Appendix A.4. The boundary-only distinctions registered in Appendix A.3 likewise remain outside the present empirical claim. The present version is frozen with explicit failure conditions but without a completed canonical test. Apart from the imported-meaning dependency conditions governed exclusively by Appendix A.5, the paper should be reopened only if a canonical artificial-system test exposes a foundational defect or if the required measurement conditions prove incoherent or infeasible.

References

- [1] A. Citri and R. C. Malenka, Synaptic plasticity: Multiple forms, functions, and mechanisms, *Neuropsychopharmacology* 33 (2008) 18-41.
- [2] G. H. Diering, R. S. Nirujogi, R. H. Roth, P. F. Worley, A. Pandey and R. L. Huganir, Homer1a drives homeostatic scaling-down of excitatory synapses during sleep, *Science* 355 (2017) 511-515.

- [3] S. Fattinger, T. T. de Beukelaar, K. L. Ruddy et al., Deep sleep maintains learning efficiency of the human brain, *Nature Communications* 8 (2017) 15405.
- [4] X. Hu, L. Y. Cheng, M. H. Chiu and K. A. Paller, Promoting memory consolidation during sleep: A meta-analysis of targeted memory reactivation, *Psychological Bulletin* 146 (2020) 218-244.
- [5] D. Ji and M. A. Wilson, Coordinated memory replay in the visual cortex and hippocampus during sleep, *Nature Neuroscience* 10 (2007) 100-107.
- [6] H. J. Kanaya et al., A sleep-like state in Hydra unravels conserved sleep mechanisms during the evolutionary development of the central nervous system, *Science Advances* 6 (2020) eabb9415.
- [7] R. D. Nath, C. N. Bedbrook, M. J. Abrams et al., The jellyfish *Cassiopea* exhibits a sleep-like state, *Current Biology* 27 (2017) 2984-2990.e3.
- [8] D. Oudiette and K. A. Paller, Upgrading the sleeping brain with targeted memory reactivation, *Trends in Cognitive Sciences* 17 (2013) 142-149.
- [9] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan and S. Wermter, Continual lifelong learning with neural networks: A review, *Neural Networks* 113 (2019) 54-71.
- [10] D. M. Raizen et al., Lethargus is a *Caenorhabditis elegans* sleep-like state, *Nature* 451 (2008) 569-572.
- [11] T. Tadros, G. P. Krishnan, R. Ramyaa and M. Bazhenov, Sleep-like unsupervised replay reduces catastrophic forgetting in artificial neural networks, *Nature Communications* 13 (2022) 7742.
- [12] K. Wang, *Connection Theory: A General Framework for Adaptive Network Dynamics - From Temporally Extended Existence to Emergent Dynamic Properties*, Version 4.6, Zenodo (2026), doi:10.5281/zenodo.21714702.
- [13] K. Wang, *Evolutionary Relativity: Frame-Qualified Organizational Change and Observer-Relative Reconstruction*, Version 3.6.1, Zenodo (2026), doi:10.5281/zenodo.21715199.
- [14] M. A. Wilson and B. L. McNaughton, Reactivation of hippocampal ensemble memories during sleep, *Science* 265 (1994) 676-679.
- [15] G. Yang, C. S. W. Lai, J. Cichon, L. Ma, W. Li and W.-B. Gan, Sleep promotes branch-specific formation of dendritic spines after learning, *Science* 344 (2014) 1173-1178.

Appendix A. Register of Imported Meanings, Bound to Connection Theory v4.6

This appendix is the sole and exhaustive register of meanings imported from Connection Theory [12], archived under version DOI 10.5281/zenodo.21714702. No meaning other than those registered below is imported from that work. Entries are grouped by the role the imported meaning plays here, not by their position in Connection Theory.

A.1 Directly load-bearing imported meanings

Determine the target class, the entry condition, or a stated boundary of the present claim.

adaptive (CT §1). Bound meaning: The capacity for consequence-sensitive organizational change; not a guarantee that any given change improves environmental fit. **Function here:** Restricts the target class to systems capable of consequence-sensitive organizational change. FSP separately requires an independently specified consequence-attribution process; the exclusion of reciprocal physical couplings in §1.3 does not follow from this term alone.

organization (CT §6). Bound meaning: The continuously coordinated, state-dependent integration of stable connections. Coordination may be expressed through routing, timing, sequencing, gain, synchronization, combination, or other relations; any dimension may be absent or degenerate. A distribution over temporally ordered traversals is one possible operational representation, not the definition. **Function here:** Provides the organizational object within which Π identifies flow-carrying configurations. Π does not reduce organization to a static connection inventory.

feedback availability (CT §4.1). Bound meaning: Feedback can act only through consequences that become causally available to the system. A consequence that never returns, leaves no retained proxy, and is not reconstructed by another process cannot influence later organization through feedback. **Function here:** The inherited boundary of §1.2. FSP's controls hold consequence availability constant while correspondence is degraded, so that FSP is not tested against this boundary (§4.4).

active manifestation (CT §2; operationalized §8.4 Pred. 1). Bound meaning: Stated in CT as one of four separable questions: whether the relevant organization is operating now. Active manifestation requires the relevant organization to operate. **Function here:** FSP applies only where the selected organization is actively operating and update-capable, and resumes only when such an organization again exists (§5.2).

Compositional shorthand. *Operating adaptive organization* is compositional shorthand for an organization that is currently operating and is adaptive, in the respective meanings bound above. Connection Theory does not name the compound as a separate term, and the phrase introduces no additional theoretical primitive.

A.2 Definition-closure imports

Required to state the A.1 meanings without leaving an undefined term; not independently load-bearing for any prediction or failure criterion here.

interaction (CT §3). **Bound meaning:** Ongoing exchange between an adaptive system and other components or its environment. **Function here:** Closes the imported definition of feedback. An FSP episode may instantiate an interaction, but the two are not defined as coextensive.

feedback (CT §4). **Bound meaning:** The history-dependent, selective influence of past interactions on a system's future organization. Names a selection criterion—which past interactions continue to influence future organization—not a rule for how the selected influence updates organization. **Function here:** Closes the imported definition of connection. The update rule is left open by CT and is not supplied by FSP either.

connection (CT §5). **Bound meaning:** A stabilized, feedback-derived relationship between components. **Function here:** Closes the imported definition of organization. FSP does not assume that Π partitions individual connections, and does not treat connection inventory as sufficient for organization.

A.3 Boundary-only imported distinctions

Imported solely to state what the present paper does not address. No prediction, entry condition, measurement constraint, or failure criterion here depends on their content.

retained capability (CT §2; §8.4 Pred. 2a). **Bound meaning:** Whether capability-relevant persistence state survives an interruption and can support later reconstruction. One of the four separable questions; distinct from the state object that supports it. **Function here:** Excluded. FSP does not model recovery of capability after interruption (§1.2).

capability-relevant persistence state (CT §2; §8.4 Pred. 2b; §8.5). **Bound meaning:** Property-relative state whose survival across interruption can support later reconstruction of a specified capability. Localization is property-relative and may require joint interventions because of distributed or redundant encoding. Requires different evidence from the continuity-bearing substrate. **Function here:** Excluded. FSP does not model the persistence or degradation of this state during inactivity (§1.2).

Organizational Continuity (CT §7; operationalized §7.1). **Bound meaning:** Given a pre-specified individuation frame—system boundary, lineage level, and adjudication interval—Organizational Continuity holds where the resumed system is connected to the earlier system by a qualifying non-forking causal lineage carried through an identifiable continuity-bearing substrate that causally contributes to resumption. Within a fixed frame the judgment is property-invariant. **Function here:** Excluded. FSP does not model lineage continuity across interruption (§1.2).

current applicability (CT §2; §7). **Bound meaning:** Whether the retained or reconstructed organization remains suited to present environmental conditions. Not guaranteed by lineage continuity or by faithful reconstruction. **Function here:** Excluded. FSP makes no claim that correspondence-correct updating improves environmental fit (§2.3).

A.4 Referenced but not imported

Connection Theory's update-non-commutativity prediction (CT §6.1; §8.4 Pred. 3) is referenced in §1.2 and §6.4 only to delimit ownership of causal order effects. It is neither imported into nor tested by the present framework, and no meaning in it is bound by this register.

A.5 Dependency-governance statement

The claims of this paper remain indexed to the meanings bound above and are not retroactively altered by later revisions of Connection Theory. A material change to an A.1 meaning requires a separately versioned dependency review. Where that review determines that a prediction or failure criterion is affected, extending the present claim to the revised meaning requires a separately versioned restatement rather than reinterpretation of this paper. A change to an A.2 term requires such review only where it materially alters an A.1 meaning. A change to an A.3 distinction does not alter the empirical predictions of this paper but may require revision of the corresponding scope statement.