THE TWO ANCHORS

Ordered Attrition, Bounded by Resilience
— A Framework for Alzheimer's Disease, with Its Grades, Provenance, and Falsification Tests Fixed in Advance

The Charter · The Grade · Anchor One
Anchor Two · The Span · The Ledger

Prepared under the Organic Network Synthesis methodology

AdultCognitiveDisease.com

Benjamin Aaron Gustafsson

July 2026

Written after an audit found that this corpus had imported its central architecture, rescored itself against it, and cited the resulting fit as evidence the architecture was correct. This framework is built so that it cannot repeat that: grades are defined before claims, every claim declares its provenance, the corpus's own coherence is inadmissible as evidence, and what would refute each claim is written down first. Every load-bearing claim rests on work this programme did not do.

"A lesion tells us only that a pathway can be broken. A reprieve tells us that a pathway can be held — and it tells us so with the one authority human genetics rarely commands: a person in whom the change came first, the outcome after, and time itself forbids the confusion of the two."

— *The Architecture of Resistance*, July 2026

The strongest evidence this corpus holds was never its own. It belongs to the neuropathology of ordered attrition and to the human genetics of those who escaped — and a framework is only as honest as its willingness to rest there.

THE SHAPE OF THE FRAMEWORK

Anchored where human evidence is strong; hypothesis where it is not

Anchor One — Ordered Attrition *(decades-long, anatomically ordered, selective — and not this programme's discovery)*

Anchor Two — Resilience *(pathology and cognition dissociate; protection is reverse-causation-free evidence)*

The Span — Mechanism → held as falsifiable hypothesis, barred from carrying weight

A model-blind audit of June 2026 found that every claim surviving blind replication was a *count*, and every claim that diverged was a *mechanism* or a *reversal* — and that a rival model fared no better. The descriptive layer of this field is human-robust for everyone; the causal layer is model-dominated and trial-failed for everyone. This framework therefore places its weight at the two ends where human evidence is strongest and treats the causal middle as the work still to be done. It claims less than its predecessor. What it buys is that its load-bearing claims are ones this programme could not have manufactured, and that the two claims evidence has already defeated are printed in the same table, at the same size, as the rest.

Contents

Abstract

This paper states a framework for Alzheimer's disease and is built so that it cannot confirm itself. It follows an audit of the programme that produced it, which found that the corpus's central architecture had been imported from a prior whitepaper, that the corpus had then been rescored against that architecture, and that the resulting fit had been cited as evidence the architecture was correct. Whatever else it does, a successor framework must not be able to repeat that.

The framework has a simple shape, and the shape is dictated by where human evidence actually is. A model-blind audit conducted in June 2026 — five reviewers given theory-free briefs, none told that any architecture existed — returned a pattern that has not been given the weight it deserves: every claim that survived blind replication was a **count**, and every claim that diverged was a **mechanism** or a **reversal**. A second blind run against a rival connectome-propagation model returned the same asymmetry, and on its keystone the rival fared worse. The conclusion is not that this programme's theory is weak. It is that the descriptive layer of Alzheimer's research is human-robust for every model, and the causal layer is model-dominated and trial-failed for all of them, the amyloid consensus included.

The framework therefore anchors at the two ends where human evidence is strongest and treats the middle as hypothesis. **Anchor One — ordered attrition:** Alzheimer's disease is a decades-long, anatomically ordered process of selective cellular attrition whose sequence is measurable in human post-mortem tissue. **Anchor Two — resilience:** pathology burden and cognitive outcome are dissociable, and the human genetics of protection supplies the field's only reverse-causation-free evidence of mechanism, because a fixed variant present from conception cannot have been installed by the disease it prevents. **The span between them — mechanism** — carries the programme's genuine intellectual contribution, the two specified inter-stage transitions, but carries it explicitly as a set of falsifiable hypotheses rather than as the structure's foundation.

Six procedural rules enforce non-circularity: grades are defined before claims; the corpus's internal coherence is inadmissible as evidence; every claim declares its provenance, and claims that originated inside the programme may not be counted as independent support; every load-bearing claim carries a falsifier fixed in advance; contradicting evidence appears in the same ledger as supporting evidence; and every demotion from the predecessor framework is recorded rather than quietly dropped. Applied honestly, these rules cost the framework several of its most quoted claims — the locus coeruleus as *the* origin, the perineuronal net as *the* terminus, three phases as a causal rather than a staging claim, and gamma entrainment as a load-bearing element, which a failed pre-registered human endpoint removes entirely.

I. Why This Paper Is Built Backwards

Ordinarily a framework paper states its thesis and then marshals evidence. This one states its rules of evidence first, then its grades, and only then its claims — because the failure it is written to correct occurred at exactly the point where those steps are usually collapsed.

The audit that preceded this paper established the following, from timestamps rather than from recollection.[24] A six-stage framework derived from one prize entrant's model existed two months before the project repository was initialized. Thirty entrants were then scored against it, using a rubric whose convergence nodes *were* that framework; the exercise returned zero contradictions, and the absence of contradiction was subsequently cited as a reason not to prune the corpus. The three-phase temporal spine arrived from a keystone whitepaper, and the synthesis report of 28 April 2026 states that its task was to realign the prize corpus to that spine — then concludes that the corpus's clean clustering into the spine is "the strongest available evidence that the architecture is correct."

That is a closed loop, and no amount of subsequent care can open it from the inside. The only remedy is a framework whose evidential warrant comes from outside the corpus entirely, and whose construction makes it structurally impossible to launder internal coherence into external support. Six rules do that work here.

Rule 1 — Grades before claims. The evidence scheme in Section II is defined without reference to any biological proposition. It is not tuned to the conclusions it will later be used to grade.

Rule 2 — Internal coherence is inadmissible. That a claim fits the architecture, that the corpus's documents agree, or that many analyses "converge" is not evidence and is not cited as evidence anywhere in this paper. Convergence among documents that share an author, a method, and a prior is a property of the author, not of the disease.

Rule 3 — Provenance is declared, and it is disqualifying. Every claim in the ledger is marked External, Assembled, Programme-invented, or Imported. Claims marked Programme-invented or Imported may be *stated*, and may be *tested*, but may never be counted as independent support for the framework that produced them.

Rule 4 — Falsifiers are fixed in advance. Every load-bearing claim carries, in the same row of the same table, the result that would refute it. Falsifiers written after a disconfirming result are worthless, so they are written now.

Rule 5 — Disconfirming evidence sits in the ledger, not in a discussion section. Contradicted claims appear in the same table, in the same format, as supported ones. There is no

separate chapter where inconvenient findings are contextualized.

Rule 6 — Demotions are recorded. Where this framework claims less than its predecessor did, Section VII names the claim, the predecessor, and the reason. Silent weakening is how a corpus launders retreat as refinement.

A note on what these rules do not do. They do not make the framework true, and they do not immunize it against being wrong — the ledger below contains items graded as contradicted, and one that has been removed outright. They make it *auditable*, which is the most any synthesis can honestly offer.

II. The Grade — Direction, Maturity, Provenance

Three orthogonal axes, defined here without reference to any claim they will later carry.

Direction asks only which way the evidence points. *Supports* — the weight of evidence favours the claim. *Mixed* — evidence is genuinely divided or the question is unresolved in the primary literature. *Contradicts* — the best available evidence tells against the claim. *Silent* — the claim has not been tested.

Maturity asks how far the evidence has travelled toward human confirmation, and is deliberately kept separate from direction. This separation is the paper's answer to a real objection raised against the June audit: that a recency-weighted standard penalizes a 2016–2024 literature for competing against a consensus assembled since 1992. The objection is legitimate, and collapsing direction into maturity is what made it possible. Graded on two axes, a recent but well-supported finding scores *Supports / low maturity* — which is an accurate description of its status, not a demerit.

Tier	Evidence class
M1	Human, population-scale or replicated across independent cohorts
M2	Human, single cohort — post-mortem series, biomarker study, or imaging cohort
M3	Human, case-level. Inferentially strong where a fixed germline variant precedes the outcome; statistically weak by sample size
M4	Animal model
M5	In vitro, or inference from adjacent findings only

Provenance asks where the claim came from, and is the axis that enforces Rule 3.

Mark	Origin	May it support the framework?
E	External — established independently of this programme	Yes
A	Assembled — external parts, combined here	Only as far as its parts allow
P	Programme-invented	No. May be tested, not counted
I	Imported from the programme's own prior framing	No

A claim graded *Supports / M1 / E* is load-bearing. A claim graded *Supports / M4 / P* is a research proposal wearing the costume of a finding, and the corpus has contained several.

III. What the Field's Evidence Profile Actually Looks Like

The framework's shape is not an aesthetic choice. It follows from a result that the programme generated and then set aside.

In June 2026 five literature reviewers were run with neutral, theory-free briefs. None was told about the methodology, the temporal architecture, the collapse trilogy, the perineuronal-net convergence node, or that any prior paper existed.[21] The audit's finding, in its own words:

> *"The quantitative spine is evidence-driven and survives blind replication. The mechanistic capstone — the PV/PNN convergence node, the microglial executioner, the ferroptosis bridge, and the 40 Hz reversibility — is model-driven, rests heavily on mouse data, and in one case has a failed human trial against it."*

And the pattern beneath it:

> *"Every load-bearing CONVERGE item is a* count*. Every DIVERGE item is a* mechanism *or a* reversal*."*

Two things about this result deserve more weight than they have been given. First, it was produced by the only instrument the programme ever built that was blind to the programme's own conclusions, which makes it the single most trustworthy internal finding in the corpus — and it was archived. Second, when the same protocol was run against a rival connectome-propagation account, that model **fared worse on its keystone.**[21] The asymmetry is therefore not a defect of this architecture. It is a property of the field: descriptive and ordinal claims about human tissue replicate; causal-mechanistic claims are model-dominated and trial-failed across every framework, the amyloid consensus included.

A framework that ignores this builds its foundations in the one stratum where nobody has load-bearing evidence. This one does the opposite: it anchors where human evidence is strong, at both ends of the disease, and treats the causal middle as the work still to be done.

IV. Anchor One — Ordered Attrition

The claim. Alzheimer's disease is a decades-long process of selective cellular attrition that proceeds in a reproducible anatomical order, and that order is measurable in human post-mortem tissue.

Its strength is that almost none of it belongs to this programme. The staging of Alzheimer-type tau pathology across the human lifespan — including the placement of the earliest lesions in brainstem nuclei in people in their twenties and thirties, decades before symptoms — is established neuropathology from large post-mortem series, and it stands whether or not any synthesis in this corpus is correct.[1] That independence is the point: under Rule 3, it is exactly the kind of claim that may carry weight.

Three components, each graded in Section VIII. That the process is **ordered** — that it appears in an anatomically stereotyped sequence rather than diffusely (M1, External). That it is **selective** — that specific cell populations fail while neighbours are spared, which the corpus's own census documents and single-nucleus atlases independently confirm (M2, Assembled).[25] And that it is **long** — that the interval between the first lesion and the first symptom is measured in decades, which human biomarker cohorts support (M1, External).

One component of this anchor requires an immediate correction to the corpus's prior framing, and it is entered here rather than buried. The programme's architecture placed the parvalbumin-positive interneuron and its perineuronal net at the terminus, as the cell on which every thread converges. The blind audit found instead that **parvalbumin involvement is late and functional, and that somatostatin interneurons precede it**[21] — a finding the second edition of the census subsequently reflects.[25] The ordering claim survives; the specific identity of the terminal cell does not, and is demoted accordingly.

What this anchor does **not** claim is equally important. It does not claim that the sequence is causal — that the failure of one population produces the failure of the next. Order is not causation, and the corpus has repeatedly written as though it were. An ordered sequence is compatible with a common upstream driver acting on populations of differing vulnerability, with independent parallel processes of differing tempo, and with a true causal relay. Nothing in Anchor One distinguishes these, and this framework does not pretend otherwise.

V. Anchor Two — Resilience

The claim. Pathology burden and cognitive outcome are dissociable; and the human genetics of protection supplies the strongest reverse-causation-free evidence of mechanism available in this disease.

The logic was stated correctly in the corpus's own atlas of protective genetics and is worth restating exactly, because it is the best epistemic argument the programme has made.[23] A risk allele shows that a pathway, when broken, permits the disease — consistent with that pathway being causal, and equally consistent with its being a bystander swept up in the wreckage. A protective allele shows something stronger. When a fixed germline change is present from conception and the disease that should have arrived does not, the change has run a controlled experiment across a human lifetime with the confound of reverse causation excluded *by the arrow of time*: the allele preceded the outcome, and nothing the disease did could have installed it. Protection is the closest thing human genetics offers to a proof of mechanism, and every protective variant is a natural trial of a target, already concluded.

The evidence divides sharply by maturity, and the framework insists on the division rather than averaging across it.

At **M1** sits the Icelandic *APP* A673T substitution, which lowers β-secretase cleavage of the amyloid precursor and cuts lifetime risk several-fold at population scale.[3] This is the single strongest item in the entire framework by both direction and maturity, and it carries an uncomfortable consequence the corpus has tended to soften: it is a clean human vindication of amyloid as a genuine contributor to the disease's early substrate. A framework that positions itself against amyloid primacy must still account for the fact that lowering amyloid production across a lifetime demonstrably protects.

At **M2** sits the dissociation itself: human post-mortem cohorts containing individuals with Alzheimer-range plaque and tangle burden who died cognitively intact, whose tissue differs from the demented.[2] Here a second correction to the corpus is entered directly. The programme has repeatedly described these brains as showing *intact perineuronal nets*. The primary study shows something more interesting and less convenient — a *distinct, homeostatic remodelling* of the net in resilient subjects, with aggrecan reduced in both resilient and diseased groups, and net-bearing neurons carrying conspicuously low tau.[2] "Intact" is not what was measured, and the framework does not repeat it.

At **M3** sit the two celebrated Antioquia cases: the *APOE3*-Christchurch homozygote who resisted an autosomal-dominant presenilin mutation for roughly three decades,[4] and the RELN-COLBOS heterozygous carrier who resisted the same mutation.[5] Their inferential

structure is exceptionally strong — fixed variant, lifetime experiment, arrow of time — and their sample size is one apiece, with a milder replication in twenty-seven Christchurch heterozygotes. Graded honestly, they are *Supports / M3*: high direction, low maturity. They are the framework's most cited items and among its least replicated, and both statements are true at once.

The therapeutic consequence of this anchor is developed in Section IX, but its shape is already visible: it does not require the causal cascade of the middle layer to be correct.

VI. The Span — Mechanism, Held as Hypothesis

Between the two anchors lies everything about *how* the ordered attrition happens, and this is where the programme did its most original work and where the evidence is weakest. Both facts are stated together, because the corpus has habitually stated the first and elided the second.

The genuinely invented material is the specification of the inter-stage transitions.[27] The programme's own architecture is right that these are what distinguish a theory from a chronology — a staging scheme that only names its stages describes the disease, while one that specifies the mechanism of each transition can be refuted. The first transition proposes that withdrawal of noradrenergic restraint from the locus coeruleus releases microglia from their homeostatic state, along the same efferents that carry seed-competent tau, through a vasculature whose barrier failure gates where they land.[11] [12] [13] The second proposes that senescent, lipid-laden microglia switch from cytokine to protease secretion, and that enzymatic and iron chemistry degrade the sulfated perisomatic matrix multiplicatively.[14]

Under Rule 3, none of this may count as support for the framework, because the programme invented it. Under Rule 4, all of it carries falsifiers. And under Rule 5, its weaknesses are entered in the same ledger as its strengths, of which the sharpest is this: **the complement-executioner component is contradicted in human tissue.** The mouse evidence that complement and microglia mediate early synapse loss is strong;[15] the sole direct human study implicates a complement-*independent* pathway.[21] The framework records that as *Contradicts*, not as "an area requiring further study."

The microglial state transition itself — the exit from a TGF-β/SMAD-maintained homeostatic programme into disease-associated, lipid-droplet-laden, and dystrophic trajectories — is better supported than the effector claims, but remains M4-dominant with human corroboration mainly morphological.[7] [8] [9] [10] It is the most promising item in the span and still not load-bearing.

One element is removed outright rather than demoted. Gamma entrainment at 40 Hz entered the corpus as a reversibility claim and was, in the blind audit's assessment, a therapeutic hope back-propagated into a structural claim; it has a **failed pre-registered human endpoint** behind it, with the widely quoted efficacy figure traceable to a subgroup and a press release.[21] A framework that retains it after that is not grading evidence, it is defending territory. It is gone.

VII. What This Framework Gives Up

Recorded explicitly, per Rule 6.

The locus coeruleus as the origin. The predecessor stated that the disease begins there, unqualified.[28] The programme's own graded edition had already retracted this three days before that statement was published, distinguishing four senses of "first" and conceding only those the neuropathology supports.[22] This framework adopts the retraction: the LC is an early and unusually vulnerable participant, and whether it is the necessary initiator is open.

The perineuronal net as the terminus. The predecessor stated that every thread terminates on one cell.[28] The graded edition demoted this to "a high-value final common bottleneck… not the single gate through which all dementia must pass,"[22] and the blind audit reached the same demotion independently by a different route.[21] Two independent routes to one conclusion is the closest thing to replication this corpus contains, and it is adopted.

Three phases as a causal claim. The audit established that the number three was fixed in a refactor named for pharmacology, whose products were three drug-class papers — three therapeutic windows, therefore three phases.[24] Staging remains useful; the framework retains ordered stages as description and abandons the claim that the count reflects a natural joint.

The unification claim. That the field's rival frameworks are "successive chapters of one structure" is an attractive proposition that the corpus tested only against itself. Under Rule 2 it is inadmissible until scored by a rubric that is not the architecture. The weaker survivor is defensible and worth keeping: many apparent disputes in this literature are disagreements about *timing* rather than about mechanism, and separating those two questions dissolves some of them.

Gamma reversibility. Removed, as above.

That ordering would strengthen as sampling improves. This framework's own first edition attached a prediction to Anchor One's framework-level falsifier: that better sampling would consolidate the anatomical order rather than fragment it. Sampling has improved, and it did the opposite. Data-driven subtype-and-stage modelling of tau-PET across 1,612 individuals resolves the population into four spatiotemporal trajectories, of which only the limbic-predominant third is described as following a Braak-like progression; the subtypes replicate across radiotracers, and event-based models trained and tested on different cohorts recover sequences differing by twenty-three adjacent-biomarker swaps. The prediction is withdrawn rather than reinterpreted, which is what Rule 4 requires of a falsifier that fails. Anchor One is narrowed accordingly: **ordered attrition is asserted at the population level, and order-invariance at the individual level is not asserted at all.** The anchor survives the narrowing — the post-mortem staging evidence is untouched, and the one element

that appears across every neuropathological subtype, and is hit hardest in the subtype that most nearly escapes the hippocampal phase, is the locus coeruleus.

That the phases can be tested in living humans. Not previously stated outright, and withdrawn here before it can be. Two of the three phases have no human in vivo instrument: the locus coeruleus sits below tau-PET resolution and inside the off-target binding profile of the dominant tracer, and no technique exists to image perineuronal nets in a living person. Phase I rests on post-mortem evidence, which is strong; Phase III rests on post-mortem and model evidence, which is partial. The framework's testable surface is narrower than its scope, and it says so rather than letting the silence read as confirmation. See phase-measurability.

VIII. The Ledger

Every load-bearing claim, graded on the scheme of Section II, with the result that would refute it. Grades read *Direction / Maturity*; provenance follows the marks of Section II.

Claim	Grade	Prov.	What would refute it
A1. AD-type pathology follows a reproducible anatomical order in human tissue	Supports / M1	E	Failure of the staging sequence to replicate in an unselected post-mortem cohort
A2. Earliest tau lesions appear in brainstem nuclei decades before symptoms	Supports / M1	E	Systematic re-examination placing the earliest lesion cortically
A3. Cell loss is selective and patterned, not diffuse	Supports / M2	A	An unbiased single-cell census showing uniform loss across populations
A4. The preclinical interval is measured in decades	Supports / M1	E	Biomarker trajectories compressing to a few years
A5. PV interneuron involvement is late and functional; SST precedes it	Supports / M2	E	Human tissue showing PV loss preceding SST
B1. Pathology burden and cognition are dissociable	Supports / M2	E	Dissociation vanishing under better-powered pathology quantification
B2. *APP* **A673T lowers lifetime risk several-fold**	Supports / M1	E	Failure to replicate outside the founder population
B3. Resilient brains show homeostatic PNN remodelling; net-bearing neurons carry low tau	Supports / M2	E	Independent cohort showing no PNN difference by cognitive status
B4. *APOE3*-**Christchurch confers resistance to autosomal-dominant AD**	Supports / M3	E	Larger carrier cohorts showing no delay in onset
B5. RELN-COLBOS confers resilience	Supports / M3	E	Additional carriers showing no protection
B6. *APOE4* **homozygosity behaves as a distinct genetic form**	Supports / M1	E	Failure to replicate the near-full-penetrance biomarker profile

Claim	Grade	Prov.	What would refute it
C1. Noradrenergic withdrawal releases microglial restraint	Supports / M4	A	Maintaining noradrenergic tone failing to slow the transition in humans
C2. Trans-synaptic tau seeding propagates along LC efferents	Mixed / M4	A	Blocking uptake failing to slow spread in human trial
C3. Blood–brain-barrier failure gates where the transition lands	Supports / M2	E	BBB integrity failing to predict decline independently of amyloid
C4. Microglia exit a TGF-β/SMAD-maintained homeostatic state	Supports / M4	E	Human single-nucleus data failing to show the state transition
C5. Complement-mediated synapse elimination is the executioner	**Contradicts** / M2	E	Already contradicted in human tissue; would require direct human evidence to restore
C6. NLRP3→IL-1β→protease switch strips the perisomatic matrix	Supports / M4	P	Protease inhibition failing to preserve matrix in human disease
C7. Oligodendrocyte ferroptosis liberates iron that amplifies degradation	Mixed / M5	P	Iron chelation failing to preserve matrix
C8. 40 Hz entrainment reverses pathology	**Contradicts** / M1	E	*Removed from the framework.* Failed pre-registered human endpoint
C9. Plaques originate predominantly from intraneuronal autophagic rupture	Mixed / M4	A	Quantification of the human plaque fraction attributable to each route

Two rows deserve comment. **C5 and C8 are graded against the framework**, and are printed here rather than omitted, because a ledger that contains only supporting evidence is a brochure. **C9** is entered at *Mixed* on the authority of the corpus's own reconciliation note, which states that the inside-out account "is not settled as a **quantity**" and that "the fraction of human plaque burden attributable to each route is not established. Any hard percentage overreaches the data."[26] Several documents in the corpus state it as settled fact; this framework does not.

Note also what the provenance column shows at a glance. Every claim in Anchors One and Two is marked **E** or **A** — external or assembled from external parts. Every claim marked **P** sits in the span, and none of them is load-bearing. That distribution is the framework's structural answer to the circularity finding: the weight rests entirely on material the programme did not generate.

IX. The Therapeutic Consequence — Phenocopy the Fortunate

A framework's therapeutic implication is the fastest test of whether its weight is correctly placed, and the difference here is sharp.

The predecessor's implication was phase-matched treatment: because the load-bearing lesion changes with the decade, a drug aimed at one phase's molecule can help only a patient in that phase.[28] It is an attractive argument and it explains trial failure elegantly. But it is entirely dependent on the causal cascade in the span being right — the layer this framework grades at M4 and below, and where one component is already contradicted in human tissue. A therapeutic programme resting there rests on the weakest stratum.

Anchor Two implies something different and better supported: **phenocopy the fortunate.**[23] Identify the human variants that demonstrably protect, and imitate what they do. This requires only that protective alleles do what human genetics says they do — the M1 and M3 evidence of Section V — and not that any proposed mechanism connecting the stages is correct. Lowering lifetime amyloid production protects, at population scale, on the A673T evidence.[3] Whatever the Christchurch and COLBOS variants are doing, they did it in humans, across decades, with reverse causation excluded.[4 5]

Two honest qualifications. The A673T result points at amyloid, which sits awkwardly with a corpus that has spent considerable effort repositioning it; the framework accepts the awkwardness rather than explaining it away. And the Antioquia cases are M3 — the correct response is to raise their maturity by finding more carriers, not to build a therapeutic programme on two people.

The timing insight survives this transition intact, because it does not depend on the causal middle. That the interval between first lesion and first symptom is measured in decades (A4, M1) is sufficient on its own to predict that interventions delivered after diagnosis will underperform, and to direct effort toward the preclinical window. That is the corpus's most valuable idea, and it is anchored.

X. What Would Refute This Framework

Per Rule 4, at the level of the framework rather than the claim.

Anchor One fails if the ordered attrition proves to be an artifact of selection in post-mortem series — if unbiased, prospectively collected cohorts show no stereotyped anatomical sequence at all. The anchor asserts ordered attrition **at the population level**, and it explicitly declines to assert order-invariance at the individual level.

This clause has been narrowed, and the narrowing is recorded in Section VII. The first edition added that "the framework predicts that ordering will strengthen, not weaken, as sampling improves." Sampling improved; ordering did not strengthen. That prediction is withdrawn rather than reinterpreted.

Anchor Two fails if the dissociation between pathology and cognition dissolves under better measurement — if apparently resilient brains turn out simply to carry less of the pathology that matters, measured properly. It also weakens substantially if the Antioquia effects fail to replicate as more carriers are identified.

The framework as a whole fails, and should be abandoned rather than patched, if the blind protocol is re-run against *this* document and returns the same asymmetry it returned in June — that is, if the anchors themselves prove to be model-driven rather than human-robust. That test is specified now, before the result is known, and running it is the first item of business.

A framework that could survive any result is not a framework. This one has already surrendered two claims to evidence collected before it was written, and it names in advance the results that would take the rest.

XI. Coda — What This Costs, and What It Buys

The framework stated here claims markedly less than its predecessor. It does not identify the origin of Alzheimer's disease, does not name its terminal structure, does not assert that the field's rival theories are chapters of one story, and does not offer a staging scheme with a principled number of stages. Set beside a document that did all four confidently, it will read as a retreat.

What it buys is that its central claims are ones the programme did not generate, cannot have smuggled in, and could not have confirmed by scoring its own corpus. Every load-bearing row in the ledger is marked External or Assembled. Every row the programme invented sits in the span and is explicitly barred from carrying weight. The two claims that evidence has already defeated are printed in the same table as the rest, at the same size.

That is a smaller structure, and a real one. The corpus's most valuable contributions survive it: the decades-long timescale and its consequence for trial design; the resilience-first therapeutic logic; the specification of transitions as testable hypotheses; and the census of which cells fail in what order. What does not survive is the claim to have unified the field — which was, on the record, the one claim the programme never tested against anything but itself.

A Note on Provenance and Method

This paper was composed by a large language model (Claude, Anthropic) and rests on two kinds of source, marked differently in the references. External primary works were verified against PubMed in a prior session and are cited with digital object identifiers; no new external citation has been introduced here, because the verification connector was unavailable at the time of writing and unverified citation is precisely the failure mode this corpus has repeatedly caught in itself. Internal sources are cited as repository artifacts, by path.

The framework's substantive debts are to work the programme did not do: the neuropathological staging of tau across the human lifespan, the human genetics of protection, and the post-mortem cohorts that establish the dissociation between pathology and cognition. Its original contributions — the specification of the two transitions — are, by its own Rule 3, barred from supporting it. That asymmetry is intentional, and it is the whole design.

One statement about this paper's own status. It supersedes the temporal architecture as the corpus's canonical framework, but it does not refute it. An imposed hypothesis can be true, and several of the demoted claims may yet be restored — by evidence, at the maturity the ledger specifies, and not by repetition.

References

References 1–19 are external primary works whose bibliographic metadata was verified against the PubMed database in a prior session; digital object identifiers are given. References 20–28 are artifacts within this repository, cited by path and by date established from build metadata or git history. No unverified external citation appears in this paper.

1. Braak H, Thal DR, Ghebremedhin E, Del Tredici K. Stages of the pathologic process in Alzheimer disease: age categories from 1 to 100 years. *J Neuropathol Exp Neurol.* 2011;70(11):960–969. DOI: 10.1097/NEN.0b013e318232a379

2. de Vries LE, Bahnerth A, Swaab DF, Verhaagen J, Carulli D. Resilience to Alzheimer's disease associates with alterations in perineuronal nets. *Alzheimers Dement.* 2024;21(2):e14504. DOI: 10.1002/alz.14504

3. Jonsson T, Atwal JK, Steinberg S, et al. A mutation in APP protects against Alzheimer's disease and age-related cognitive decline. *Nature.* 2012;488(7409):96–99. DOI: 10.1038/nature11283

4. Arboleda-Velasquez JF, Lopera F, O'Hare M, et al. Resistance to autosomal dominant Alzheimer's disease in an APOE3 Christchurch homozygote: a case report. *Nat Med.* 2019;25(11):1680–1683. DOI: 10.1038/s41591-019-0611-3

5. Lopera F, Marino C, Chandrahas AS, et al. Resilience to autosomal dominant Alzheimer's disease in a Reelin-COLBOS heterozygous man. *Nat Med.* 2023;29(5):1243–1252. DOI: 10.1038/s41591-023-02318-3

6. Fortea J, Pegueroles J, Alcolea D, et al. APOE4 homozygozity represents a distinct genetic form of Alzheimer's disease. *Nat Med.* 2024;30(5):1284–1291. DOI: 10.1038/s41591-024-02931-w

7. Butovsky O, Jedrychowski MP, Moore CS, et al. Identification of a unique TGF-β-dependent molecular and functional signature in microglia. *Nat Neurosci.* 2014;17(1):131–143. DOI: 10.1038/nn.3599

8. Keren-Shaul H, Spinrad A, Weiner A, et al. A unique microglia type associated with restricting development of Alzheimer's disease. *Cell.* 2017;169(7):1276–1290.e17. DOI: 10.1016/j.cell.2017.05.018

9. Marschallinger J, Iram T, Zardeneta M, et al. Lipid-droplet-accumulating microglia represent a dysfunctional and proinflammatory state in the aging brain. *Nat Neurosci.* 2020;23(2):194–208. DOI: 10.1038/s41593-019-0566-1

10. Streit WJ, Braak H, Xue QS, Bechmann I. Dystrophic (senescent) rather than activated microglial cells are associated with tau pathology and likely precede neurodegeneration in Alzheimer's disease. *Acta Neuropathol.* 2009;118(4):475–485. DOI: 10.1007/s00401-009-0556-6

11. Heneka MT, Nadrigny F, Regen T, et al. Locus ceruleus controls Alzheimer's disease pathology by modulating microglial functions through norepinephrine. *Proc Natl Acad Sci U S A.* 2010;107(13):6058–6063. DOI: 10.1073/pnas.0909586107

12. Nation DA, Sweeney MD, Montagne A, et al. Blood–brain barrier breakdown is an early biomarker of human cognitive dysfunction. *Nat Med.* 2019;25(2):270–276. DOI: 10.1038/s41591-018-0297-y

13. Montagne A, Nation DA, Sagare AP, et al. APOE4 leads to blood–brain barrier dysfunction predicting cognitive decline. *Nature.* 2020;581(7806):71–76. DOI: 10.1038/s41586-020-2247-3

14. Crapser JD, Spangenberg EE, Barahona RA, Arreola MA, Hohsfield LA, Green KN. Microglia facilitate loss of perineuronal nets in the Alzheimer's disease brain. *EBioMedicine*. 2020;58:102919. DOI: 10.1016/j.ebiom.2020.102919

15. Hong S, Beja-Glasser VF, Nfonoyim BM, et al. Complement and microglia mediate early synapse loss in Alzheimer mouse models. *Science*. 2016;352(6286):712–716. DOI: 10.1126/science.aad8373

16. Frost B, Hemberg M, Lewis J, Feany MB. Tau promotes neurodegeneration through global chromatin relaxation. *Nat Neurosci*. 2014;17(3):357–366. DOI: 10.1038/nn.3639

17. Lee JH, Yang DS, Goulbourne CN, et al. Faulty autolysosome acidification in Alzheimer's disease mouse models induces autophagic build-up of Aβ in neurons, yielding senile plaques. *Nat Neurosci*. 2022;25(6):688–701. DOI: 10.1038/s41593-022-01084-8

18. Gouras GK, Tsai J, Näslund J, et al. Intraneuronal Aβ42 accumulation in human brain. *Am J Pathol*. 2000;156(1):15–20. DOI: 10.1016/S0002-9440(10)64700-1

19. Alzheimer's Association Calcium Hypothesis Workgroup. Calcium Hypothesis of Alzheimer's disease and brain aging: a framework for integrating new evidence into a comprehensive theory of pathogenesis. *Alzheimers Dement*. 2017;13(2):178–182.e17. DOI: 10.1016/j.jalz.2016.12.006

20. research/methodology-history/The_Map_That_Became_a_Clock.md — superseded methodological account, July 2026.

21. _archive/bias-audit-diversion-2026-06/order-of-dying/BIAS_AUDIT_MEMO_2026-06-25.md — the model-blind audit: five theory-free reviewers; the CONVERGE/DIVERGE pattern; the failed 40 Hz pre-registered endpoint; the rival-model comparison. Archived and rejected; see the accompanying README.md.

22. research/unified-architecture/PhD_Thesis_Unified_Architecture.md, Second Edition, 2026-07-07 — the four senses of "first"; the net as "a high-value final common bottleneck."

23. research/protective-genome/PhD_Thesis_Architecture_of_Resistance.md, 2026-07-15 — the logic of protection as reverse-causation-free evidence; "phenocopy the fortunate."

24. research/methodology-history/The_Order_of_Discovery.md, 2026-07-20 — the dated reconstruction establishing the imported spine, the pharmacological phase count, and the circular confirmation.

25. research/cellular-architecture/ONS_CellularArchitecture_Thesis_SecondEdition.md, 2026-06-25 — the cellular census; SST-early and the revised soma threshold.

26. kb/wiki/concepts/plaque-origin-what-is-established.md, 2026-07-17 — "not settled as a quantity"; "any hard percentage overreaches the data."

27. research/collapse-trilogy/homeostatic-microglial/LocusCoeruleus_Transition_Paper.md, 2026-05-21, and research/collapse-trilogy/convergent-synaptic/Proteolytic_Turn_Paper.md, 2026-05-22 — the two specified transitions.

28. research/ofp-entry/ONS_OFP_Entry.md, 2026-07-10 — the predecessor framework as submitted.