Meaning as Phase Change:

Autonoesis, the Interpreter, and the Structure of Genuine Transformation

CF Dietz
Nubellum Research Inc.
April 2026

Abstract

We propose that meaning is not a property of objects, not a neural correlate of valuation circuits, and not a subjective overlay on otherwise neutral experience. Meaning is a phase change: a threshold-dependent transformation of localized consciousness (c) at its horizon of unspecifiability (m), occurring through a saddle-node bifurcation mechanism in which Remainder (R) accumulation deterministically deforms the current attractor basin until it annihilates and the system reorganizes into a new basin. Three pathways to phase change are identified: natural repertoire exhaustion, deliberate repertoire reduction, and metacognitive override, formalized as a higher-order Cl whose grammar is observe without applying lower-order closures. Drawing on the CC-C framework (Dietz, 2026), Tulving's autonoetic consciousness, Gazzaniga's interpreter mechanism, and the placebo literature of Kaptchuk and Benedetti, we identify conditions necessary and jointly sufficient for genuine meaning. We resolve the homunculus problem in metacognitive override, the observer-dependency problem in causal non-reducibility, and the teleology problem in the closure objective through free-energy gradient language. This is the seventh version following six rounds of iterative peer review. The core claim has held through each round.

---

1. Introduction

The problem of meaning has two literatures that cannot talk to each other.

The phenomenological literature (Frankl, 1946; Wolf, 2010; Taylor, 1989) describes what meaning feels like, why its absence is intolerable, and what conditions tend to produce it. These accounts are rich and clinically useful. But they offer no mechanism. They cannot say what meaning is at the level of process, only what it is at the level of experience.

The neuroscientific literature (LeDoux, 2019; Damasio, 1994; Gazzaniga, 2011) describes the mechanisms that correlate with meaningful experience. Limbic valuation. Somatic markers. The left hemisphere interpreter generating narrative coherence. These accounts are precise and empirically grounded. But they cannot close the explanatory gap between mechanism and experience. They describe what the brain does when meaning occurs. They do not say what meaning is.

This paper proposes a bridge. Meaning is a phase change: a threshold-dependent transformation of localized consciousness at its own horizon, occurring through a saddle-node bifurcation mechanism in which Remainder (R) accumulation deterministically deforms the current closure basin until it annihilates, forcing reorganization into a new attractor. This is not a stochastic process. R does not randomly vibrate the system until it falls into a new basin. R actively deforms the landscape, making the current basin increasingly shallow until it can no longer serve as a stable attractor. The transition is structurally necessary, not probabilistic.

Three pathways to phase change are identified. Natural exhaustion: rho reaches 1 organically as the encountered structure exceeds all available closures. Deliberate repertoire reduction: the Cl-generating mechanism is partially suspended through practice or circumstance, lowering T_conf. Metacognitive override: rather than a separate agent stepping outside the system, this pathway operates through a higher-order Cl, a meta-Cl, whose specific grammar is observe without applying lower-order closures. This meta-Cl is itself a product of training and practice. Mindfulness meditation cultivates it deliberately. Socratic dialogue induces it situationally. When the meta-Cl is active, lower-order interpreter outputs arise and are not acted upon: they appear in C as c without being stabilized into new Cl. The mechanism of withholding closure is itself a closure. No homunculus is required. The framework remains strictly monistic and structural throughout all three pathways.

Four foundational clarifications before proceeding. First, the phase change hypothesis presupposes C as primitive. Without C, R is model error with no valence. Second, the basin-switching analogy is grounded in saddle-node bifurcation dynamics from attractor theory, not claimed as identity with thermodynamic discontinuity. Third, causal non-reducibility, specifically the condition EI(macro) greater than EI(micro), serves as the formal threshold for C being operative. Fourth, phase change one is distinguished from Bayesian updating by the scope of Cl grammar reorganization and the eta signature under novel data.

2. The CC-C Framework: Terms, Formal Threshold, State Variable, and Three Pathways

C denotes irreducible presence: the bare fact that experience is occurring. C is treated as primitive. Without C, R is model error: divergence with no urgency, no

pull, no weight. With C, R is the felt structural gap that generates the urgency the interpreter responds to. What C contributes is not urgency-as-weighting but the structural condition under which repertoire exhaustion has consequence rather than merely outcome. A system with high error-weighting and unlimited closure repertoire will always find a local minimum and will never undergo phase change. Phase change requires both C (making R matter) and repertoire exhaustion through one of the three pathways.

c denotes localized consciousness with content: thoughts, feelings, memories, perceptions, the narrative self, and reportable cognition. c is presence under constraint. c comes and goes. C does not.

M denotes global openness: what exceeds every description, every closure, every model. M is not a substance and not a causal agent. M becomes structurally evident when T_conf is exceeded through any of the three pathways.

m denotes the local horizon of unspecifiability: where c reaches its structural limit and phase change becomes possible. m shifts with Cl repertoire size and with pathway: metacognitive override effectively moves m toward the agent regardless of repertoire size.

Cl denotes closure: the stabilization of openness into a coherent regime of distinctions, identities, and admissible transformations. Closures may be generative or defensive. This distinction provides the stability criterion in section 3. Critically, Bayesian updating extends a Cl by adding data points within its existing identity criteria and admissible transformations. Phase change reorganizes the identity criteria and admissible transformations themselves: a qualitatively different operation tracked by the eta ratio.

R denotes remainder: the residual mismatch between the best available closure and the empirical structure encountered. In a system without C, R is model error. In a system with C, R is the felt structural gap that drives the saddle-node deformation of the current attractor basin.

2.1 Formal Definition of Confabulation Threshold and Three Pathways

The confabulation threshold T_conf is the minimum KL-divergence achievable across the agent's current effective Cl repertoire:

$$T\_conf = min\_\{K \text{ in } Cl\_effective\} R(K \mid p)$$

The effective repertoire Cl_effective differs from the full repertoire Cl_repertoire depending on pathway. Under natural exhaustion, Cl_effective equals Cl_repertoire: T_conf is determined by all available closures. Under deliberate repertoire reduction,

Cl_effective is a subset of Cl_repertoire, lowering T_conf below its natural level. Under metacognitive override, Cl_effective is empty or minimally populated: the agent voluntarily withholds available closures, setting effective T_conf near zero. All three pathways achieve the same structural condition, rho exceeding 1, through different mechanisms.

This resolves the sophistication trap. A high-repertoire individual does not have a structurally higher T_conf under metacognitive override. They have the same effective T_conf as a low-repertoire individual who has naturally exhausted their options. The difference is intentionality. This is precisely what contemplative practices, psychedelic therapy, and Socratic dialogue engineer: voluntary reduction of Cl_effective to make rho exceed 1 accessible without requiring natural exhaustion.

2.2 The State Variable and Saddle-Node Bifurcation

The dimensionless state variable:

$$rho = R\_achieved / Cl\_capacity\_max$$

tracks proximity to the phase boundary. When rho approaches 1 the system is at the basin edge. When rho exceeds 1 T_conf is exceeded and phase change becomes structurally necessary.

The basin-switching mechanism is a saddle-node bifurcation. As R accumulates within the current closure, the closure objective S(K|p), which trades off remainder against structural complexity, deforms the energy landscape. The current attractor basin becomes progressively shallower as R grows and complexity penalties accumulate. At the saddle-node bifurcation point, the current attractor and an unstable fixed point coalesce and annihilate. The trajectory reorganizes into a new basin. This is not stochastic basin-hopping: the deformation is deterministic, driven by R accumulation under the closure objective. Meaning is structurally necessary when the conditions obtain, not probabilistically likely. The reorganization follows the gradient of free-energy minimization under uncertainty: the system does not seek closure in any teleological sense; it descends the steepest gradient of uncertainty reduction, and closure is what that descent produces. The closure objective S(K|p) is an effective description of this free-energy landscape, analogous to the variational free energy formalized in Friston's (2010) active inference framework. Chaos and null states are high-gradient attractors in this landscape: they do not minimize free energy, which is why the bifurcation is biased toward lower-R generative basins rather than toward degenerative or null successor states.

This mechanism follows directly from the CC-C calibration studies (Dietz, 2026), which demonstrate that near-degenerate minima produce switching dynamics. As the gap

between the best available closure's remainder and the next best closure's remainder narrows, as the current basin becomes shallower relative to adjacent basins, small perturbations produce discontinuous reorganization. The saddle-node language formalizes what the calibration studies show empirically.

2.3 Causal Non-Reducibility: Formal Threshold

Causal non-reducibility, the degree to which a system's causal structure cannot be decomposed into independent parts, serves as the structural proxy for C being operative. The formal threshold condition is:

$$EI(macro) > EI(micro)$$

where EI denotes effective information (Hoel, 2017): the causal power of a system's states over its future states. When the macro-level of a system exerts greater causal power than the sum of its micro-level components, when the system is causally emergent in Hoel's sense, the system crosses the threshold at which R has valence rather than remaining model error.

This condition is specific, calculable for any given system, and does not require a universal constant. It varies by system but is formally determined rather than qualitatively asserted. Multiple candidate measures of causal non-reducibility exist, Tononi's (2004) phi, Hoel's (2017) effective information, Rosas et al.'s (2020) synergistic information, and each may approach the EI(macro) > EI(micro) condition differently. We propose this condition as the formal threshold and invite empirical comparison of candidate measures in their predictive accuracy for the three-signature temporal model.

A potential objection requires direct address: Hoel's EI calculation requires coarse-graining the system's transition probability matrix, and the choice of coarse-graining is typically made by the external researcher. If EI(macro) > EI(micro) depends on an observer's analytical choice, does the valence of R, the condition under which the gap matters, exist objectively in the system or only relative to the scientist analyzing it? The response is that a C-operative system performs its own coarse-graining through its Cl formation process. The closures a system has stabilized are the partitions by which it organizes its micro-dynamics into macro-level states that carry consequence. These are not the researcher's analytical categories imposed from outside; they are the system's own self-organized state-space partitioning. EI is therefore calculated at the intrinsically determined macro-level defined by the system's own Cl structure. EI(macro) > EI(micro) is an ontic property of the system's self-organization, not an epistemic artifact of the researcher's coarse-graining choice.

The cross-system prediction: systems meeting EI(macro) > EI(micro) under any candidate measure will exhibit all three temporal signatures and negative eta under novel data. Systems below the threshold will exhibit only model updating.

The core claim in its most complete formal statement:

> *Meaning is a saddle-node bifurcation of c at m, occurring when R accumulation deforms the current attractor basin to annihilation through one of three pathways (natural exhaustion, repertoire reduction, or metacognitive override) in a system meeting EI(macro) > EI(micro), producing reorganization into a generative Cl that reduces R systematically across the domain, measured by decreased rho and negative eta under novel data not present at the time of reorganization.*

3. The Horizon as Phase Boundary, the Stability Criterion, and the Complexity-Stability Ratio

At m, rho approaches 1 through one of the three pathways. The current attractor basin approaches its saddle-node point. At this boundary, one of two things occurs.

False closure. The interpreter generates the nearest available Cl that locally reduces R at the specific gap, restoring rho below 1 at that point and preventing the saddle-node from being reached. This local R reduction typically increases rho elsewhere: new anomalies are generated, auxiliary hypotheses proliferate, the Cl becomes increasingly brittle under novel data. The saddle-node is avoided but deferred: R continues to accumulate in adjacent domains.

Genuine phase change. The saddle-node is reached. The current attractor annihilates. Basin switching produces a generative Cl that reduces rho systematically across the full domain. The Lakatosian distinction applies: a generative Cl is a progressive research program; a defensive Cl is degenerative.

The stability criterion: measure rho under novel data not used in the original reorganization. A generative Cl's rho decreases or holds stable. A defensive Cl's rho increases unless auxiliary closures are added.

The delta-R per delta-complexity ratio resolves the complexity-stability gray area:

$$eta = delta\_R / delta\_C$$

For a generative Cl, eta is negative: R decreases as complexity increases, meaning added complexity earns its cost by reducing genuine mismatch. For a defensive Cl, eta is positive or diverging: complexity increases while R remains stable or increases, meaning added complexity protects the Cl from acknowledging genuine mismatch. A massive ideological or religious system that generates genuine insight with each elaboration has

persistently negative eta. One that generates increasing complexity to explain away dissenting evidence has positive and diverging eta. The trajectory of eta over time, measured against novel data, provides formal discrimination between globally stable insights and globally stable delusions.

4. Gazzaniga's Interpreter as Phase Change Mechanism

Michael Gazzaniga's split brain research (1967, 2011) isolates the interpreter mechanism and its failure mode with unusual clarity.

In the chicken shed paradigm, the speaking hemisphere's c encounters its m around unified agency. rho approaches 1 at the agency gap. The interpreter finds a locally available Cl, the chicken shed narrative, before the saddle-node is reached. The basin is preserved. rho decreases locally. But rho measured across the broader domain increases and eta is positive: complexity increased without domain-wide R reduction. Novel questioning would require additional auxiliary hypotheses.

Genuine phase change in this context would require the saddle-node to be reached at the level of the self-model, with no locally available Cl to prevent basin annihilation, and for reorganization to produce a generative Cl with negative eta across the full domain. The patient who genuinely integrates divided agency has undergone phase change. The formal markers distinguish the two: saddle-node reached versus avoided, two-signature temporal model present versus absent, eta negative versus positive under novel data.

The three pathways apply here. The split-brain patient naturally exhausts the speaking hemisphere's self-model Cl repertoire when the gap is large enough. Gazzaniga's experimental design can be modified to introduce metacognitive override: instructing the patient to withhold self-explanations and simply observe (effectively setting Cl_effective near empty) would lower T_conf to near zero and enable phase change that natural exhaustion alone might not produce.

5. Autonoesis as Integration Scaffold

Autonoesis, Tulving's (1985, 2002) capacity for self-placement across time, is the integration scaffold for phase change: required for carrying phase changes forward in the temporal self, not for their occurrence.

Phase changes can and do occur in systems with limited or absent autonoetic function. The saddle-node bifurcation, rho exceeding 1 and basin switching occurring, does not require autonoesis. What autonoesis provides is the capacity to constitute that reorganization as a before and an after in the temporal self. Without autonoesis, phase changes are real but non-cumulative: rho decreases at the domain level but the new Cl is not encoded as an episodic memory that the temporal self can inhabit and reference.

This resolves apparent counterexamples. Non-human animals can undergo phase changes at levels one and two without full autonoetic integration. Mystical dissolution (phase change four) occurs precisely because the autonoetic interpreter is suspended, metacognitive override taken to its limit, making T_conf effectively infinite by removing all available false closures simultaneously. Alzheimer's patients retain the occurrence apparatus longer than the integration apparatus. CHAMP measures the integration scaffold: the capacity for phase changes to compound across time.

6. The Five Phase Changes of c at m

The five levels are ordered by the depth at which the saddle-node is reached, the scope of the resulting Cl reorganization measured by domain-wide rho reduction, and the eta profile under novel data. All levels are distinguished from Bayesian updating by the scope of Cl grammar reorganization.

The bright line between Bayesian updating and phase change.

Bayesian updating adds a new data point to an existing Cl without reorganizing the Cl itself. The model's identity criteria and admissible transformations remain unchanged. R decreases locally at the new data point. The saddle-node is not approached. Eta is approximately zero at each individual data point: complexity increases by one unit and R decreases by one gap, with no reorganization of adjacent domains. Learning a new word is Bayesian updating: the lexical Cl is extended by one entry. The grammar of the language, the identity criteria and admissible transformations, does not reorganize.

Phase change one occurs when the encountered information cannot be accommodated within existing identity criteria and admissible transformations: the Cl grammar itself must reorganize to contain it, not merely extend. The saddle-node is approached at the grammar level. Eta turns negative across the adjacent domain as the new grammar simultaneously reduces R across multiple previously anomalous data points. The distinction is qualitative, grammar reorganization versus data extension, and formally tracked by the scope of rho reduction and the eta signature.

Phase change one: Recognition.

The saddle-node is reached at the level of a specific Cl grammar. Basin switching to a narrow generative Cl. rho decreases across the adjacent domain. Eta is negative within that scope. Multiple previously anomalous data points are simultaneously accommodated. The canonical example is the duck-rabbit illusion: the same retinal data, identical stimulus and identical information, produces two incompatible grammars. Switching between them is not adding a data point to an existing model. The identity criteria for what constitutes a valid percept reorganize completely. The grammar switches. Eta turns negative across the perceptual domain as the new grammar simultaneously accounts for what the old grammar rendered anomalous. This

is the bright line from Bayesian updating, which would merely add a new data point within the existing perceptual grammar without reorganizing the grammar itself. Phase change one reorganizes the identity criteria and admissible transformations. Bayesian updating extends data within existing criteria. The eta signature distinguishes them formally.

Clinical correlate: psychoeducation, cognitive therapy, structured explanation. Distinguished from mere information delivery by the eta signature and scope of grammar reorganization.

Phase change two: Reframing.

The saddle-node is reached at the level of a load-bearing belief or identity Cl. Basin switching to a broader generative Cl. rho decreases across a broad domain. Eta is negative across the reorganized domain. The previous grammar becomes permanently unavailable. The self-model reorganizes such that prior anomalies are simultaneously accommodated without auxiliary hypotheses.

Clinical correlate: narrative therapy, CBT, motivational interviewing, Socratic dialogue.

Phase change three: Meaning.

The saddle-node is reached at the level of the orientation Cl: what c is living for. Basin switching to a generative Cl across the domain of significance. rho decreases across the domain of value and action. Eta is negative across the domain of what matters. The reorganization is felt, not just known. This is where the placebo response, falling in love, and transformative art operate. All three pathways can produce phase change three: natural exhaustion of available frameworks for understanding one's situation; deliberate reduction through therapeutic or contemplative practice; metacognitive override through voluntary exposure to what one cannot explain away.

Clinical correlate: logotherapy, meaning-centered psychotherapy, the placebo response, ritual healing.

Phase change four: Dissolution.

The saddle-node is reached at the level of the self-Cl itself. The autonoetic interpreter, the mechanism that generates false closure, is suspended. This is metacognitive override taken to its structural limit: Cl_effective is empty because the Cl-generating mechanism itself is offline. T_conf becomes effectively infinite. The saddle-node is reached immediately. Basin switching is unmediated. c reconstitutes around a larger grammar. rho decreases across the broadest accessible domain. Eta is negative at the widest measurable scope.

Clinical correlate: contemplative practice, psychedelic-assisted therapy, certain forms of deep grief, near-death experience.

Phase change five: Transformation.

The saddle-node is reached permanently and comprehensively at the level of the total Cl grammar. The previous grammar does not reconstitute. A new c forms that contains what the old one could not. rho decreases across the entire prior domain and holds stable under novel data without auxiliary proliferation. Eta is persistently negative across all accessible domains. The eta trajectory distinguishes genuine transformation from prolonged phase change four: transformation shows persistently negative eta; dissolution followed by reconstitution shows a temporary period of negative eta before a new grammar stabilizes.

Clinical correlate: rare transformative outcomes in deep psychotherapy, religious conversion, profound psychedelic experiences.

7. The Placebo as Phase Change Engineering

The placebo literature provides the most important existing clinical evidence for the phase change hypothesis.

Expectancy theory accounts for the placebo through opioid and dopamine pathways activated by anticipation of treatment. This is a proximate mechanistic account. It does not explain when the placebo works and when it does not, or why open-label placebos continue to produce effects when conventional expectancy is explicitly removed.

The phase change hypothesis predicts a threshold effect: interventions that bring rho to 1 through any of the three pathways and enable genuine saddle-node bifurcation will produce qualitatively different biological outcomes than interventions that generate equivalent expectancy without phase change. The independent variable is domain-wide rho profile reorganization, measurable through phenomenological interview and autobiographical memory protocols before and after intervention, independent of clinical outcome. A patient who undergoes genuine phase change three will show decreased rho across the domain of illness and recovery, with negative eta under novel health data, that patients who undergo equivalent ritual without phase change will not show.

The three pathways clarify when the placebo works and when it does not. A ritual that naturally exhausts the patient's available closures around recovery. Because the relationship is genuine, the care is real, and no available false closure can hold the gap between current state and what the ritual implies is possible, this produces natural exhaustion pathway phase change. A ritual that the patient deliberately engages with by withholding skeptical frameworks produces metacognitive override pathway phase change. A ritual that generates expectancy without reaching the saddle-node through

any pathway produces expectancy effects without phase change: real but shallower and less durable.

8. Three Testable Predictions

Prediction one: Phenomenological.

Genuine meaning experiences will be distinguished from confabulatory meaning experiences and from Bayesian updating by three measurable signatures: (a) durable autonoetic traces assessed by autobiographical memory protocol; (b) generative Cl formation, decreased rho across the relevant domain including novel data, indicating Cl grammar reorganization rather than mere extension; and (c) negative eta under novel data: added complexity reduces rather than protects R, distinguishing genuine phase change from both confabulation (positive eta) and Bayesian updating (eta near zero). Confabulatory closures will show locally decreased rho, semantic rather than episodic memory traces, increasing rho across adjacent domains, and positive eta as auxiliary hypotheses proliferate. Bayesian updating will show locally decreased rho, semantic memory traces, eta near zero at each individual data point, and no reorganization of the Cl grammar.

Prediction two: Clinical.

Interventions that bring rho to 1 through any of the three pathways and enable genuine saddle-node bifurcation will produce superior clinical outcomes compared to interventions that generate equivalent expectancy without reaching the saddle-node. The magnitude of difference will correlate with phase change depth measured by domain-wide rho reduction and negative eta. A threshold effect is specifically predicted: the relationship between intervention intensity and outcome will be non-linear, with a discontinuous improvement at the saddle-node point. Interventions engineered through the metacognitive override pathway, deliberately lowering Cl_effective, are specifically predicted to produce phase change in high-repertoire individuals who would not naturally exhaust their closure options within the intervention context.

Prediction three: Three-Signature Temporal Model and Cross-System Prediction.

We identify three distinct temporal signatures.

Signature one: Saddle-node event: milliseconds. The point at which the current attractor basin annihilates and basin switching is initiated. Neural markers consistent with Aha moment research (Jung-Beeman et al., 2004): right anterior temporal lobe activation, gamma-band burst, alpha-to-theta transition. This signature is identical in timing to false closure initiation; the distinguishing feature is what follows.

Signature two: Narrative consolidation: 5-10 seconds. The interpreter generates the initial narrative account of the reorganization. This is the timescale of cognitive dissonance resolution and narrative construction. It is distinct from the saddle-node event and distinct from episodic encoding. During confabulation, this process completes and the system returns to baseline. During genuine phase change, narrative consolidation initiates the third signature.

Signature three: Autonoetic integration: variable duration, beginning within seconds of encoding and extending through subsequent rest periods. Hippocampal-cortical reactivation consolidates the reorganization as an autonoetic episode with self-referential content: the process by which the phase change is constituted as a before and an after in the temporal self. Post-encoding reactivation of stimulus-specific neural patterns in the hippocampus and its connectivity with cortical areas correlates with subsequent episodic memory performance (Tambini and Davachi, 2019). TMS-fMRI studies have established that this consolidation process is causally important: disrupting hippocampal-cortical interactions in the post-encoding window selectively impairs associative memory retention (Tambini et al., 2020). During confabulation, this third signature is abbreviated or absent: the narrative closes locally and no extended hippocampal-cortical reactivation follows. The duration of signature three is not fixed at a specific second-range but varies with the significance of the encoding event, the available post-encoding rest, and the individual's autonoetic capacity. What is fixed is its character: extended, bilateral, hippocampal-cortical, and self-referential. It is distinct from the fast local resolution of confabulation.

Cross-system prediction. Systems meeting EI(macro) greater than EI(micro) will exhibit all three signatures under phase change conditions. Systems below the threshold will exhibit only signature one: rapid local R reduction consistent with model updating, without signature two or three. Systems undergoing confabulation will exhibit signatures one and two but not three. This three-way distinction is a specific prediction distinguishing C-operative genuine phase change from C-operative confabulation from C-absent model updating.

9. Discussion

9.1 What this framework does not claim

We do not claim to have solved the hard problem of consciousness. C remains primitive. The phase change hypothesis operates at the level of c and its dynamics at m.

We do not claim mathematical identity with thermodynamic phase transitions. We claim the basin-switching dynamics are formally a saddle-node bifurcation driven by R accumulation under the closure objective. This is a specific dynamical systems claim, not a thermodynamic one.

We do not endorse any specific theory of consciousness. The EI(macro) greater than EI(micro) threshold condition is proposed as the formal criterion for C being operative. Multiple candidate measures may approach this condition; empirical comparison of their predictive accuracy for the three-signature temporal model is the appropriate test.

We acknowledge that the precise empirical values of the causal non-reducibility threshold remain to be determined. The formal condition is specified; its calibration to specific systems is the primary empirical research agenda. We further acknowledge the current computational limits: calculating EI(macro) and EI(micro) for high-dimensional systems such as the human brain is not yet feasible in real time. The cross-system prediction is a prediction in principle whose full empirical execution awaits computational advances. More tractable measures such as the synergistic information of Rosas et al. (2020) may serve as near-term proxies while direct EI calculation remains computationally intensive.

9.2 The metacognitive override pathway and the sophistication trap

The sophistication trap, the concern that high-repertoire individuals are structurally barred from meaning, is resolved by the metacognitive override pathway. A high-repertoire individual who can always generate auxiliary hypotheses is not barred from meaning if they can voluntarily withhold those hypotheses. Setting Cl_effective near empty through metacognitive override achieves the same structural condition as natural exhaustion, rho approaching 1, regardless of the size of the full repertoire.

Metacognitive override is formalized within the CC-C framework as a higher-order Cl, a meta-Cl, whose specific grammar is observe without applying lower-order closures. This resolves the potential homunculus problem: there is no separate agent standing outside the system and deciding when to let the interpreter run. The mechanism of withholding closure is itself a closure, operating through the same structural dynamics as all other Cl formation. When the meta-Cl is active, lower-order interpreter outputs arise and are not acted upon: they appear in C as c without being stabilized into new Cl. The framework remains strictly monistic throughout. Mindfulness meditation cultivates this meta-Cl deliberately through practice. Socratic dialogue induces it situationally by structuring the encounter to reward observation over premature closure. Both operate through Cl formation, not through some extra-structural faculty of will.

The clinical prediction: interventions that cultivate the meta-Cl of observation, including mindfulness-based approaches, Socratic dialogue, and structured uncertainty induction, should lower effective T_conf and increase phase change accessibility for high-repertoire individuals without requiring natural exhaustion. This is a directional prediction distinguishing the phase change account from accounts that treat cognitive complexity as uniformly beneficial or uniformly limiting for meaning.

9.3 Falsification criteria

The hypothesis is falsifiable at multiple levels. The three-signature temporal model generates specific neural predictions distinguishing C-operative genuine phase change from C-operative confabulation from C-absent model updating. The eta ratio distinguishes genuine from defensive Cl and from Bayesian updating in formally testable terms. The EI(macro) greater than EI(micro) condition is calculable and falsifiable. The metacognitive override prediction is testable through intervention design.

Specific disconfirmation: an experience that produces all three simultaneously would require significant revision: subjectively felt as deeply meaningful, rho decreased across the full domain under novel data with negative eta, and yet no durable autonoetic traces and no extended episodic encoding signature.

9.4 The relationship to existing theories

Susan Wolf (2010) requires both subjective attraction and objective worth. Subjective attraction is c approaching m with rho increasing toward 1. Objective worth is the structural condition: saddle-node reached, generative Cl formed with negative eta, that distinguishes genuine transformation from sophisticated false closure.

Victor Frankl (1946) was engineering phase change three under conditions that forced the saddle-node absolutely. His clinical method replicates those structural conditions therapeutically, primarily through the metacognitive override pathway: confronting patients with what they cannot close.

Lakatos (1978) on progressive versus degenerative research programs provides the formal precedent for the generative versus defensive Cl distinction and the eta ratio.

The Aha moment literature (Jung-Beeman et al., 2004) provides the neural signature for the saddle-node event. The post-encoding consolidation literature (Tambini and Davachi, 2019; Tambini et al., 2020) provides the empirical grounding for the autonoetic integration signature.

9.5 The CHAMP implication

The Cognitive Home Assessment and Monitoring Platform (Dietz and Mueller, in preparation) assesses cognitive health longitudinally. Episodic memory assessment is assessment of the meaning-integration scaffold. A decline signals that phase changes are becoming harder to hold across time, not necessarily harder to have.

The three pathways have direct clinical implications. Natural exhaustion becomes harder as Cl repertoire shrinks with cognitive decline. Deliberate repertoire reduction becomes harder as metacognitive capacity declines. Metacognitive override becomes harder as executive function declines. The EASE library of evidence-based activities may

work through all three pathways depending on the individual's current capacity. Clinical autonarrative design for CHAMP should be calibrated to the patient's current pathway accessibility as measured by rho, eta, episodic memory function, and metacognitive capacity.

10. Conclusion

We have proposed that meaning is a saddle-node bifurcation of c at m, occurring through one of three pathways (natural exhaustion, repertoire reduction, or metacognitive override) in a system meeting EI(macro) greater than EI(micro), producing reorganization into a generative Cl that reduces R systematically across the domain, measured by decreased rho and negative eta under novel data.

Seven rounds of iterative formalization have progressively strengthened the argument. The mechanism is now specified as saddle-node bifurcation driven by R accumulation interpreted as free-energy gradient descent: deterministic, not teleological. The metacognitive override pathway is formalized as a higher-order meta-Cl, resolving the homunculus problem structurally. The EI(macro) greater than EI(micro) threshold is specified as an ontic property of intrinsic coarse-graining through Cl formation, resolving the observer-dependency problem. The bright line between Bayesian updating and phase change is drawn by the scope of Cl grammar reorganization and the eta signature, illustrated by the duck-rabbit case. The temporal model distinguishes three signatures with different neural substrates and timescales.

The core claim has held through seven rounds of review. The bridge between phenomenology and neuroscience stands. The clinical application in CHAMP is specific, targeted, and testable across three pathways. The primary empirical agenda is clearly defined.

The question meaning asks of any finite consciousness is the same question m always asks of c: what happens when R accumulation deforms the current basin to its saddle-node, whether by exhaustion, reduction, or voluntary surrender of available closures, and the gap becomes structurally impossible to close from within the current grammar?

The answer, when the conditions are right and C is operative, is meaning.

References

Benedetti, F. (2009). Placebo Effects: Understanding the Mechanisms in Health and Disease. Oxford University Press.

Chalmers, D. J. (1995). Facing up to the problem of consciousness. Journal of Consciousness Studies, 2(3), 200-219.

Damasio, A. (1994). Descartes' Error: Emotion, Reason, and the Human Brain. Putnam.

Dietz, C. F. (2026). Consciousness, Closure, and the Cosmos: A consciousness-first research posture via nested closure regimes. Version 3.3.

Frankl, V. E. (1946). Man's Search for Meaning. Beacon Press.

Friston, K. (2010). The free-energy principle: A unified brain theory? Nature Reviews Neuroscience, 11(2), 127-138.

Gazzaniga, M. S. (1967). The split brain in man. Scientific American, 217(2), 24-29.

Gazzaniga, M. S. (2011). Who's in Charge? Free Will and the Science of the Brain. Ecco Press.

Hoel, E. P. (2017). When the map is better than the territory. Entropy, 19(5), 188.

Jung-Beeman, M., et al. (2004). Neural activity when people solve verbal problems with insight. PLOS Biology, 2(4), e97.

Kaptchuk, T. J., et al. (2015). Open-label placebo treatment in chronic low back pain. Pain, 157(12), 2766-2772.

Kuhn, T. S. (1962). The Structure of Scientific Revolutions. University of Chicago Press.

Lakatos, I. (1978). The Methodology of Scientific Research Programmes. Cambridge University Press.

LeDoux, J. (2019). The Deep History of Ourselves. Viking.

Nagel, T. (1974). What is it like to be a bat? The Philosophical Review, 83(4), 435-450.

Noe, A. (2009). Out of Our Heads: Why You Are Not Your Brain. Hill and Wang.

Rosas, F. E., et al. (2020). Reconciling emergences: An information-theoretic approach to identify causal emergence in multivariate data. PLOS Computational Biology, 16(12), e1008289.

Tambini, A., and Davachi, L. (2019). Awake reactivation of prior experiences consolidates memories and biases cognition. Trends in Cognitive Sciences, 23(10), 876-890.

Tambini, A., et al. (2020). Causal contribution of awake post-encoding processes to episodic memory consolidation. Current Biology, 30(18), 3533-3543.

Taylor, C. (1989). Sources of the Self: The Making of Modern Identity. Harvard University Press.

Tononi, G. (2004). An information integration theory of consciousness. BMC Neuroscience, 5, 42.

Tulving, E. (1985). Memory and consciousness. Canadian Psychology, 26(1), 1-12.

Tulving, E. (2002). Episodic memory: From mind to brain. Annual Review of Psychology, 53, 1-25.

Wolf, S. (2010). Meaning in Life and Why It Matters. Princeton University Press.

Yalom, I. D. (1980). Existential Psychotherapy. Basic Books.