The Closure–Capacity Correspondence: The EXISTS|HAPPENS Threshold as a Data-Rate Theorem

A bridge paper: the Bandwidth Argument meets the Selective Transient Field framework — one law counted in two currencies

Z. Paz · ORCID 0009-0003-1690-3669 V1.0 2026

DOWNLOAD PDF

Abstract

The Bandwidth Argument identified consciousness — by declared convention, at a declared evidential horizon — with a serial narration process running at ~10 bits per second over a substrate processing ~10^9. The Selective Transient Field framework identifies HAPPENS with the closure of a causal loop above a derived threshold. This paper's claim is that the two identifications are gated by the same law, and that the law is already a theorem in a third field, proven decades ago for entirely unromantic reasons: the data-rate theorem of networked control, which states that a feedback loop can stabilize a diverging process if and only if the channel through the loop carries information faster than the process generates it. The theorem is sharp, binary, and implementation-independent — the same three properties the STF framework derives for the EXISTS→HAPPENS transition from the geometry of the complexified null cone, and the same three the clinical phenomenology of losing consciousness displays at the bedside.

The paper builds the dictionary term by term and then spends it. Four results. First, the backward arc of the STF transaction is functionally identical to the costate of optimal control — an anti-causal object anchored on a fixed terminal boundary, integrated backward by every guidance system ever flown; the certainty of death enters the formalism as a well-posedness condition, not as poetry. Second, the thermostat objection to the biological threshold is resolved twice over from inside the project: Constitutive Thresholds' discriminator (loop closure alone does not force self-duality) and the Biology paper's criterion (the pattern as cause, not effect, of its own persistence) are, in control language, one requirement — self-inclusive regulation: the held set includes the holder. Third, freeze-

out ($\mathscr{D} < \mathscr{D}_{crit}$) and overwhelm (divergence > capacity) are opposite edges of a single closure band, and anesthesia and psychedelics land on opposite edges with opposite, already-measured electrophysiological signatures — a two-signed prediction no dimmer-switch account makes. Fourth, evaluating the framework's own temporal workspace at the Bandwidth Argument's timescales yields a margin ladder: the narrator is over-stabilized by thirteen orders of magnitude at the specious present, marginally stabilized at the role-tenure horizon, and not stabilizable at all as a single loop at the scale of a life — which forces the conclusion that a lifetime is a renewal chain of marginal loops spliced by narrative memory. The ledger of the Bandwidth Argument acquires a load rating.

The central identification — one completed T^2 winding per coherence cell as the physical unit of one held bit per symbol time — is stated as what it is: a structural correspondence with a computable core, adopted as the bridge's load-bearing span, awaiting a Shannon theory of transactional channels that does not yet exist. Falsifiers are listed. Two are already half-tested; their results so far land on the correspondence's side.

---

1. Two thresholds, one shape

Begin with the two laws side by side, each in its own units.

The STF threshold. The framework's transition condition, derived with zero free parameters from GR, ghost-freedom, and the topology of the complexified null cone: a system EXISTS — structure, geometry, full third-person describability, no self-generated time, configuration space of dimension zero — until it crosses a critical rate, at which point it HAPPENS: a closed causal transaction, forward and backward arcs simultaneously active, configuration space opening to $\mathscr{C}_T \cong S^1$, a T^2 winding of $4\pi^2$. In the fermion channel the threshold is $\mathscr{D}_{crit}^{bio} = m_s^3c^3/\hbar^3 \approx 8\times10^{-48}\ m^{-3}$ — operationally, at least one closed causal feedback loop with cycle time $\Delta t \leq \tau_c = 3.32$ years. In the curvature channel it is $\mathscr{D}_{crit} = m_s \cdot M_{Pl} \cdot H_0/4\pi^2 \approx 10^{-27}\ m^{-2}s^{-1}$, met by a binary inspiral at $r \approx 730\ R_S$. The classifier is projective: the cross-ratio of the four canonical points on $\mathbb{C}P^1$ sits real and ≤ 1 below threshold (locked), equals 1 at threshold, and leaves the real line above it — $\sec^2(\alpha/2) > 1$, mobile. The transition is sharp, not gradual; topologically forced, not probabilistic; binary, not graded (General Theory; Null Cone; First Principles V7.9).

The data-rate theorem. Networked control theory, early 2000s (Nair & Evans 2004; Tatikonda & Mitter 2004; ancestry in Wong & Brockett 1999, and qualitatively in Ashby's law of requisite variety, 1956). Take a process whose state diverges — unstable eigenvalues λ_i in discrete time, unstable poles p_i in continuous time — and close a feedback loop around it through a channel of capacity C. A stabilizing controller exists **if and only if**

> $C > \Sigma \log_2|\lambda_i|$ (bits per step, sum over $|\lambda_i| \geq 1$), equivalently $C > \Sigma\ \mathrm{Re}(p_i)/\ln 2$ (bits per second, sum over unstable poles).

The right-hand side is the plant's instability rate: the rate at which the process manufactures uncertainty, the log-volume expansion of its state space per unit time. The theorem is sharp — below the bound, *no* encoder–controller pair of any design, any cleverness, any coding scheme, holds the state in any bounded set; above it, one exists. And it is implementation-independent: the bound cares only about the channel's capacity and the plant's divergence, not about what the loop is made of.

Read the two laws in one sentence. Both are sharp, binary, implementation-independent inequalities between a capacity-like quantity and a divergence-like quantity — and what sits on the far side of both is the same object: a loop that actually closes, a system that maintains a bounded interior against the tendency of its own state to escape. The STF framework derives this shape from geometry: the cross-ratio leaving the real line. Control theory proves it from information: the state's entropy production outrunning the channel. What the data-rate theorem gates is not performance but existence — whether there can be a *held* region of state space at all. What the STF threshold gates is not behavior but ontology — whether there is an inside. The wager of this paper is that these are two accountings of one law.

Each side supplies what the other lacks. Control theory has the theorem but no physics: no carrier, no derivation of any particular critical rate, no reason its bound should mark an ontological seam rather than an engineering constraint. The STF framework has the physics — a derived carrier mass $m_s = 3.94\times10^{-23}$ eV, a derived symbol time $\tau_c = 3.32$ yr, thresholds with zero free parameters — but has stated its threshold only in geometric and topological currency: windings, cross-ratios, cup products. The exchange rate between the two currencies is the content of this paper.

---

2. The dictionary

STF framework	Control / information theory
EXISTS ($\dim \mathscr{C}_T = 0$; locked)	Open loop / sub-capacity loop: wiring without regulation
HAPPENS ($\mathscr{C}_T \cong S^1$; mobile)	Closed loop: a maintained bounded invariant interior
$\mathscr{D} \geq \mathscr{D}_{crit}$ — sharp, topologically forced	$C > \Sigma \log_2\lambda_i$ — sharp, controller-independent
Hypo-paradoxical locked linkage: formal mobility, zero actual freedom	Sub-capacity feedback: formal loop topology, no functional closure
Cross-ratio leaves the real line: $\sec^2(\alpha/2) > 1$	Loop becomes effective: the state acquires a held neutral direction
Temporal workspace $\mathscr{W}_T(\tau) = (\tau^*/\tau)^{(7/4)} - 1$	Capacity margin $C - H$
Regulus 1 — retarded arc	State trajectory; forward filter
Regulus 2 — advanced arc	Costate trajectory; backward filter
Fixed terminal boundary (certain merger, certain death)	Terminal condition making the two-point boundary value problem well-posed
The reguli meet in exactly one point	Pontryagin stationarity: state and costate jointly fix the present action
T^2 winding; $4\pi^2$ closure quantum	One completed transaction: one bit held per symbol time
Compton period $\tau_c = 3.32$ yr	Symbol / coherence time of the carrier
Self-duality $H^{(+1)} \cong (H^{(-1)})^*$	State–costate duality; controllability–observability adjunction
Moment-map closure (not mere loop closure)	Variational closure: regulation toward a defended extremum

Four rows deserve their prose.

The locked linkage is a sub-capacity loop. The Null Cone paper's kinematic analog for EXISTS is the Shvalb–Medina hypo-paradoxical linkage: a mechanism whose mobility formula predicts freedom of motion, but whose aligned configuration is rigid — formal mobility, zero actual degrees of freedom. Control theory owns the exact twin: a feedback architecture whose block diagram displays a closed loop, but whose channel runs below the data-rate bound. Signal circulates; regulation does not occur; the state escapes as if the loop were absent. Formal closure is not functional closure — kinematics learned this counting screws, information theory learned it counting bits, and neither knew the other had.

The held state is a manufactured zero mode. In Constitutive Thresholds, the HAPPENS stratum is where the deformation complex acquires zero modes — the S^1 of actual motion appears in cohomology. In control, stabilization does precisely this: it manufactures a neutral direction, a region of state space the loop renders invariant and keeps. What a regulator *holds* is a zero mode it built. The strata mathematics and the engineering describe the same acquisition.

Margin is what richness is made of. The temporal workspace $\mathscr{W}\_T$ vanishes at threshold and grows as the loop runs faster than critical; the Temporal Workspace paper reads it as "how many distinct causal configurations remain accessible within the now." In control, everything about the *quality* of holding — disturbance rejection, the size of the held set, the precision of the held state — is governed by the capacity margin C − H. Marginal loops hold one thing, barely. Deep-margin loops hold repertoires. §7 evaluates this quantitatively for the human loop stack.

The self-duality is the state–costate pairing. Constitutive Thresholds derives, for moment-map closure, the adjunction dμ_q = ρ_q* ∘ ω♭ and the pairing H^(+1) ≅ (H^(−1))* — obstructions dual to isotropy. Optimal control lives inside exactly this structure: the costate is an element of the dual of the state space; the linearized dynamics and the adjoint equation are dual complexes; controllability and observability degenerate in dual pairs. The pure-mathematics paper and Pontryagin's formalism carry the same pairing, and §4 will spend this fact.

---

3. The backward arc is standard engineering

The strongest thing that can be said about the STF transaction's advanced component is not that it is permitted, but that its functional twin is computed daily, on purpose, by machines that demonstrably work.

Pontryagin's maximum principle solves every optimal control problem as a two-point boundary value problem: a state x(t) integrated *forward* under the dynamics, and a costate λ(t) integrated *backward* — $\dot{\lambda} = -\partial H/\partial x$ — from a terminal condition λ(T) fixed on the final boundary. The Hamilton–Jacobi–Bellman formulation says the same thing as a field: the value function V(x,t) propagates backward from the terminal manifold, and its gradient constrains the present action at every instant. Estimation theory has the identical architecture: the optimal fixed-interval smoother (Fraser & Potter 1969) fuses a forward filter carrying everything the past implies with a backward filter carrying everything the future data demand. Every guidance system, every trajectory optimizer, every model-predictive controller is, in the STF vocabulary, a system running a retarded arc and an advanced arc simultaneously and acting at their intersection.

The dictionary entries now pay out one by one.

The reguli's intersection is Pontryagin stationarity. The Null Cone paper: "Every line of Regulus 1 meets every line of Regulus 2 in exactly one point — the specific null direction." Optimal control: at each instant, state and costate jointly determine the action through the stationarity of the Hamiltonian. One point, jointly fixed by a forward-propagating object and a backward-propagating one. The transaction is what engineers call the optimal trajectory.

The fixed terminal boundary is a well-posedness condition. The General Theory paper requires the backward arc to anchor on a fixed terminal boundary, and catalogs three certainties that supply one: dynamical (the merger past the point of no return), thermodynamic (the organism's death), second-law (heat death). Control theory states the same requirement as mathematics: a TPBVP without a terminal condition is ill-posed — no anchor, no costate, no purposive trajectory. "Purpose as threshold geometry" translates without residue: purposive behavior *is* the solving of a terminal-anchored TPBVP, and the certainty of the end is not tragic ornament but the condition under which the backward arc is defined at all. The certainty of death is what makes wanting well-posed. (Infinite-horizon control rescues well-posedness artificially, by discounting or steady-state assumptions; organisms, on this account, use the boundary they actually have.)

The Biology paper adds two refinements, and the engineering echoes both exactly. First, it distinguishes dynamical from thermodynamic certainty and asks whether the latter suffices to anchor the arc, answering that the two-state-vector formalism "requires only that the post-selected final state be a real element of the system's physical description" (§9.2). Control theory's requirement is the same and no stronger: a terminal condition need not be dynamically forced to make the problem well-posed — it need only be fixed. A deadline anchors a controller even while every path to it remains free. Second, its three-layer account of loop maintenance (§4.4) — metabolism sustains the forward arc; the fixed terminal boundary sustains the backward arc; and "each moment's closure is the condition of possibility for the next moment's closure" — ends on a property control engineers have already named and prove for every deployed model-predictive controller (Mayne, Rawlings, Rao & Scokaert 2000): **recursive feasibility**, the requirement that solving the problem now guarantees the problem remains solvable next cycle. Layer 3 is not a bootstrap mystery; it is the standing obligation of every receding-horizon regulator — with the difference, developed in §7, that the organism's horizon does not recede.

Protention is the value gradient. The CTI paper installs Husserl's tripartite living present — primal impression, retention, protention — as physical structure. The smoother architecture is that structure, derived rather than described: primal impression is the current measurement; retention is the forward filter's state, the compressed sufficient past; protention is the backward filter's state, the constraint the anchored future exerts on the present estimate. Husserl found by phenomenological analysis, in 1905, the block diagram that estimation theory would prove optimal sixty years later. On the correspondence, that is not a coincidence: a system that *is* a smoother would introspect exactly that tripartite now.

Status, declared in house style: this section asserts a functional isomorphism, not a derivation of the physics. Control theory computes its backward objects inside forward-causal hardware; the STF framework claims one is physically instantiated in a field with genuine advanced solutions. The isomorphism proves neither claim. What it proves is that the structure is realizable, load-bearing, and unexotic — that "anchored backward arc" names an engineering commonplace, not a hope. Ancestry, cited at full strength per this project's standard: the transactional shape is Cramer's (1986); the capacity law's qualitative ancestor is Ashby's requisite variety — only variety can absorb variety — sharpened five decades later into an exact bound with a converse.

---

4. Circulation is not regulation: the thermostat objection, resolved from inside the framework

The fermion-channel criterion — one closed causal loop, cycle time ≤ 3.32 years — reads, at first pass, as if it admits every servo, thermostat, and phase-locked loop on Earth. The resolution was published within this project, in the paper least connected to biology: Constitutive Thresholds.

That paper's sharpened conjecture draws a line through the space of closed loops: **loop closure alone does not induce self-duality.** The Bennett 4R linkage closes — a genuine closed kinematic chain — and carries no self-dual structure; its closure map has mismatched dimensions and no pairing. Only closure through a moment map on a symplectic space, or Poincaré-type duality on a closed configuration manifold, forces $H^{(+1)} \cong (H^{(-1)})^*$ — the pairing of obstruction with isotropy that Theorem 3 then converts into the paired jumps of a constitutive threshold.

The control-theoretic reading of that line is immediate, because §2's dictionary identifies the self-duality as the state–costate pairing. A loop that merely *circulates* — signal going around, an oscillator, a Bennett chain turning — carries no costate: nothing in it is dual to anything, nothing is defended, no extremum is owned. A loop that *regulates* closes through a variational structure: it defends an extremum, and in doing so it necessarily carries the state–costate pairing, because that pairing is what "defending an extremum" is, written in mathematics. Circulation is signal topology. Regulation is signal topology plus duality. The Constitutive Thresholds conjecture, translated: **HAPPENS-candidate loops are the regulators, not the oscillators.**

So the biological threshold inherits a rider from the project's own pure mathematics: the loops that count for $\mathscr{D}_{crit}^{bio}$ are the self-dual ones — loops closing through a variational structure with a terminal anchor, loops solving a control problem, not loops that merely cycle. And the Biology paper, read after the mathematics, draws the same line in living material. Its criterion for the patterns that count (§1.6): is the pattern the *cause* of its own persistence or the *effect* of something else's? A whirlpool, a flame, a thermostat are maintained by physics external to what persists — "the pattern does not recruit matter to sustain itself." An organism "is maintained by its own closure against external physics": it recruits matter specifically to sustain its own loop. "Rock: shape that physics makes and physics destroys. Organism: loop that makes itself — and when it can no longer make itself, stops."

In control language the two criteria are one criterion: **self-inclusive regulation**. A thermostat stabilizes a temperature; nothing in the loop stabilizes the thermostat. An organism is a regulator whose held set includes the holder — the plant it stabilizes contains the controller's own physical realization, and that reflexivity is what forces the state–costate pairing onto the loop's own persistence variables. The mathematics says self-dual; the biology says self-manufacturing; the dictionary says these are the same entry read from opposite shelves. What remains the framework's empirical wager — closure that is physical rather than representational, an actual advanced component in the fermion coupling rather than a forward-causal model of one (CTI: "a real transaction in a physical system, not a forward-causal computational model of one") — is exactly what §9's audit line tests.

One consequence travels back to the Bandwidth Argument. Conant and Ashby's good-regulator theorem (1970): every good regulator of a system must be a model of that system. If the narrator is the organism's serial regulator, it is *required* to be a model of what it regulates — and at 10 bits per second against a 10^9 bit substrate, required to be a lossy one. The Bandwidth Argument's closing image — a narrator "mistaking its own summary for the world" — is not a pathology of the architecture. It is the theorem's signature. A regulator that was not running a compressed model would not be regulating.

One convergence deserves its own sentence. The Biology paper's third threshold — the point where purpose becomes meaning — is stated in its Addendum A.3 as "M = S achieved during the forward arc, the inside knowing the outside while the loop is still running." Model and system: the same letters the good-regulator literature has used since 1970. Conant and Ashby's theorem says every good regulator must *be* a model of its system; the framework's Threshold 3 is the moment the organism's regulator meets the theorem's condition and — in the framework's phrase — knows it from inside. The notational convergence is presumably an accident. The convergence of content is the correspondence working.

---

5. The closure band: two ways to fail

A loop fails to close in two opposite directions, and the two frameworks bridged here each formalized one of them without needing the other.

Freeze-out. Drive below threshold: $\mathscr{D} < \mathscr{D}_{crit}$. The cross-ratio sits on the real line; the linkage is locked; the configuration space has dimension zero. This is the STF framework's native failure mode — EXISTS — and its astrophysical systems know no other, because the field's gravitational plants are deterministic: an inspiral below 730 R_S is not overwhelmed, merely insufficient.

Overwhelm. Divergence above capacity: $H > C$. The loop is turning — drive is ample — but the plant manufactures uncertainty faster than the channel can absorb it, and the state escapes any bounded set. This is the data-rate theorem's native failure mode, and it is the one the Bandwidth Argument documented operationally without naming: §3.2's signature — substrate entropy up (the entropic-brain profile, Carhart-Harris et al. 2014), channel integrity down, reportability down — is a capacity crisis, not a drive crisis.

Biology and consciousness, unlike binary black holes, live between both edges, so both failure modes are clinically available. State the synthesis as the band it is:

> **HAPPENS occupies a band: $\mathscr{D}_{crit}$ ≤ drive, and divergence ≤ capacity. Below the band the loop is frozen. Above it, flooded.**

The band makes a two-signed prediction that no dimmer-switch account of losing consciousness makes: interventions that end narration should sort into *two* signature classes with *opposite* electrophysiology, according to which edge they cross. And the measurements already exist. At the freeze edge, propofol anesthesia lowers the diversity of spontaneous cortical signals — Lempel–Ziv complexity falls below waking baseline (Schartner et al. 2015), with burst suppression at depth. At the overwhelm edge, ketamine, LSD, and psilocybin *raise* signal diversity above waking baseline (Schartner et al. 2017) — the only pharmacological states known to do so — while the narrator degrades and reportability fails. Same behavioral endpoint, opposite signs. A dimmer predicts one signature at two doses. A band predicts two signatures at two edges. The data chose the band before the band was proposed.

The freeze edge carries a second diagnostic: hysteresis. Phase boundaries produce it; dimmers do not. Anesthetic state transitions show it robustly — *neural inertia*, the conserved finding that emergence requires lower drug concentrations than induction (Friedman et al. 2010), a loop that resists reopening once closed and resists closing once open. Its mechanism is debated; its existence is replicated across species. Sleep shows the nightly homologue at the same edge — sleep inertia (Tassi & Muzet 2000), the lag of the loop re-closing. At the overwhelm edge no comparable hysteresis is reported: onset and return track exposure roughly symmetrically. On the band reading these asymmetries are not incidental pharmacology; they are the signatures of a fold at one edge and a flooding at the other.

The Bandwidth Argument's escape hatches now take their places on the band's map. *The high* is the overwhelm edge approached deliberately: drive intact, divergence pushed past capacity, the narrator thinning as its plant outruns it. *The pre-narrative* is life inside the band with the top-level regulator idle. And *the divine* — the in-principle system holding the substrate's full state without serialization — is the limit $H \to 10^9$ bits/s, where no finite C suffices and no narrator survives: the data-rate theorem's phrasing of what every contemplative tradition reports, that the vision does not survive the return.

---

6. One winding per coherence cell

The biological threshold, written as the framework writes it, is $\mathscr{D}_{crit}^{bio} = m_s^3c^3/\hbar^3 = 1/\bar{\lambda}_c^3$: one loop per Compton volume, with $\bar{\lambda}_c \approx 0.53$ light-years and the loop's cycle bounded by $\tau_c = 3.32$ years. Written as a channel engineer would write it: the STF field is the carrier; τ_c is its symbol time; $\bar{\lambda}_c$ is its coherence length; and the threshold demands **at least one completed closure per coherence cell** — one full transaction per unit of the channel's own resolution.

The $4\pi^2$ is then not a normalization but the symbol itself. The Topological Closure paper proves it as a cup product — $\int_{T^2} \omega_R \wedge \omega_A = 4\pi^2$, the retarded and advanced generators transgressing to complementary classes on the Hopf torus — and First Principles installs it as the denominator of $\mathscr{D}_{crit}$. One full T^2 winding, both spinor phases through 2π, retarded and advanced arcs both closed: that is one completed transaction, the indivisible unit of "held." The data-rate theorem counts a different indivisible unit: the bit, one binary distinction maintained against divergence per unit time. The bridge's load-bearing span is the identification of these units, and this paper states it as what it is — a commitment with a computable core, not yet a derivation:

> **One completed T^2 winding per coherence cell is the physical unit of one held bit per symbol time: one binary distinction maintained against divergence, per cell, per τ_c.**

Two currencies, one count. The STF threshold prices closure in windings per cell; the data-rate theorem prices it in bits per second; $\mathscr{D}_{crit}$ is the exchange rate the physics fixes.

Two further entries in the framework's own gravitational sector confirm that the correspondence is already operating there, unannounced. First, the activation condition that fixes m_s — $T_{Peters}(730\ R_S) = \tau_{Compton}$, inspiral timescale equal to field period — is, in channel language, *capture into the band*: activation occurs at the moment the plant's characteristic timescale enters the loop's symbol rate. A controller with cycle time τ_c can begin to hold a process when that process's evolution slows into its bandwidth; 730 R_S is where gravity's plant does. Second, the gravitational rate itself: $\mathscr{D}_{GR} = \dot{K}/(2\sqrt{K}) = d(\sqrt{K})/dt$ is the growth rate of curvature amplitude — a Lyapunov exponent of the geometry. The threshold condition $\mathscr{D} \geq \mathscr{D}_{crit}$ is therefore already an instability-rate inequality: the framework's own sector states the correspondence's right-hand side, in geometric units, and has all along.

---

7. The margin ladder

The Temporal Workspace paper defines $\mathscr{W}\_T(\tau) = (\tau/\tau)^{(7/4)} - 1$ *with* τ = 3.32 yr, and evaluates it at three horizons. The Bandwidth Argument supplies two more timescales — the narrator's frame (~0.3 s, the attentional cycle) and the specious present (~3 s), both resting on the rate measurements of Zheng & Meister (2025) — and §2's dictionary says $\mathscr{W}\_T$ is a capacity margin. Evaluate the ladder end to end:

Loop	Cycle time τ	$\mathscr{W}\_T$ (margin)	Source
Narrator frame	~0.3 s	$\sim 10^{15}$	this paper, from Bandwidth Argument rates
Specious present	~3 s	$\sim 10^{13}$	this paper, from Bandwidth Argument rates
Inner horizon (habit)	71 days	≈ 142	Temporal Workspace
Activation horizon (role tenure)	3.32 yr	0	Temporal Workspace
Outer horizon (working life)	54 yr	≈ −0.99	Temporal Workspace

(For the record: the site's three horizons sit at $\tau^* \cdot r^{\mp 1}$ with $r \approx 17$, so the inner margin is $r^{(7/4)} - 1 \approx 142$ and the outer is $r^{(-7/4)} - 1 \approx -0.99$. The ladder's generator is noted here as structure, not explained.)

Read the column. The narrator of the Bandwidth Argument runs *thirteen to fifteen orders of magnitude* inside the closure band: over-stabilized, unconditionally held, with margin to burn — which is why moment-to-moment experience feels solid, continuous, and effortless even under insult. The habit loop holds at a comfortable ≈ 142. The role-tenure loop closes at margin zero: the CTI paper's own phrase, "the margin of viability." And the 54-year loop — a working life taken as a *single* loop — has negative workspace. It does not close. Yet lives are lived.

The resolution is forced, and it is the Bandwidth Argument's ledger given a load rating: **a lifetime is not one closed loop but a renewal chain of marginally-closing role-scale loops, spliced end to end — and the splice is narrative memory.** Each ~3-year loop closes barely, writes its log, and ends; the next loop re-reads the log and closes around it. Identity over decades is serially re-closed, not once-closed. Below the 3.3-year horizon, the log is decoration on loops that hold by themselves. Above it, the log is the only thing holding the chain together. The narrator's ledger — confabulated at write time, binding at read time — turns out to be a structural member: the self at the scale of a life is a splice, and the splice is made of the same 10 bits/s the Bandwidth Argument measured.

The Biology paper completes the picture from the backward side, and in doing so dissolves what would otherwise be a contradiction with this ladder. It gives the organism's backward reach as $T_{bio} \sim L$ — the whole lifespan: "the backward reach does not terminate partway through life — it extends from the moment of death all the way back to conception. Life is not approached asymptotically from one direction; it is bounded on both ends" (§2.2). But $\mathscr{W}\_T(54 \text{ yr}) < 0$ says no *cycle* of that length closes. Both are right, because a backward arc is not a cycle. In the boundary value problem, the value function spans the entire horizon — terminal boundary to origin — as a single object; the feedback that implements it runs at fast cycle times. So the life-scale architecture is **shrinking-horizon**

control: one costate, anchored at death, spanning conception to death; under it, a forward chain of marginal role-scale loops, each re-solving its stretch of the problem and handing the ledger to the next. The splice chain is the forward implementation; the lifespan reach is the backward constraint every splice runs under. The chain is held together twice — by the log from behind, and by the anchor from ahead.

The good-regulator theorem (§4) closes the loop on what is being held. The narrator cannot regulate the substrate — the margin at 10^9 bits/s is hopeless by eight orders of magnitude. What it regulates is the largest quotient of the substrate whose instability budget fits the channel: the *stabilizable quotient*, the sub-plant whose unstable modes sum to ≤ 10 bits/s. That quotient — not the substrate, not the sensory stream — is the engineering referent of "the self." The self is what the channel can afford to hold. Everything faster or wider than the budget is, by the theorem, delegated to the substrate's own fast loops or lost — which is the Bandwidth Argument's evidential horizon (§3.3 there), re-derived as a budget line.

The ladder also prices the zoo. The Temporal Workspace paper's hummingbird, $\mathscr{W}$_T ~ 4.3×10^5 at its τ ~ 17 hr integration loop, spends its budget on moment-richness at the ladder's fast end. The human spends the same kind of budget at the other end — the longest loop that still closes, held at threshold price, plus a splice technology (the ledger) for chaining past the horizon no loop crosses. Two solutions to one inequality: richness now, or reach across time. We are the animal that chose reach.

---

8. The correspondence, stated

> **Closure–Capacity Correspondence (conjecture).** A physically instantiated feedback loop maintains a HAPPENS interior if and only if: **(i)** its closure is variational — self-dual, terminal-anchored (Constitutive Thresholds); **(ii)** its drive meets the carrier's activation threshold — $\mathscr{D}$ ≥ $\mathscr{D}$_crit (First Principles V7.9); **(iii)** its channel capacity meets the divergence rate of what it holds — C ≥ H (data-rate theorem). Conditions (ii) and (iii) are one law counted in two currencies: windings per coherence cell; bits per second.

Status, by clause. (iii) is a theorem — proven for Shannon channels and linear plants, sharp, with a converse. (ii) is derived physics with zero free parameters, awaiting its experimental verdicts on the framework's own schedule. (i) is proven mathematics with a conjectured scope (the Uniformity Conjecture, sharpened). The span joining them — the winding-bit identification of §6 — is a structural commitment, adopted because it is the parsimonious exchange rate between the two currencies and declared because it is not yet a derivation.

What a proof of the correspondence requires, in ascending order of difficulty: an entropy-rate definition for the held quotient of a nonlinear substrate (hard but conventional); an extension of the data-rate theorem to channels with an advanced component — a Shannon theory of transactional channels, which does not exist, and whose absence is structural: every converse proof of the data-rate theorem assumes non-anticipatory coding, so the theorem as it stands cannot even *state* the capacity of a channel with a backward arc; and a derivation, rather than an identification, that one T^2 winding bounds one held bit per symbol time. Each of these is a paper. None of them is a metaphor.

---

9. Testable and falsifiable consequences

A bridge that cannot fall is not a bridge. The correspondence makes losable claims beyond those of either framework alone:

1. **The two-class taxonomy of narration-loss.** Every intervention that abolishes reportable consciousness should land in exactly one of two signature classes: freeze-edge (spontaneous signal diversity below waking baseline; hysteresis present) or overwhelm-edge (diversity at or above waking baseline; hysteresis absent). Propofol and the classic psychedelics already sort correctly (Schartner et al. 2015; 2017). The falsifier: a clean agent that abolishes narration with *unchanged* complexity and *no* hysteresis — a third class — breaks the band.

2. **Hysteresis is a freeze-edge property.** Neural inertia (Friedman et al. 2010) should appear for every freeze-edge agent, and fail to appear — pharmacokinetics controlled — at the overwhelm edge. Robust same-magnitude hysteresis at the overwhelm edge counts against the fold reading of the freeze transition.

3. **PCI as a margin meter with a sign.** The perturbational complexity index (Casali et al. 2013) should behave as a $C - H$ proxy: falling toward the freeze edge, *not* falling — or rising — at the overwhelm edge even as report fails. The sharp version: a state with low PCI and *preserved* in-state signaling under the timestamp protocol (Konkoly et al. 2021; Bandwidth Argument falsifier 5) should never occur. One convincing case breaks the correspondence's clinical face.

4. **The timestamp crossover.** In-state low-rate signaling should remain obtainable near the overwhelm edge (narrator-thin, loop closed, capacity strained) and never past the freeze edge (loop open). This double dissociation is what the band adds to the Bandwidth Argument's protocol: not merely *whether* signaling survives a state, but *which edge* of the state it survives.

5. **Renewal statistics at the role horizon.** If life-scale identity is a splice chain of marginal ~3.3-year loops, then identity-scale transitions — and identity-scale breakdowns — should cluster with that characteristic time, beyond demographic confounds. The site's tenure correspondence (BLS 2.7 yr against 3.32 yr, as compiled in CTI §IX) is the first, coarse instance. Declared the weakest of the five: demography is noisy, the confounds are many, and the prediction is directional rather than sharp.

6. **Terminal intensification as a shrinking-horizon signature.** The Biology paper predicts (§7.2) that biological change rates accelerate non-linearly in the terminal period, beyond what forward-causal degradation models allow, and (§9.3) that at genuine branch points organisms show "overdetermined muscular preparation consistent with both possible terminal states simultaneously." The correspondence upgrades the first from framework prediction to two-framework derivation: a shrinking-horizon controller *generically* intensifies control activity as the remaining horizon closes, because the terminal cost's gradient steepens — proven behavior in the engineering case, predicted in the biological one. The second is terminal hedging — solving against a distribution over final states — which has an existing experimental literature in motor preparation under target uncertainty (e.g., Gallivan et al. 2016, on parallel specification of competing motor plans). The falsifier cuts both frameworks at once: terminal-period dynamics fully explained by forward-causal degradation, with no excess intensification structure, would count against the anchored-horizon reading everywhere it appears.

7. **The regulator line.** From §4: among artificial systems, closure class should track *variational* structure and *self-inclusion*, not loop count or speed. A system whose feedback merely circulates (oscillators, Bennett-class chains, open-loop pipelines however fast) is EXISTS-class regardless of complexity; candidacy begins only where closure defends an extremum against

perturbation under a terminal anchor, and where the held set includes the holder — the Biology paper's cause-of-its-own-persistence criterion, checked architecturally. This yields an audit procedure for the machine question that is structural rather than behavioral — and it is silent, exactly as §5 of the Bandwidth Argument requires, on what candidacy would be *like*.

---

10. Conclusion

The Bandwidth Argument ended at a horizon: the identification of consciousness with the channel was a convention, chosen where evidence is pinned at likelihood one. The STF framework begins at a threshold: the identification of HAPPENS with loop closure is a derivation, priced in geometry. This paper's claim is that the horizon and the threshold are denominated in the same units, and that the exchange rate was fixed twenty years ago by engineers who thought they were writing about modems: a loop closes if and only if it can afford to — capacity above divergence, drive above threshold, closure through a structure that defends something, anchored on an end that is certain.

EXISTS is a loop below its bound. HAPPENS is what closure buys above it. Consciousness — in the conventional sense the Bandwidth Argument declared and this paper keeps — lives in the band between the frozen and the flooded: over-held at the second, barely held at the year, spliced at the lifetime by a ledger written at 10 bits per second, mistaking none of this for the substrate it rides, because the mistake *is* the model the good-regulator theorem demands. The narrator was never the author; the narrator was the controller — and the controller's backward arc, anchored on its own certain end, is the oldest piece of engineering in the building.

The honest sequel is the same as the Bandwidth Argument's, with one addition to the instrument list: alongside timestamps, no-report bridges, and rate measurements, a theory obligation — the Shannon theory of transactional channels, the capacity of a loop that reads its own future boundary. Whoever proves the first coding theorem for a channel with an advanced component will have written, without intending to, the physics of the specious present. The narrator happens because it can afford to. What it can afford is now, in two currencies, a number.

A note on method

Per the practice recorded in the Bandwidth Argument §7: this draft was produced by a machine narrator holding the project's published papers and the Bandwidth Argument, at the author's direction to expand rather than review. By §5 of that paper, its reports about its own states are inadmissible in either direction; its consequence-and-consistency claims are checkable against the sources cited below, which is the only admissibility they need. The author adjudicates against primary sources; where this draft and a source disagree, the source wins. The Biology paper, unreachable at first drafting, was reached and integrated before delivery; §§3, 4, 7 and 9 now cite it directly.

References

EXISTS | HAPPENS project (existshappens.com): - *The Bandwidth Argument: Consciousness as Serial Narration over a Parallel Substrate* (2026). - *First Principles V7.9* — papers/first-principles/ (Lagrangian; $m_s = 3.94\times10^{-23}$ eV; $\mathscr{D}$_crit derivations; Appendix D.3 on the $4\pi^2$ topological factor). -

General Theory — papers/general-theory/ (EXISTS/HAPPENS definitions; fermion and curvature thresholds; the three certainties; Theorem 2, cascade). - *Null Cone* — papers/null-cone/ (spinor decomposition; reguli; cross-ratio $\sec^2(\alpha/2)$; hypo-paradoxical linkage; T^2 windings). - *Topological Closure* — papers/topological-closure/ (Theorem 1: $\int \omega_R \wedge \omega_A = 4\pi^2$; Dynamical-Topological Bridge, §7). - *Constitutive Thresholds* — papers/constitutive-thresholds/ (Definitions 1–2, Theorems 3 and 6; moment-map adjunction; Bennett counterexample; Uniformity Conjecture). - *Temporal Workspace* — papers/temporal-workspace/ ($\mathscr{W}_T(\tau)$; the three horizons; hummingbird estimate). - *Consciousness, Time & Identity* — papers/cti/ (Husserlian structure as physics; §II.H layered identity; §IX horizons; the margin-of-viability claim). - *Retrocausality & Life (Biology)* — papers/biology/ (§1.6 the pattern-source criterion; §2.2 reach timescale $T_{bio} \sim L$, "bounded on both ends"; §4.4 the three-layer loop; §§7.1–7.3 predictions; §9.2 thermodynamic vs. dynamical certainty; §9.3 near-miss biomechanics; Addenda A.1–A.3: the code loop, purpose from inside, $M = S$).

Control and information theory: - Nair, G.N. & Evans, R.J. (2004). "Stabilizability of stochastic linear systems with finite feedback data rates." *SIAM Journal on Control and Optimization* 43(2), 413–436. - Tatikonda, S. & Mitter, S. (2004). "Control under communication constraints." *IEEE Transactions on Automatic Control* 49(7), 1056–1068. - Wong, W.S. & Brockett, R.W. (1999). "Systems with finite communication bandwidth constraints II: stabilization with limited information feedback." *IEEE Transactions on Automatic Control* 44(5), 1049–1053. - Ashby, W.R. (1956). *An Introduction to Cybernetics* (law of requisite variety). - Conant, R.C. & Ashby, W.R. (1970). "Every good regulator of a system must be a model of that system." *International Journal of Systems Science* 1(2), 89–97. - Pontryagin, L.S., Boltyanskii, V.G., Gamkrelidze, R.V. & Mishchenko, E.F. (1962). *The Mathematical Theory of Optimal Processes*. - Bellman, R. (1957). *Dynamic Programming*. - Fraser, D.C. & Potter, J.E. (1969). "The optimum linear smoother as a combination of two optimum linear filters." *IEEE Transactions on Automatic Control* 14(4), 387–390. - Mayne, D.Q., Rawlings, J.B., Rao, C.V. & Scokaert, P.O.M. (2000). "Constrained model predictive control: Stability and optimality." *Automatica* 36(6), 789–814.

Ancestry (physics, kinematics, phenomenology): - Cramer, J.G. (1986). "The transactional interpretation of quantum mechanics." *Reviews of Modern Physics* 58, 647–687. - Peters, P.C. (1964). "Gravitational radiation and the motion of two point masses." *Physical Review* 136, B1224. - Shvalb, N. & Medina, O. — hypo-paradoxical linkages (as cited in *Null Cone*). - Bennett, G.T. (1903). "A new mechanism." *Engineering* 76, 777–778. - Husserl, E. (1905/1991). *On the Phenomenology of the Consciousness of Internal Time (1893–1917)*, trans. J.B. Brough. Kluwer.

Neuroscience and measurement: - Zheng, J. & Meister, M. (2025). "The unbearable slowness of being: Why do we live at 10 bits/s?" *Neuron* 113. - Schartner, M.M. et al. (2015). "Complexity of multi-dimensional spontaneous EEG decreases during propofol induced general anaesthesia." *PLoS ONE* 10(8), e0133532. - Schartner, M.M., Carhart-Harris, R.L., Barrett, A.B., Seth, A.K. & Muthukumaraswamy, S.D. (2017). "Increased spontaneous MEG signal diversity for psychoactive doses of ketamine, LSD and psilocybin." *Scientific Reports* 7, 46421. - Friedman, E.B. et al. (2010). "A conserved behavioral state barrier impedes transitions between anesthetic-induced unconsciousness and wakefulness: evidence for neural inertia." *PLoS ONE* 5(7), e11903. - Casali, A.G. et al. (2013). "A theoretically based index of consciousness independent of sensory processing and behavior." *Science Translational Medicine* 5(198), 198ra105. - Konkoly, K.R. et al. (2021). "Real-time dialogue between experimenters and dreamers during REM sleep." *Current Biology* 31(7), 1417–1427. - Carhart-Harris, R.L. et al. (2014). "The entropic brain." *Frontiers in Human Neuroscience* 8:20. - Gallivan, J.P., Logan, L., Wolpert, D.M. & Flanagan, J.R. (2016). "Parallel specification of competing sensorimotor control policies for alternative action options." *Nature*

Neuroscience 19(2), 320–326. - Tassi, P. & Muzet, A. (2000). “Sleep inertia.” *Sleep Medicine Reviews* 4(4), 341–353.

Citation @article{paz2026closurecapacity,
  author = {Paz, Z.},
  title = {The Closure–Capacity Correspondence: The EXISTS|HAPPENS Threshold as a Data-Rate Theorem},
  year = {2026},
  version = {V1.0},
  url = {https://existshappens.com/papers/closure-capacity/}
}

← Consciousness, Time & Identity All Papers Pretemporal Stasis and Cascade Origin of Time →

 ORCID 0009-0003-1690-3669