The Load-Bearing Neuron

What a cortical cell spends across ninety years, and where the bill is settled — the architecture of Alzheimer's synaptic failure traced to the loss of governed control at cofilin serine 3

Five loads · One residue · The writer and the reader · The shield · The program re-opened

Benjamin Aaron Gustafsson

AdultCognitiveDisease.com

August 2026

A setting, not a signal. A signal arrives, acts, and leaves, and the value it leaves behind is a message; a setting arrives and stays, and the value it holds is only a level. The residue is the same residue either way. What the disease takes is not the phosphate but the difference between the two.

Contents

Abstract

Synapse loss is the strongest structural correlate of cognitive impairment in Alzheimer's disease. In the study that established this, neocortical synapse density carried a multivariate correlation of 0.96 against the Dementia Rating Scale, while plaque density contributed only 26 per cent of that model's strength (Terry et al., 1991). Any complete account of the disease must therefore say, in mechanical terms, what happens at a dendritic spine — and a dendritic spine is an actin structure whose volume is the running balance between filament severing and filament stabilisation. That balance is set by one protein, cofilin-1, at one residue: **serine 3**.

This paper makes three claims and grades each.

First, that five independent laboratories have converged on that residue in this disease and disagree about the direction of the change in human tissue, and that the disagreement has never been stated in print. Reelin, through ApoER2 and Disabled-1, *phosphorylates* serine 3 and stabilises the cytoskeleton (Chai et al., 2009). Amyloid-β, through the immune receptor LilrB2 and its murine orthologue PirB, produces enhanced cofilin signalling — *de*phosphorylation — reported in mouse and detected in human Alzheimer brain (Kim et al., 2013), a direction independently reproduced with a natural receptor antagonist (Kawaguchi et al., 2022) and, by a third route, through β1-integrin and the phosphatase Slingshot-1, with activated and oxidised cofilin found in mitochondria from Alzheimer brain (Woo et al., 2015a). And amyloid-β, through Rho-associated kinase, *increases* serine-3 phosphorylation in the post-synaptic fraction of human Alzheimer cortex, where that phosphorylation is reported necessary and sufficient for synaptic impairment (Rush et al., 2018). Three groups describe amyloid removing the phosphate; one describes amyloid adding it; all four touch human tissue. The 2018 paper states the field's position on the actin cytoskeleton to be "unknown and contentious" and does not resolve it. No subsequent paper has.

Second, that the resolution is not a direction but a loss of range, and that the loss of range has four material preconditions which a small class of neurons cannot meet. Cofilin's action on actin is non-monotonic: it severs filaments at low cofilin-to-actin ratios and stabilises them at high ratios (Bamburg and Bernstein, 2016). An effector that inverts across its own range has no pathological direction; it has a pathological loss of governance. But governance is not free. To drive a residue transiently and locally and then release it, a neuron must have (i) a *source* of the placed signal, (ii) *energy* to pay for the filament turnover that following the signal requires, (iii) *receptors delivered back to the surface* on which the signal lands, and (iv) an *exposure regime* in which chronic ligand does not occupy those receptors continuously. This paper argues that these four preconditions are precisely what a **load-bearing neuron** — a cell high in geometric, electrical, metabolic, proteostatic and exposure load, and never replaced — progressively loses across a human lifetime, and that their loss is what converts a signal into a setting.

Third, that every molecule which executes synapse loss in the ageing cortex is a molecule that built the circuit in youth, and that this is the deepest regularity in the material. Reelin phosphorylates serine 3 to arrest a migrating neuron before birth and phosphorylates the same residue to stabilise a spine in the adult. Complement component 4 prunes synapses in development (Sekar et al., 2016) and returns as C4d to strip spines in the aged and Alzheimer cortex (Brott et al., 2025). PirB closes the critical period (Syken et al., 2006) and executes plasticity failure in the Alzheimer model (Kim et al., 2013). The perineuronal net closes the critical period (Pizzorusso et al., 2002) and degrades in disease. The disease is not a novel process. **It is the closing program of cortical development, re-opened in a cortex that no longer has the signals, the energy or the shielding to govern it.**

The paper traces the lifetimes of the cells this argument names — the layer III and layer V pyramidal neuron that carries the residue; the somatostatin-expressing interneuron that writes the reelin signal it cannot itself read; the parvalbumin interneuron whose net regulates the exposure of others; the net-bearing minority of excitatory neurons that appear to resist; and the coerulean neuron, the extreme case that establishes the load rule and, on the argument advanced here, is *not* required to be the origin of anything. The keystone connection is stated as an inference and given a falsifier: **the cell that supplies the only characterised placed, transient source of serine-3 phosphorylation is a cell that dies early in this disease, and its death leaves the residue driven by tone alone.**

A hard, cheap, falsifying prediction follows and nobody has tested it: **in Alzheimer cortex the synapse-to-synapse variance of phospho-serine-3 cofilin will be elevated relative to age-matched control even where the mean is unchanged, will separate into modes on local receptor content, and will be lower on the minority of pyramidal neurons that carry a perineuronal net.** The method exists and has been applied to these exact synapses.

Four consequences follow for practice. Bulk phospho-cofilin is disqualified as an endpoint. Monotonic Rho-kinase therapy is predicted to help one stratum of patients and harm another. The therapeutic target is not the kinase and not the phosphatase but the restoration of signal-driven control. And because reelin's receptor activation obligately requires N-sulfated heparan sulfate (Pan et al., 2025) — the same polymer through which pathological tau enters neurons (Holmes et al., 2013) — a heparin mimetic deployed against tau propagation would silence the reelin brake by the identical chemistry. That conflict is verified at drug level and has not been stated.

Note on evidence and grading

Every load-bearing claim in this paper carries a maturity grade, assigned on a five-point scale, and the grades are collected in a single ledger in Part XVII rather than distributed through a discussion section.

M1 — human, population-scale, or replicated across independent cohorts. **M2** — human, single cohort or tissue series. **M3** — animal, replicated across laboratories. **M4** — animal or primary culture, single laboratory. **M5** — in vitro biochemistry, or an inference drawn from adjacent established findings.

Two further labels are used where neither applies. **Observation** marks a claim about what the published literature does or does not contain — for example, that two groups report opposite signs and do not cite each other. Such a claim is not an experimental result; it is verifiable directly from the sources, and the reader can check it in an afternoon. **Prediction** marks a statement about a measurement that has not been made.

The paper's own original claims are graded **M5** or **Prediction** without exception. There is no result here. What is offered is a reading of results other people obtained, a mechanism that would explain why those results conflict, and a set of measurements that would show the reading to be wrong. Readers who want to attack the argument should begin with the four entries in the ledger marked as this paper's own.

References were checked against the source record — title, authorship, year, journal, and, critically, the *direction* of the reported effect. The directional disagreement documented in Part IV is the reason that check was necessary, and it is the reason no claim in this paper about the sign of a phosphorylation is made from memory.

How this paper is built, and why it is not a chain

Accounts of Alzheimer's disease are commonly cast as chains: this causes that, which causes the next thing, and the last link is the dementia. A chain has the virtues of narrative — it is easy to follow, easy to summarise, and easy to remember — and it has one fatal vice, which is that it purchases its coherence with a claim about *time* that the evidence rarely supports. To say that A precedes B in a chain is to assert a temporal ordering; to assert a temporal ordering in a disease whose preclinical phase runs for decades and whose only human evidence is cross-sectional autopsy material is to assert something that cannot, in most cases, be checked. The result is a class of account that reads as explanation and functions as narrative.

This paper is built the other way. It makes no claim about what starts the disease, and it makes no claim that the events it describes occur in a particular order. What it claims is that a set of independently established mechanisms **converge on one physical location**, and it names that

location precisely: serine 3 of cofilin-1, on the dendritic spines of excitatory pyramidal neurons of neocortex and hippocampus. The connections between those mechanisms and the residue are called **bridges** rather than links, and the distinction is not cosmetic. A link asserts that one thing happens and then another does. A bridge asserts only that two bodies of work, developed independently and usually in ignorance of one another, terminate at the same molecule — and that if you stand at that molecule you can see both of them.

Bridges can be graded; temporal orderings, in this disease, mostly cannot. That is the whole of the methodological argument, and it governs the paper's shape. Part X sets out seven bridges, grades each, and states what would demolish it. One of those bridges — the noradrenergic one — is the kind of claim that in a chain-form account would be promoted to the origin of everything. Here it is one bridge among seven, it is the most inferential of the seven, and Part IX states explicitly that the argument does not need it.

The order of the parts is therefore logical, not chronological. Parts I through V build the argument at the residue. Parts VI through IX trace the lifetimes of the cells on which the residue sits, because the paper's central positive proposal is about what those cells can no longer afford. Part X assembles the bridges. Part XI states the synthesis. Parts XII through XVII state the predictions, the experiments, the prohibitions, the scope, the refutation conditions and the ledger.

PART I — The Claim, and the Kind of Claim It Is

1.1 What is asserted

The assertion of this paper can be written in one sentence, and everything that follows is either a justification of a term in it or a statement of what would falsify it.

> *Alzheimer's synaptic failure is executed at cofilin serine 3 on the dendritic spines of excitatory glutamatergic pyramidal neurons of neocortex and hippocampus; the lesion at that residue is not a displacement of its mean in either direction but a loss of the capacity to drive it transiently and locally and then release it; and that capacity is lost because the four material preconditions of governed control — a source of placed signal, energy for filament turnover, receptor return to the surface, and a shielded exposure regime — are exactly the four things that a load-bearing cortical neuron, across a human lifetime, progressively cannot supply.*

Four terms in that sentence do work and must be defined before they can be defended.

Executed. The claim is about the mechanical step by which an upstream cause becomes a lost spine. It is not a claim about which upstream cause. A theory of execution is compatible with amyloid-first, tau-first, vascular-first, inflammatory-first and metabolic-first accounts of origin, and that compatibility is a weakness as much as a strength: an account consistent with every origin story constrains none of them. Part XV states the boundary in full.

Loss of range, not displacement of mean. This is the paper's central positive proposal, and it is an inference, not a result. It follows from a single established biochemical fact — that cofilin's effect on actin inverts across cofilin's own concentration range (Bamburg and Bernstein, 2016) — combined with a single observation about the literature, that laboratories measuring the residue in human tissue report opposite signs. Part V develops it; Part XII states the measurement that would refute it.

Load-bearing. Used here in a specific and defensible sense, not as a metaphor for "important." Part II gives the definition, the five component loads, and the census of which cortical cells satisfy it. The term earns its place only if it picks out the same cells that the disease picks out, and if it does so on grounds that were not derived from the disease. Part II is written to that standard.

Preconditions of governed control. The most consequential and the least conventional part of the claim. The literature at this residue has treated control as a signalling question: which

kinase, which phosphatase, which ligand. This paper argues that control is also a *budgetary* question. A neuron that cannot pay for actin turnover cannot make the excursion; a neuron whose receptors are not returned to the membrane has nothing for the signal to land on; a neuron whose surface is continuously occupied by a ligand that encodes nothing has no baseline to return to; and a neuron whose signal source is dead has nothing to follow. Parts VI through IX trace how each precondition fails, cell by cell, across a lifetime.

1.2 Convergence claims and origin claims

Claims about this disease fail more often for being of the wrong *kind* than for being unsupported. A claim about what initiates Alzheimer's disease requires evidence about the earliest affected people, prospectively identified. A claim about what the disease converges on requires evidence that multiple, independently established upstream processes arrive at one place. These are different claims requiring different evidence, and they are conflated with a regularity that has damaged the field: a convergence result is presented as an origin story, and a therapeutic programme is then built on the origin story rather than on the convergence.

The distinction matters here because the argument of this paper is unusually easy to inflate. Once one has shown that reelin, amyloid, complement, integrin signalling and Rho-kinase all write to one residue, the temptation to say *therefore the disease begins at that residue* is considerable, and it would be wrong. Convergence at a residue is compatible with the residue being entirely passive — a place where things happen to meet, with no causal priority of its own. What raises it above passivity is not the convergence but the second claim: that the residue is *non-monotonic*, so that both directions of departure from the healthy state are pathological, and therefore that no upstream cause can push it anywhere safe.

That is a substantive property of the node, not an artefact of how many arrows point at it. It is also the reason the therapeutic reading in Part XIV is counterintuitive.

1.3 What was wrong with the chain

It is worth being explicit about the alternative this paper declines, because the alternative is attractive and is frequently right in other diseases.

A chain-form account of Alzheimer's disease selects an early, distal, plausible perturbation — a shift in gut ecology, a rise in peripheral inflammation, a failure of sleep, a change in vascular compliance — and traces it forward through a sequence of intermediate steps to the cortex. Each step is individually supportable. The account's persuasiveness comes from its continuity, and continuity is exactly what cannot be tested in a disease that unfolds over five decades in a species that cannot be biopsied.

Three specific failures follow, and all three are avoidable.

The temporal claim absorbs the evidential weight. In a chain, the claim that A precedes B does more argumentative work than any of the individual mechanisms, and it is the least supported claim in the account. Human evidence for ordering in this disease comes almost entirely from cross-sectional autopsy series stratified by stage, which establish co-occurrence and correlation with severity, not sequence within an individual.

Distal links dominate proximal ones. Because the chain's rhetorical energy comes from the surprise of its starting point, the distal end receives disproportionate attention. The result is accounts in which a great deal is said about the first link and comparatively little about the mechanism at the tissue where the symptom is generated.

Interventions are placed at the wrong end. If the account is a chain, the natural therapeutic move is to intervene at the earliest link, which is usually the one furthest from the failing synapse and the one with the largest number of intervening steps at which the intervention can be lost.

None of this makes the distal mechanisms false. The noradrenergic and matrix arguments developed in Parts IX and X are retained here in full, and one of them supplies the most specific pharmacological prediction in the paper. What changes is their *status*: they are bridges to a residue, gradeable individually, and the argument survives the failure of any one of them. That is what a convergence architecture buys and a chain does not.

1.4 The three propositions, stated so they can be attacked

Proposition 1 — the disagreement. Five laboratories have written to cofilin serine 3 in Alzheimer's disease. Their reported directions are inconsistent. Three of the five report human tissue. None addresses the others on the point of sign. *Status: Observation. Verifiable from the sources tabulated in Part IV. If a reader finds a published paper that states and resolves this conflict, Proposition 1 fails and the paper's motivation weakens considerably.*

Proposition 2 — the reconciliation. The residue's effector is non-monotonic; therefore both directions of change are pathological; therefore a bulk mean across synapses with differing receptor complements is uninterpretable; therefore the lesion is best read as a loss of governed range rather than a displacement of the mean. *Status: Inference, M5. Untested. Part XII states its falsifier, which is cheap.*

Proposition 3 — the load argument. Governed range has four material preconditions. The neurons that fail earliest and hardest in this disease are the neurons that carry the most load by an independent definition, and load is precisely what erodes those four preconditions across a lifetime. *Status: Inference, M5, assembled from graded components. Parts VI through IX supply the components with their individual grades; Part XVI states what would refute the assembly.*

Proposition 1 can be true and 2 false. Propositions 1 and 2 can both be true and 3 false — in which case the residue account survives and the lifetime account is decoration. Proposition 3 is the paper's most ambitious claim and the one most likely to be wrong, and it is stated third for that reason.

1.5 A note on the word "governance"

The paper uses "governance," "governed control" and "range" repeatedly, and a reader is entitled to ask whether these are doing analytic work or supplying atmosphere. The intended meaning is narrow and can be stated in four properties. A residue is **governed** when its phosphorylation state is:

Driven, not merely determined. Something writes to it in response to an event, rather than its value being the passive resultant of two constitutive activities.

Placed. The write is local — to one synapse, or a small number — rather than uniform across the neuron. A neuron with ten thousand spines that can only move them together has one degree of freedom, not ten thousand.

Transient. The write has a termination. It returns toward a baseline, and the return is itself an event with a mechanism rather than a passive decay.

Reversible in both directions. The residue can be moved up and down on demand, which is what distinguishes a control variable from a threshold.

A residue that has lost governance in this sense may still show a perfectly normal average phosphorylation. That is the whole operational point, and it is why every distinctive prediction in Part XII is about a distribution, an excursion or an interaction rather than about a level. **Ungoverned does not mean deranged; it means uninformative.** The residue continues to have a value; that value has stopped carrying a message.

The word is a metaphor from control engineering and the metaphor should not be pushed further than the four properties above. In particular, no claim is made that the cortex implements a controller with a set-point in any formal sense, and no model of that controller is offered. What is claimed is that these four properties are individually measurable — three of them by the experiments in Part XIII — and that a system that has lost them behaves in a way a system that merely has a shifted mean does not.

1.6 What this paper does not do

It does not explain who gets Alzheimer's disease, or why one person with a heavy amyloid burden becomes demented and another dies cognitively intact. Those are the important questions and this is not an answer to them.

It does not explain selective vulnerability at the level of *why*. It offers a framework — load — within which vulnerability can be described and, in principle, ranked, and Part II is explicit that the framework is descriptive. Saying that the earliest-failing cells are the most heavily loaded is not the same as demonstrating that load causes the failure, and the paper does not claim to have demonstrated it.

It does not speak to the fraction of dementia that is not attributable to Alzheimer pathology, nor to the substantial fraction of people meeting neuropathological criteria who are not demented (Boyle et al., 2019; Crary et al., 2014; Nelson et al., 2019). It is a theory of a mechanism, not a theory of a population, and Part XV states the arithmetic of that restriction rather than leaving it to a limitations paragraph.

And it does not offer a result. It offers a reading, a mechanism, and a set of measurements. The most useful thing a reader can do with it is perform one of them.

PART II — Load

2.1 What the term has to mean if it is to do work

"Selectively vulnerable" is the phrase the field uses, and it is a description masquerading as an explanation. It says that certain neurons die first. It does not say what property of those neurons makes them die first, and because it does not, it cannot be used to predict which neuron will fail in a circuit nobody has yet examined, nor to rank interventions by which cell they protect.

The term proposed here — **load-bearing** — is offered as a replacement that can do those things, and it is subject to a discipline that "vulnerable" is not. It must be definable without reference to Alzheimer's disease. If the definition contains any clause about tangles, plaques, or degeneration, the argument is circular: the cells will have been selected because they fail, and the discovery that they fail will be presented as a finding.

So the definition is built entirely from properties a neuron has in a *healthy* brain, measurable in a healthy brain, and known independently of any disease. Only after the definition is fixed is the resulting list of cells compared against the list the disease produces. The comparison is the test. If the two lists diverge, the framework is wrong.

A **load-bearing neuron** is a cell that scores high on at least three of five independent burdens, and that is never replaced.

2.2 The five loads

Load 1 — geometric. The physical extent of the structure the cell must build and hold: total membrane area, total axonal and dendritic length, total number of synaptic contacts. A neuron does not merely possess its arbor; it maintains it continuously against turnover, and the maintenance cost scales with the extent. Layer III and layer V pyramidal neurons of association cortex carry arbors bearing thousands to tens of thousands of spines. The noradrenergic neuron of the locus coeruleus carries a projection of extraordinary extent for a cell of its size, largely unmyelinated. Long, thin, unmyelinated processes are the expensive case in both directions: they cost more per unit length to maintain and they are less protected.

Load 2 — electrical. The rate and pattern of activity, and with it the calcium that activity admits. A tonically firing pacemaker cell admits calcium continuously and must extrude it continuously. A fast-spiking interneuron sustains rates an order of magnitude above a pyramidal neuron. The relevance of calcium to the argument of this paper is direct and unusual: calcium is not

merely a stressor here, it is a *writer to the residue*. Calcium-dependent activation of calcineurin dephosphorylates and thereby activates Slingshot-1, which dephosphorylates cofilin at serine 3 (Wang et al., 2005). Electrical load is therefore, in the most literal sense, load on the switch this paper is about.

Load 3 — metabolic. ATP consumption per unit time, and the fraction of it that is obligatory. The brain's energy budget divides into signalling and non-signalling components, and just under half is spent on processes not directly attributable to information transmission (Engl and Attwell, 2015; Harris et al., 2012). Within that non-signalling half sits the item that matters here, and it is worth quoting the state of knowledge exactly rather than rounding it: **the energetic cost of actin treadmilling is contested, with published estimates ranging from under 1 per cent of the brain's global energy budget to roughly half of neuronal energy use** (Engl and Attwell, 2015). The experimental work at the high end of that range comes from the cofilin literature itself, where ATP hydrolysis associated with actin filament turnover was taken to be responsible for approximately 50 per cent of neuronal energy consumption (Bernstein et al., 2006).

That the estimate spans two orders of magnitude is not a reason to set it aside. It is a reason to be careful about what is claimed. What the argument here requires is not the high figure but a much weaker statement: that actin filament turnover is a *material* and *obligatory* energy cost, that it scales with the amount of dynamic actin a cell maintains, and that a cell with a large spine complement therefore pays more for structural maintenance than a cell with a small one. That statement is safe at any point in the published range.

Load 4 — proteostatic. The burden of making, delivering, retrieving and degrading protein at distance from the soma. A spine is a functional compartment tens or hundreds of micrometres from the nucleus; the receptors on its surface must be delivered, endocytosed, sorted and returned. Endosomal recycling is the machinery that performs the return, and its failure is among the best-established molecular lesions in this disease: retromer components are reduced in the Alzheimer brain, and the affected region relies on a distinct retromer core dedicated to endosomal recycling (Small et al., 2005; Simoes et al., 2021). Degradation has a spine-local arm too: a nervous-system-specific proteasome complex resides in the neuronal plasma membrane and degrades nascent protein in an activity-dependent manner (Ramachandran and Margolis, 2017; Ramachandran et al., 2018).

Load 5 — exposure. How much of the cell's surface is accessible to the extracellular ligand pool, and whether that access is regulated. This is the least conventional of the five and, for the argument here, the most consequential, because the ligands that write to serine 3 are extracellular: reelin, amyloid-β oligomers, the complement fragment C4d, myelin-associated inhibitors. A neuron sheathed in a condensed perineuronal net has restricted lateral mobility of surface components and limited access to its membrane. A neuron without one is, in the relevant sense, undefended — its

receptors are exposed to whatever the interstitial fluid contains, at whatever concentration it contains it, for as long as it contains it.

Exposure is where the paper's central mechanism gets its bite. A signal is a ligand that arrives, acts, and leaves. A setting is a ligand that arrives and stays. The difference between them is not chemical; it is a matter of the exposure regime.

2.3 The sixth condition: non-replacement

The five loads are quantities. The sixth condition is categorical: the cell is not replaced.

This is what converts load into accumulated load. A hepatocyte carrying a heavy synthetic burden is replaced; a cortical pyramidal neuron born before the animal was is the same cell at ninety. Every failure of proteostasis it has ever suffered, every oxidised protein it did not clear, every episode in which it could not pay for turnover, is carried forward in the same cell. Non-replacement is why lifetime is the right unit of analysis for this argument and why Parts VI through IX are organised as biographies rather than as mechanisms.

It also means that the relevant quantity for each of the five loads is not the instantaneous rate but the **integral of the rate over the lifetime**. Two cells with the same present-day firing rate are not equally loaded if one has been firing at that rate for eight decades and the other has not. The disease is a disease of old people; the framework must be one in which age is a term, not a footnote.

2.4 The census

Applying the definition to cortical and brainstem cell classes, without reference to disease, produces the following. The load columns are ordinal judgements from healthy-brain anatomy and physiology, not measurements; they are stated so a reader can dispute them individually.

cell class	geometric	electrical	metabolic	proteostatic	exposure	load-bearing?
Layer III / V pyramidal neuron, association cortex	very high	moderate	high	very high	high (mostly netless)	**yes**
Entorhinal layer II projection neuron	high	moderate	high	high	high	**yes**
Somatostatin / NPY interneuron (Martinotti, bitufted)	high	high	high	moderate	high (netless)	**yes**
Parvalbumin fast-spiking interneuron	moderate	very high	very high	moderate	**low** (net-bearing)	partly — shielded
Locus coeruleus noradrenergic neuron	extreme	high (tonic)	extreme	high	extreme (netless)	**yes, maximally**
Net-bearing pyramidal minority	high	moderate	high	high	**low** (net-bearing)	partly — shielded
Cortical astrocyte	high	n/a	moderate	moderate	high	not applicable (replaceable pool, different biology)

Two things about this table are worth stating plainly, because both are places where the framework could have failed and did not.

The list was not assembled from the disease and it matches the disease. The cells that score highest — deep-layer and entorhinal projection neurons, somatostatin interneurons, the locus coeruleus — are the cells the neuropathology has independently identified as failing earliest and most severely. Entorhinal layer II is the classical origin of cortical tangle pathology and the site of the earliest cortical neuron loss; single-nucleus work has since localised the vulnerability within that layer to a specific excitatory subpopulation (Leng et al., 2021). The locus coeruleus carries the earliest detectable tau in the brain, present in a majority of people by the fourth decade (Braak and Del Tredici, 2011; Braak et al., 2011; Ehrenberg et al., 2017). Somatostatin is the oldest non-cholinergic neurochemical deficit reported in this disease (Davies et al., 1980).

The one cell that scores high on the first four loads and low on the fifth behaves differently, and the difference is informative. The parvalbumin interneuron is metabolically

and electrically the most extreme cell in the cortex and is nevertheless not among the first to accumulate tangle pathology. What distinguishes it is the net. In subcortical regions, neurons associated with aggrecan-based perineuronal nets are protected against tau pathology (Morawski et al., 2010); in human frontal cortex, excitatory neurons bearing a net carry conspicuously low phospho-tau (de Vries et al., 2024). Exposure, in other words, is not a decorative fifth column. It appears able to override the other four.

That single observation carries a large share of the paper's therapeutic argument, and Part VIII is devoted to it.

2.5 What a load-bearing cell cannot do

The framework becomes predictive at the point where one asks what a heavily loaded, unreplaced cell is *unable* to do that a lightly loaded one can.

It cannot hold a large reserve. Reserve is unspent capacity, and a cell operating near its ceiling has none. A cell whose obligatory maintenance consumes most of its ATP production cannot fund a large transient excursion in filament turnover on demand.

It cannot tolerate a slow leak. A modest, chronic inefficiency — a small fraction of receptors not returned to the surface per cycle, a small fraction of oxidised protein not cleared per day — is trivial over a year and decisive over sixty. Non-replacement converts rates into totals.

It cannot distinguish a signal from a setting once its surface is chronically occupied. This is the sharpest of the three and it is the one that connects load to the residue. Signal-driven control requires a baseline to return to. A cell whose receptors are continuously engaged by a ligand that carries no information has lost the baseline, and with it the meaning of the excursion.

2.6 The claim that makes load matter

Everything above is preparatory. The claim that gives it force is this: **the five loads are not settled in five different places. They are settled in one.**

- **Geometric load** is spine number and spine volume, which is dynamic actin mass, which is cofilin's substrate.
- **Electrical load** admits calcium, which activates calcineurin, which activates Slingshot-1, which dephosphorylates cofilin at serine 3 (Wang et al., 2005).
- **Metabolic load** is dominated, among non-signalling costs, by actin filament turnover — the process cofilin governs — and the neuron's emergency response to energy failure is itself a cofilin event: rod formation slows filament turnover and, with it, the associated ATP

hydrolysis, transiently retarding the decline of mitochondrial potential (Bernstein et al., 2006).

- **Proteostatic load** determines whether the receptors that write to serine 3 — ApoER2, LilrB2, β1-integrin — are present on the membrane to be written through.
- **Exposure load** determines whether those receptors are engaged transiently by a signal or continuously by a setting.

Five burdens, one settlement point. This is the structural reason the paper is about a residue rather than about a pathway, and it is the reason Part III describes the biochemistry of that residue in more detail than a paper of this scope would ordinarily require. If the settlement point were not non-monotonic, the convergence would be an interesting coincidence. Because it is non-monotonic, the convergence has a consequence: **there is no direction in which a heavily loaded neuron can be pushed at this residue that is safe**, and therefore no monotonic intervention at it that is safe either.

2.7 Is the framework circular?

The obvious objection to any vulnerability framework is that it was reverse-engineered from the answer. It is worth meeting the objection directly rather than trusting the reader to notice that Section 2.1 anticipated it.

The strong form of the objection. Nobody selects five arbitrary properties and finds that they rank cortical cells in the order the disease does. The five were chosen, consciously or not, because they distinguish the cells already known to fail. The framework therefore predicts nothing; it summarises.

What can be said in reply. Three things, of decreasing strength.

First, the framework produces a ranking that is testable on populations it was not built from. The cells in the census of Section 2.4 are the well-studied ones. The framework applies without modification to cell classes whose fate in this disease is not well characterised — layer 6b neurons, claustral projection neurons, the several molecularly distinct subclasses within the somatostatin family — and it makes ordinal predictions about each. A framework that can be run on new populations is not merely a summary of old ones, and prediction P11 in Part XII states the general form of the test.

Second, the framework has one internal near-miss that a reverse-engineered account would not have produced. The parvalbumin interneuron scores at the top on four of the five loads and is not among the first cells to fail. A framework designed backwards from the pathology would have found a way to score it low on all five. This one scores it high on four, low on one, and is forced to claim that the fifth column can override the other four — which is an exposed, specific, refutable claim (falsifier 11 in Part XVI) rather than a comfortable one.

Third, the loads are conventional quantities, individually measured by other people for other reasons. Arbor extent, firing rate, ATP consumption, recycling flux and matrix ensheathment are all standard measurements in cellular neuroscience with literatures that predate the framework and were not assembled with it in mind.

What cannot be said in reply. That the framework has been validated. It has not. The correspondence in Section 2.4 is a correspondence across a small number of cell classes, judged ordinally, by the author. The honest status is: *a candidate organising principle that survived its first two tests — the extreme case and the near-miss — and has not been subjected to a quantitative one.*

2.8 How load could actually be measured

A framework that stays ordinal stays rhetorical. It is worth saying what the quantitative version would look like, because the availability of the measurements is the difference between a metaphor and a variable.

Geometric load is directly measurable from reconstruction: total dendritic and axonal length, membrane area, spine count per cell. Single-neuron reconstruction supplies all three.

Electrical load is measurable as time-averaged firing rate in vivo and, more usefully for this argument, as time-averaged calcium load, for which genetically encoded indicators give a per-cell readout.

Metabolic load is the hardest and the most interesting. What the argument wants is not total ATP consumption but the *fraction of production committed to obligatory maintenance* — the reserve. That is not routinely measured in identified neurons, and the technique to measure it in a single identified cortical cell in vivo does not currently exist. This is a genuine gap and it is the reason Precondition B is the weakest-supported of the four.

Proteostatic load has proxies: endosomal recycling flux, surface receptor turnover rate, and the transcriptomic signatures of the recycling and degradation machinery, all of which are accessible in single-cell data.

Exposure load is the easiest: perineuronal net status by lectin histochemistry is a binary, per-cell, in-tissue measurement that costs one channel.

A composite load index built from these five, computed per cell class, and correlated against per-synapse dispersion at serine 3 across those classes, is the quantitative test of Proposition 3. It requires no new method except the metabolic one, and it could be approximated without that by treating the four measurable loads as the index and stating the omission.

PART III — The Residue

3.1 A spine is a bag of actin

The dendritic spine is not a structural afterthought of the synapse; it is the synapse's mechanical substance. Its head volume correlates with the area of the post-synaptic density, with the number of AMPA-type glutamate receptors it holds, and with the strength of the connection it carries. To change the strength of a cortical synapse is, in the end, to change the volume of a spine, and to change the volume of a spine is to change the balance of its actin.

That balance runs continuously. Actin in a spine is not a scaffold laid down once and left; it is a treadmill, with filaments polymerising at one end and depolymerising at the other, so that the structure persists while its material does not. The persistence of a spine over months (Grutzendler et al., 2002) is therefore not the persistence of its molecules but the persistence of a *rate balance* maintained by a cell that is paying for it, every second, for the life of the animal.

Two consequences follow immediately and both matter later.

The first is that spine stability is an active, funded state, not a passive one. A spine does not decay when the neuron stops paying attention to it; it decays when the neuron stops paying for it. This is the precise sense in which structural memory is metabolically expensive, and the precise sense in which a metabolic failure becomes a structural one.

The second is that any signal that changes synaptic strength must, in the end, change the treadmill. There are many receptors and many second messengers, and they differ enormously in their upstream biology; but the number of ways to change the actin balance in a spine is small, and one of them dominates.

3.2 The switch

Severing is performed by cofilin — in neurons principally n-cofilin, cofilin-1 — the major actin-depolymerising factor of the mammalian nervous system. Cofilin binds cooperatively along ADP-bound actin subunits and, at appropriate occupancy, introduces a twist that fractures the filament, liberating new ends and accelerating turnover.

Cofilin is switched at a single residue near its amino terminus: **serine 3**. Phosphorylation at serine 3 sterically prevents cofilin from binding filamentous actin. Phosphorylated cofilin is therefore inert with respect to the filament, and the network it would otherwise dismantle is stable. Dephosphorylated cofilin binds and acts (Chai et al., 2009; Bamburg et al., 2021).

The residue has writers and erasers, and both classes have neuronal representatives with independent literatures.

Adding the phosphate. LIM-domain kinases 1 and 2 phosphorylate serine 3. LIM kinase is itself activated by the Rho-family effectors, principally Rho-associated kinase (ROCK) downstream of RhoA, and p21-activated kinase downstream of Rac and Cdc42. The pathway RhoA → ROCK → LIMK → phospho-cofilin is the canonical route by which a receptor that raises RhoA activity stabilises actin.

Removing it. The Slingshot family of phosphatases (SSH1, SSH2, SSH3) and chronophin dephosphorylate serine 3. Slingshot-1 is the neuronal workhorse, and its regulation is the point at which the electrical load of Part II enters the chemistry: calcium-mobilising stimuli activate SSH1L and produce cofilin dephosphorylation, the activation is blocked by calcineurin inhibitors or a dominant-negative calcineurin, and knockdown of SSH1L abolishes the response (Wang et al., 2005). Calcineurin dephosphorylates SSH1L directly and raises its cofilin-phosphatase activity in cell-free assay.

So the arithmetic of the residue is: **calcium activates the eraser; Rho-family tone drives the writer.** A neuron that is electrically busy is, other things being equal, pushing serine 3 toward the dephosphorylated, severing state; a neuron under Rho tone is pushing it toward the phosphorylated, stable state. Both are physiological. Both are used. The healthy spine visits both.

3.3 Why the residue is the natural place to look

Three properties, and each is independent of the others.

It is a convergence point by construction. Essentially every receptor system that alters spine structure must eventually alter the actin network, and the great majority do so through Rho-family GTPases and thence through this residue. A signal that changes spine volume and does *not* pass through cofilin is the exception, not the rule. This is a statement about cell biology, not about Alzheimer's disease, and it would be true if the disease did not exist.

It is directly implicated in this disease in human tissue. Rod-shaped inclusions containing cofilin and actin are prominent in hippocampal and cortical neurites of post-mortem Alzheimer brain, and are most prominent in neurites contacting amyloid deposits (Minamide et al., 2000). Such rods are present in Alzheimer brain and not in normal brain (Bamburg and Bernstein, 2016). Whatever else is true, this residue's effector forms visible pathological structures in the disease and does not form them in health.

It is druggable now. Rho-kinase inhibitors exist, at least one is clinically available, and the class is under active investigation for this indication (Zheng et al., 2025). Whatever is true at serine 3 has immediate consequences for what is given to patients, which is why Part XIV is written in the

imperative rather than the conditional.

3.4 Which neuron

Cofilin is ubiquitous and Alzheimer's disease is not, so a claim at this residue is worthless until it names a cell. Reading the five preparations of Part IV side by side, they are found to have converged on **one cell type without any of them saying so**: the excitatory glutamatergic pyramidal neuron of neocortex and hippocampus, at its dendritic spines.

The spine-density experiments are performed on layer 5 pyramidal neurons, and the receptor and its complement ligand colocalise at *excitatory* synapses of human cerebral cortex (Brott et al., 2025). The phospho-cofilin measurement is made in a post-synaptic-density-enriched synaptosome fraction with a functional readout of GluA1 insertion — a glutamatergic synapse by definition (Rush et al., 2018). The plasticity deficits are hippocampal long-term potentiation and visual-cortical ocular dominance (Kim et al., 2013). The integrin work reports depletion of F-actin-associated post-synaptic proteins in hippocampal neurons and rescue of long-term potentiation deficits by cofilin reduction (Woo et al., 2015a; Woo et al., 2015b). And the reelin arm terminates on the same cell: Disabled-1, the adaptor through which the entire reelin signal is transduced, is expressed predominantly in pyramidal neurons (Pesold et al., 1999), and ApoER2 sits in the post-synaptic density of excitatory synapses in complex with the NMDA receptor (Beffert et al., 2005).

This convergence on one cell is a result of reading the preparations together, not a premise. It is worth stating because it is what makes them commensurable at all — and because it fixes what this paper is not about. The parvalbumin-positive fast-spiking interneuron, the cell most often nominated as the cortex's vulnerable inhibitory element, is characteristically aspiny or sparsely spiny and receives excitatory input on dendritic shafts and soma. The mechanics described here are the mechanics of a spine and do not transfer to it directly. Part VIII sets out how that cell nevertheless bears on the argument, and it is not by having its own serine.

3.5 The fact the paper turns on

Cofilin's relationship to filament stability is not a line. It is a curve that turns over.

At low ratios of cofilin to actin, cofilin severs: sparse decoration of a filament introduces mechanical discontinuities at the boundaries between decorated and undecorated segments, and the filament fractures there. At high ratios, cofilin stabilises: a fully decorated filament has no such boundaries, and dense cofilin binding both saturates the severing mechanism and, at the extreme, bundles filaments into ordered aggregates (Bamburg and Bernstein, 2016).

Severs at low ratios; stabilises and bundles at high ratios. This is established biochemistry, it is stated in the same review that summarises the disease pathology, and it is the single fact from which the whole of Part V is derived.

The consequence for a signalling argument is severe and, so far as the published record shows, unabsorbed by the field. If an effector's action inverts across its own range, then the question "does this disease raise or lower the effector's activity?" is malformed. Raising it past the turnover point produces one pathology; lowering it produces another; and the *sign* of a bulk measurement tells you which side of the turnover point the sampled population happened to sit on, not what the disease is doing.

3.6 The rod, and what it is for

The cofilin–actin rod deserves separate treatment because it is usually presented as pure pathology and is, on the evidence, something more interesting.

Rods are bundled aggregates of cofilin and actin in a roughly 1:1 molar ratio, formed in neurites under stress. They form spontaneously in neurons overexpressing *active* cofilin, which establishes that dephosphorylation is sufficient to induce them (Minamide et al., 2000). Their persistence disrupts microtubules and degenerates the distal neurite without killing the neuron. Mature rod formation additionally requires oxidation of cofilin to disulfide-linked species, which is why oxidative and energetic stress are the classical inducers (Bamburg and Bernstein, 2016).

Now the part that changes how they should be read. When cultured neurons are transiently stressed by inhibition of ATP synthesis, they form rods rapidly and disassemble them when the insult is removed. For roughly an hour after rods form, **neurites containing them lose mitochondrial membrane potential and ATP more slowly than neurites without them** — because actin in rods is far less dynamic than filamentous actin elsewhere, and sequestering it slows filament turnover and the ATP hydrolysis that turnover costs (Bernstein et al., 2006).

The rod is, in its first hour, an **energy-saving device**. It is the neuron shutting down its most expensive structural process to keep its mitochondria alive. Only if the stress persists, or recurs within a day, does the rod become the lesion that degenerates the distal neurite.

Three things follow, and they run through the rest of the paper.

The relationship between energy and this residue is not incidental but architectural. The neuron's emergency response to energy failure is executed *at cofilin*, by the same dephosphorylation that in gentler circumstances remodels a spine. Metabolic load and structural control are not two systems that happen to interact; they are one system read at two time-scales.

Pathology here is a matter of duration, not of kind. The transient rod is adaptive; the persistent rod is destructive; the molecular event is identical. This is the same structure as the paper's central claim about the residue — a transient excursion is physiology, a maintained displacement is disease — and it is encouraging that the two arguments have the same shape when

they were derived from different data.

A cell that spends its life near its energy ceiling will make rods more readily and disassemble them less completely. That is the load argument arriving at the residue, and it is stated here as an inference (M5) with the falsifier given in Part XVI.

PART IV — Five Laboratories at One Serine

What follows sets out each finding in its own terms, with the direction of the change stated explicitly, and with the maturity grade attached. The reader should keep one column in mind throughout: *does this preparation include human tissue?* Three of the five do, and it is the human ones that conflict.

4.1 Reelin adds the phosphate, and this stabilises

Chai and colleagues showed that reelin signalling leads to serine-3 phosphorylation of n-cofilin; that phosphorylation at serine 3 renders n-cofilin unable to depolymerise filamentous actin and thereby stabilises the cytoskeleton; and that the chain runs through the lipoprotein receptor ApoER2, the adaptor Disabled-1, Src-family kinases and PI3-kinase. Phosphorylation was localised to the leading processes of migrating neurons as they approached the reelin-containing marginal zone, and immunostaining for phospho-cofilin in dissociated reeler neurons rose significantly after incubation in reelin-containing medium (Chai et al., 2009).

The physiological reading the authors offer is important and is retained here in full, because it is the seed of this paper's central claim: reelin-induced stabilisation *anchors* the leading process. It is a stop signal delivered at a place. Its value lies in being local, and in being transient — a leading process that were permanently anchored would be a migration that never finished.

Two structural facts about this arm must be recorded, because both are routinely lost when reelin is discussed as a generic protective factor.

The cell that makes reelin is not the cell that responds to it. In adult cortex reelin is expressed primarily in GABAergic interneurons and secreted extrasynaptically into the perineuronal matrix, but only by a defined subset: bitufted, horizontal and Martinotti cells expressing neuropeptide Y or somatostatin are reelin-positive, a small number of calbindin cells are, and **none of the parvalbumin-expressing cells are** — basket and chandelier cells are, in the authors' words, often immunopositive to parvalbumin but never to reelin. Disabled-1, through which the signal is read, is expressed predominantly in pyramidal neurons (Pesold et al., 1999). The signal is written by one interneuron class into a shared extracellular compartment and read by a different, excitatory cell. The residue in dispute is on the reader.

The arm is not confined to development. Reelin signalling promotes the development of dendritic spines in hippocampal pyramidal neurons (Niu et al., 2008); ApoER2 resides in the post-synaptic density of excitatory synapses in a functional complex with NMDA receptors, and an

activity-dependently spliced exon in its intracellular domain is required for reelin-induced potentiation and for normal learning (Beffert et al., 2005); and reelin binding to ApoER2 and VLDLR induces Disabled-1 tyrosine phosphorylation and modulates tau phosphorylation (Hiesberger et al., 1999). The adult synaptic role is established independently of the migration role.

Direction: phosphorylation up. Consequence: stabilisation. Context: developmental for the cofilin measurement; mouse and culture. Maturity M4.

4.2 Amyloid removes the phosphate, through an immune receptor

Kim and colleagues reported that murine PirB and its human orthologue LilrB2, present in human brain, are receptors for soluble amyloid-β oligomers with nanomolar affinity; that the first two extracellular immunoglobulin domains mediate the interaction; and that engagement leads to **enhanced cofilin signalling, also seen in human Alzheimer brains**. In mice the deleterious effect of amyloid oligomers on hippocampal long-term potentiation required PirB, and in a transgenic model PirB contributed to adult memory deficits and mediated loss of synaptic plasticity in juvenile visual cortex (Kim et al., 2013).

Enhanced cofilin signalling means more active cofilin, which means **less** phosphate at serine 3.

The same laboratory has since shown that C4d, a complement cleavage product of previously unknown function, binds LilrB2 and PirB with nanomolar affinity; that C4d and LilrB2 colocalise at excitatory synapses in human cerebral cortex and with amyloid-β in Alzheimer's disease; that both C4 and C4d increase with age and more so in Alzheimer's; and that infusing C4d into wild-type mouse cortex significantly reduces dendritic spine density, with the loss **completely prevented** by knockout of PirB (Brott et al., 2025).

That last result matters beyond its own frame. It identifies a ligand at this receptor that is not amyloid, that rises with normal ageing, and that is sufficient on its own to strip spines. The receptor is therefore not an amyloid receptor that happens to bind other things; it is a receptor with at least two endogenous ligands whose combined occupancy rises across a lifetime.

Direction: phosphorylation down. Consequence: destabilisation. Context: mouse, with human tissue corroboration. Maturity M3 for the mechanism, M2 for the human colocalisation and elevation.

4.3 An independent laboratory reproduces the dephosphorylation

Kawaguchi and colleagues, working on an unrelated problem — an endogenous antagonist of the Nogo receptor and PirB called LOTUS — showed that LOTUS inhibits amyloid-β binding to PirB; that in cultured hippocampal neurons from LOTUS-overexpressing transgenic mice **amyloid-induced dephosphorylation of cofilin** and amyloid-induced loss of PSD-95 were both suppressed; that the amyloid-induced fall in dendritic spine density was improved; and that human LOTUS inhibits amyloid binding to human LilrB2 in the same way (Kawaguchi et al., 2022).

This matters as corroboration of the *direction*, not as therapy. A different laboratory, a different country, a different tool — a natural competitive antagonist rather than a knockout — reproduced the specific chain: block the ligand at the receptor and the dephosphorylation does not occur.

Direction: phosphorylation down. Consequence: destabilisation. Independent of the originating group. Maturity M4.

4.4 A third route to removal: β1-integrin and Slingshot

A fourth laboratory reached the same direction by an entirely separate receptor, and this line has not, to the author's knowledge, been read alongside the PirB work.

Woo and colleagues found that amyloid-β42 oligomers bind with high affinity to low- or intermediate-activation conformers of **β1-integrin**, producing loss of surface β1-integrin and activation of cofilin via **Slingshot homology-1**. Conditional loss of β1-integrin prevented amyloid-induced cofilin activation; allosteric modulation or activation of β1-integrin reduced oligomer binding to neurons and blocked oligomer-induced reactive oxygen species production, mitochondrial dysfunction, depletion of filamentous actin and focal vinculin, and apoptosis. Cofilin, in turn, was required for the amyloid-induced loss of surface β1-integrin and for depletion of actin-associated post-synaptic proteins. Reduction of SSH1 prevented mitochondrial translocation of cofilin, and — the human observation — **Alzheimer brain mitochondria contained significantly increased activated and oxidised cofilin**. In the mouse model, brains contained increased SSH1–cofilin and decreased SSH1–14-3-3 complexes, and genetic reduction of cofilin rescued synaptic protein loss, gliosis, long-term potentiation deficits and contextual memory (Woo et al., 2015a). A companion study localised an upstream scaffold: endogenous RanBP9 positively regulates SSH1 levels and mediates amyloid-induced cofilin translocation to mitochondria and cofilin–actin pathology in cultured cells, primary neurons and in vivo (Woo et al., 2015b).

The review from the same group states the two-armed structure explicitly: amyloid signals to cofilin through LIM kinase-1 *and* through Slingshot homolog-1, with additional inputs from β-arrestin, RanBP9, chronophin, phospholipase D1 and 14-3-3 (Kang and Woo, 2019).

Direction: phosphorylation down. Consequence: destabilisation, mitochondrial translocation, rod pathology. Independent receptor, independent laboratory. Maturity M4 for the mechanism, M2 for the human mitochondrial observation.

4.5 Amyloid adds the phosphate, in human Alzheimer cortex

Rush and colleagues, working in Grenoble with no shared authorship with any of the above, reported **elevated** phospho-cofilin-1 in the post-synaptic-enriched fraction of synaptosomes from cortical samples of APP/PS1 mice **and of human Alzheimer cases**. In primary cortical neurons, amyloid-β oligomers induced rapid actin **stabilisation** and **increased** phospho-cofilin within thirty minutes. Fluorescence recovery after photobleaching and calcium imaging in neurons expressing active or inactive cofilin mutants indicated that cofilin phosphorylation is **necessary and sufficient** for amyloid-induced synaptic impairment via actin stabilisation, occurring *before* the formation of cofilin–actin rods. The clinically available Rho-kinase inhibitor fasudil prevented the actin stabilisation, the synaptic impairment and the synaptic loss by blocking cofilin phosphorylation; amyloid also blocked potentiation-induced insertion of the AMPA-receptor subunit GluA1 in a fasudil-sensitive manner (Rush et al., 2018).

The paper opens by stating that how amyloid-β affects the actin cytoskeleton "remains unknown and **contentious**," and notes that others have described increased cofilin phosphorylation in Alzheimer patients. It does not cite or address the opposite finding from the receptor literature.

Direction: phosphorylation up. Consequence: stabilisation. Context: human Alzheimer cortex plus mouse and culture. Maturity M2 for the human measurement.

4.6 And the rods require the phosphate to be absent

Minamide and colleagues showed that mediators of neurodegeneration induce rod-like inclusions of cofilin and actin in axons and dendrites, that rods form spontaneously in neurons overexpressing **active** cofilin — indicating that activation by dephosphorylation is sufficient to induce them — and that persistent rods disrupt microtubules and degenerate the distal neurite without killing the neuron (Minamide et al., 2000). Rods are present in Alzheimer brain and not in normal brain, require oxidation of cofilin for their mature form, sequester cofilin, block transport, and exacerbate mitochondrial membrane-potential loss (Bamburg and Bernstein, 2016; Bamburg et al., 2021).

Direction: requires phosphorylation down. Consequence: severing first, then bundling and blockade. Maturity M2 for rods in human Alzheimer brain, M4 for the induction mechanism.

4.7 The contradiction, tabulated

	writes to Ser3	direction	actin result	human data?	source
Reelin → ApoER2 / Dab1	yes	**phosphorylate**	stabilise	no	Chai 2009
Aβ → LilrB2 / PirB	yes	**dephosphorylate**	destabilise	yes, AD brain	Kim 2013
Aβ → PirB, antagonist test	yes	**dephosphorylate**	destabilise	human LilrB2 binding	Kawaguchi 2022
Aβ → β1-integrin → SSH1	yes	**dephosphorylate**	destabilise, rods	yes, AD mitochondria	Woo 2015a
Aβ → ROCK → LIMK	yes	**phosphorylate**	stabilise	yes, AD cortex	Rush 2018
Cofilin–actin rods	require low pSer3	**dephosphorylate**	bundle, block transport	yes, AD brain	Minamide 2000

Rows two and five are the same ligand, the same residue, the same disease, and opposite signs, each supported by measurement in human Alzheimer tissue. Rows three and four are independent corroborations of row two's direction by two different receptors and two different laboratories. Row one is the physiological system that produces the same direction as row five, in a healthy brain, for a purpose.

The count is worth stating because it changes what a fair reading looks like. **Three independent laboratories, using four different tools and three different receptors, report that amyloid removes the phosphate. One laboratory reports that amyloid adds it, and reports the addition to be necessary and sufficient for the functional deficit.** That is not a lopsided contest that can be settled by majority: the single dissenting result is the one with the cleanest functional causality and the most direct human synaptosome measurement, and it has not been challenged.

4.8 The four possible resolutions

There are only four ways this can be resolved. Three of them are commonly assumed, and none of the three is sufficient.

One of the measurements is wrong. Possible; not demonstrated. Both use appropriate methods — a fractionated synaptosome preparation from human cortex on one side, cofilin signalling in human Alzheimer brain and activated cofilin in Alzheimer mitochondria on the other. Neither has been retracted or failed replication, and the dephosphorylation direction has independent support from two further laboratories.

They are measuring different compartments. Almost certainly true in part. A post-synaptic-density-enriched synaptosome fraction is not the same object as whole neurites, bulk

cortex, or a mitochondrial fraction. But this is a description of the problem rather than a resolution: it says the answer depends on where you look, which is precisely the thing that needs explaining. Why should the answer depend on where you look, if the disease has a direction?

They are measuring different disease stages. Rush and colleagues explicitly place their phosphorylation event *before* rod formation, and rods require dephosphorylation. A temporal sequence — phosphorylation early, dephosphorylation later — is the most attractive of the conventional resolutions and may well be part of the answer. It does not by itself explain why a bulk measurement in end-stage human cortex should come out on the phosphorylated side when the rods are demonstrably present by then. Nor does it explain the third and fourth rows, in which the dephosphorylation is produced acutely, in culture, on the same timescale as the phosphorylation in the fifth.

Or the residue has no pathological direction. This is the possibility the field has not taken up, and it follows directly from a fact already established in the same literature by the same people.

4.9 Why the disagreement was not noticed

A conflict of this size, at a named residue, in a heavily funded disease, going unremarked for a decade requires explanation. The explanation is not that anybody was careless. It is structural, and it is worth setting out because the same structure hides other conflicts.

The five groups do not share a problem. The Stanford line arrived at cofilin from *developmental plasticity* — MHC class I, ocular dominance, the closure of critical periods. The Grenoble line arrived from *synaptotoxicity pharmacology* — what amyloid does to a spine and what fasudil does about it. The Yokohama line arrived from *axon regeneration*, via an endogenous antagonist of the Nogo receptor. The Tampa line arrived from *integrin and mitochondrial biology*. The Fort Collins line arrived from *actin biochemistry*, and had been describing rods for fifteen years before amyloid entered the picture. Five entry points, five vocabularies, five sets of reviewers.

The direction is often stated indirectly. The 2013 paper reports "enhanced cofilin signalling." That phrase is correct and it is not the string a reader searching for phospho-cofilin direction would find. To know that it means dephosphorylation, one has to know the biochemistry and read the phrase as a claim about it. A search-based literature review does not retrieve it as a contradiction of "increased phospho-cofilin-1."

The measurements are of different objects with the same name. "Cofilin phosphorylation in Alzheimer brain" denotes, across these papers, a western blot of a post-synaptic-density-enriched synaptosome fraction; an immunohistochemical impression of cofilin signalling in tissue; and a measurement in a mitochondrial fraction. These are three quantities. The literature calls them one.

And the compartment defence is available to everyone. Any author encountering a conflicting direction can attribute it to fraction, region or stage, and be partly right, and move on. That defence is individually reasonable and collectively fatal: it lets a field accumulate opposite results indefinitely without anyone being obliged to reconcile them.

The reason to state this is not to allocate blame but to make a methodological point that the rest of the paper depends on. **A convergence architecture is worth building precisely because it forces adjacent literatures into the same frame, and adjacency is where unnoticed contradictions live.** The contradiction in Section 4.7 was not discovered by a new experiment. It was discovered by putting five preparations in one table and insisting that the direction column be filled in.

PART V — The Reconciliation, and What Governance Costs

5.1 A non-monotonic effector under lost governance

The fact is Bamburg's, and it is not in dispute: cofilin **severs filaments at low cofilin-to-actin ratios and stabilises them at high ratios** (Bamburg and Bernstein, 2016). The relationship between active cofilin and actin stability is a curve that turns over.

Three consequences follow, and together they dissolve the contradiction of Part IV.

First, both directions of change are pathological, by different routes. Drive serine-3 phosphorylation up and cofilin is withdrawn from the filament: the network over-stabilises, turnover stops, and the structural plasticity a spine needs in order to potentiate is lost — which is precisely the phenotype Rush and colleagues measured, including the failure of GluA1 to insert after potentiation. Drive phosphorylation down and active cofilin rises: filaments sever; then, as occupancy climbs past the turnover point and oxidation cross-links the protein, the same molecule bundles actin into rods that block transport and strangle the distal neurite — precisely the phenotype Minamide, Bamburg, Woo and colleagues described. **There is no direction of travel from the healthy state that is safe.**

Second, a bulk measurement of a non-monotonic system is uninterpretable. Suppose two populations of excitatory synapses in the same cortex. At one, LilrB2 is occupied — by oligomeric amyloid-β, and with age and disease by C4d — and cofilin is driven toward the dephosphorylated, active state. At another, Rho-kinase tone dominates and the residue is driven the other way. A homogenate, a synaptosome fraction, or a field-average fluorescence intensity reports the mean of the two. The mean moves in whichever direction the more abundant population dictates, and which population is more abundant depends on region, fraction, stage and case series.

This is not a hypothetical about hypothetical populations. The sorting has been imaged directly: LilrB2, C4d and oligomeric amyloid-β colocalise at a **subset** of excitatory synapses, particularly near plaques, and at synapses further from plaques amyloid is not detected (Brott et al., 2025). The geometry was quantified independently for synapse loss: a roughly 60 per cent deficit within the oligomeric halo, recovering toward control by about 50 micrometres (Koffie et al., 2009). If the ligand is distributed unevenly across synapses, and the effector inverts across its own range, then the mean is not merely noisy. **It is the wrong statistic**, and two careful laboratories sampling differently will publish opposite signs and be unable to reconcile them.

Third, and this is the claim: the lesion is the loss of range, not the displacement of the mean. In the healthy spine, serine 3 is driven in both directions on demand, transiently and locally — which is exactly what the reelin result describes, a stop signal delivered at a place and then released. Potentiation requires a brief local excursion into severing and a return. The pathological state is one in which the residue can no longer be driven both ways: it is pinned, in different directions at different synapses, by chronic ligand occupancy that no longer encodes anything.

The disease is not that cofilin is too active or too inactive. It is that cofilin has stopped being a signal and become a setting.

This reading has the property a good reconciliation should have. It makes both disputing groups right about their own preparations, and both wrong about the disease.

5.2 What range requires

The proposal above, stated in Part I as Proposition 2, is a signalling argument, and on its own it is incomplete. It says the residue has lost its range without saying why a residue should lose its range, or why in these cells, or why late in life. A mechanism that can be lost must have preconditions that can fail. This section names them, and it is where the paper departs furthest from what has been said before.

To drive a residue transiently and locally and then release it, four things must be true at once.

Precondition A — a source of placed signal. Something must arrive at a particular synapse, at a particular time, and instruct it. The only characterised physiological system that delivers serine-3 phosphorylation as a *placed, transient* event rather than as tone is reelin acting through ApoER2 and Disabled-1 (Chai et al., 2009). Reelin is written into the extracellular matrix by a defined subset of interneurons and read by pyramidal cells (Pesold et al., 1999). If the writing cells fall silent or die, the reader retains its receptors and its adaptor and loses its instruction. Part VII traces that cell's lifetime and states this as the paper's keystone inference.

Precondition B — energy to make the excursion. Following an instruction to sever means paying for a burst of filament turnover, and filament turnover is a material and obligatory ATP cost whose magnitude in neurons is contested but never trivial (Engl and Attwell, 2015; Bernstein et al., 2006). A cell operating near its ceiling can hold a structure but cannot afford to remodel it. Worse, its emergency response to shortfall is itself a cofilin event that removes actin from the dynamic pool (Bernstein et al., 2006) — so energy failure does not merely prevent the excursion, it actively drives the residue toward one end of its range.

Precondition C — receptor return. The instruction lands on a receptor, and the receptor must be on the surface. Surface receptor availability is not a constant; it is the output of an endosomal recycling system whose components are reduced in the Alzheimer brain (Small et al.,

2005; Simoes et al., 2021), and which is specifically impaired for ApoE receptors by the ε4 allele (Chen et al., 2010). A neuron whose ApoER2 is not returned to the membrane cannot be instructed however much reelin the matrix holds.

Precondition D — an exposure regime with a baseline. A signal is defined against a background. If a receptor is continuously occupied by a ligand that carries no information — amyloid-β oligomers accumulating near plaques, C4d rising with age — then the residue is held at a value rather than driven to it, and the excursion has nothing to return to. Exposure is regulated by the extracellular matrix, and Part VIII sets out the two sugars that regulate it and the consequences of their conflation.

5.3 One excursion, described

The argument to this point is abstract, and abstraction is where a claim about "range" can hide. It is worth describing a single healthy excursion concretely, so that the reader can see what the disease is proposed to remove. What follows is a composite from the sources cited throughout Parts III and IV; it is a description of a mechanism, not a report of a measurement, and it is offered as such.

Before. A stable spine on the apical dendrite of a layer III pyramidal neuron. Its actin is treadmilling at a rate that holds volume constant. Serine 3 on the local cofilin pool sits at some intermediate phosphorylation, held there by a balance between the Rho–LIMK arm and the resting activity of Slingshot-1. The spine's surface carries ApoER2 in the post-synaptic density in complex with NMDA receptors (Beffert et al., 2005), an amount of LilrB2, and β1-integrin. The extracellular matrix immediately outside holds heparan-sulfate proteoglycans presenting a low ambient concentration of reelin secreted by an interneuron whose axon passes nearby (Pesold et al., 1999).

Seconds 0–1: the instruction. Coincident pre- and post-synaptic activity opens NMDA receptors. Calcium enters. Calcineurin is activated. Slingshot-1 is dephosphorylated and its cofilin-phosphatase activity rises (Wang et al., 2005). Serine-3 phosphorylation on the local cofilin pool falls. Cofilin binds the filament network.

Seconds 1–30: the severing, and what it is for. At the ratio produced, cofilin severs. Severing creates free barbed ends. Free barbed ends are polymerisation sites. This is the counterintuitive core of the biology and it is why cofilin is a growth factor as much as a demolition enzyme: **to enlarge a spine you must first cut its network, because that is how you make places for new filament to grow.** The spine expands. The expansion costs ATP, and it costs it precisely in the currency of Part II's third load.

Seconds 30 onward: the stop. The excursion must terminate, or severing continues past the point where new ends can be used and the network fragments. The stop is where reelin belongs on this reading: an extracellular ligand, presented on the matrix at that location, engaging ApoER2 and Disabled-1, driving serine-3 phosphorylation back up and withdrawing cofilin from the filament

— anchoring the newly enlarged structure (Chai et al., 2009). It is the same operation the same ligand performed on the leading process of this cell before birth, applied to a different structure for the same purpose: *stop here, and hold.*

Minutes to hours: consolidation. GluA1 is inserted; the post-synaptic density enlarges to match the new head volume; the treadmill re-establishes a balance at a higher set-point. The spine is now larger and stable, and the cell has paid for the transition once and now pays a slightly higher standing cost forever.

What each precondition supplies. Precondition A supplies the stop. Precondition B supplies the polymerisation and the pumping. Precondition C supplies the ApoER2 on which the stop lands. Precondition D supplies the fact that the receptors were unoccupied before the event, so that the calcium-driven fall in phosphorylation was a departure from a baseline rather than a nudge to a value already pinned.

And what the failure modes look like. Remove A and the excursion does not terminate cleanly: severing runs long, the network fragments, and, with oxidation, rods form. Remove B and the excursion cannot be funded: the cut is made and the regrowth is not, or the cell pre-emptively sequesters actin into rods to protect its mitochondria (Bernstein et al., 2006). Remove C and the stop signal arrives at a membrane with nothing to receive it. Remove D and there is no departure to make, because chronic LilrB2 or integrin occupancy has already set the residue low, or chronic Rho tone has already set it high.

Four different-looking pathologies. One lost capacity. **The excursion, described this way, is what "range" means in this paper, and its four preconditions are what the load-bearing neuron progressively cannot supply.**

5.4 The synthesis

Put the four preconditions beside the five loads of Part II and the correspondence is exact, one term at a time.

precondition of governed range	load that erodes it	how it fails
A — source of placed signal	(a property of the *writer* cell's load)	the somatostatin interneuron that secretes reelin is itself heavily loaded and netless; its peptide is the oldest non-cholinergic deficit in the disease
B — energy for the excursion	metabolic load	obligatory maintenance consumes the reserve; shortfall drives rod formation, which is a cofilin event
C — receptor return to surface	proteostatic load	endosomal recycling declines; ApoE4 selectively impairs ApoE-receptor recycling
D — a baseline to return to	exposure load	matrix shielding fails; chronic ligand occupancy rises with age
—	geometric load	sets the *size* of the bill for all of the above
—	electrical load	writes to the residue directly, through calcium → calcineurin → Slingshot-1

The claim of this paper, in its complete form, is that **the loss of range is not a mysterious property of a diseased residue. It is the arithmetic consequence of a heavily loaded, never-replaced neuron running out of the four things governed control requires.** The residue is where the shortfall becomes visible, because the residue is where the loads are settled.

This is why the paper is organised as biography. If the argument is right, the interesting object is not the residue at a moment but the residue across ninety years, in cells that have been paying for it the whole time.

5.5 The claim restated, with its status

> *The load-bearing cortical neuron loses governed control of cofilin serine 3 because the four material preconditions of that control — placed signal, metabolic reserve, receptor return, and a shielded exposure baseline — are progressively consumed by the loads the cell carries and cannot shed. The signature of that loss is dispersion, not displacement.*

Status: inference (M5). Untested. Its components are graded individually in Part XVII; the assembly is not a result and is not presented as one. Part XII gives the measurement that would show it to be wrong, Part XIII gives the experiments in the order they should be run, and Part XVI lists the five findings that would remove it.

One honest note on the shape of the argument. A reader may reasonably observe that a proposal invoking four preconditions is harder to refute than one invoking a single mechanism, because the failure of any one precondition can be absorbed by the others. That objection is correct and is met in two ways. First, each precondition is stated with its own falsifier, so they can be

attacked separately rather than as a bloc. Second, and more importantly, the *central* prediction of the paper — elevated per-synapse dispersion where the mean is unchanged — does not depend on any of the four. It follows from Proposition 2 alone. If that prediction fails, the four preconditions are irrelevant, because there will be nothing for them to be preconditions of.

PART VI — The Reader: ninety years of a pyramidal neuron

The cell that carries the residue is a projection neuron of layer III or layer V of association neocortex, or its hippocampal equivalent in CA1. It is born before the person is. It is never replaced. Everything the disease will eventually do to it is done to a cell that has already been running for eight decades.

This part follows one such cell from the ventricular zone to the autopsy table, and it is organised around a single question asked repeatedly: *at each stage, what does this cell's control of serine 3 cost, and who is paying?*

6.1 The first thing this cell ever does with serine 3 is stop moving

Before it has a spine, before it has an axon, before it has a name in any circuit, the future pyramidal neuron is a migrating cell climbing a radial glial fibre toward the cortical plate. Its leading process must extend, and then, at the right moment, it must stop.

The stop signal is reelin, secreted by Cajal–Retzius cells into the marginal zone. It acts through ApoER2 and VLDLR, through the adaptor Disabled-1 and Src-family kinases and PI3-kinase, and its mechanical effect is **phosphorylation of n-cofilin at serine 3**, which withdraws cofilin from filamentous actin and stabilises the cytoskeleton of the leading process — anchoring it (Chai et al., 2009).

It is worth pausing on this, because it is the deepest fact in the paper and it is usually treated as a piece of developmental trivia.

The very first use this cell makes of cofilin serine 3 is to receive a placed, transient instruction from an extracellular protein written into the matrix by another cell. Not a cell-autonomous program; not a tonic set-point; an instruction, delivered at a location, that stabilises actin *there* and is then released so that the process can be remodelled again. The grammar of the residue — placed, transient, externally written, matrix-mediated — is established before birth, and it never changes. What changes, over ninety years, is whether the grammar can still be spoken.

The same ligand, the same receptor, the same adaptor and the same residue reappear in the adult synapse. Reelin signalling promotes dendritic spine development in hippocampal pyramidal neurons (Niu et al., 2008). ApoER2 is a component of the post-synaptic density of excitatory

synapses, in complex with the NMDA receptor, and an alternatively spliced exon in its cytoplasmic tail — spliced in an activity-dependent manner — is required for reelin-enhanced potentiation and for normal performance in learning and memory tasks (Beffert et al., 2005). Reelin acting through these receptors also modulates the phosphorylation of tau (Hiesberger et al., 1999).

So the adult synapse is not merely analogous to the migrating growth cone. It is the same apparatus, redeployed, with the source of the ligand handed from the Cajal–Retzius cell to a class of cortical interneuron (Pesold et al., 1999). Part VII follows that handover, because when the second writer fails there is no third.

6.2 Childhood: the overproduction

By the first years of life the cell has its arbor and is building spines at a rate it will never approach again. In human prefrontal cortex, dendritic spine density on layer IIIC cortico-cortical and layer V cortico-subcortical projecting pyramidal neurons **exceeds adult values by two- to threefold in childhood** (Petanjek et al., 2011).

Two implications for the argument.

The geometric load is front-loaded and then partially discharged. The cell builds far more structure than it will keep. Every one of those spines is dynamic actin; every one costs turnover; and the childhood cortex is, on any reading of the energetics, spending heavily on structure it is going to throw away.

The cell's control system is calibrated during a period of maximal turnover. Whatever set-points, receptor densities and phosphatase-to-kinase ratios the cell settles on, it settles on them in a regime of rapid remodelling. It will then run those settings for eighty years in a regime of near-total stability. Nothing in the biology guarantees that a control system tuned for growth degrades gracefully into a control system for maintenance.

6.3 The long prune, and who does the cutting

Overproduction is followed by elimination. It was long held that cortical synaptic pruning is complete by early adolescence. In human prefrontal cortex it is not: spine density begins to decrease at puberty and **overproduction and developmental remodelling, including substantial spine elimination, continue beyond adolescence and throughout the third decade of life** before stabilising at the adult level (Petanjek et al., 2011).

The cell is therefore under active, externally directed structural subtraction for roughly its first twenty-five years. What performs the subtraction is the part that matters here, because it is the same machinery that will reappear at the other end of life.

Complement. The classical complement cascade mediates central-nervous-system synapse elimination; complement components localise to developing synapses and are required for their pruning (Stevens et al., 2007). Complement component 4 in particular localises to neuronal synapses, dendrites, axons and cell bodies in human brain, and mediates synapse elimination during postnatal development in mice; common structural alleles of C4 associate with schizophrenia in proportion to their tendency to generate greater C4A expression (Sekar et al., 2016). Microglia execute the engulfment.

Major histocompatibility class I and its receptors. Class I MHC molecules are expressed by central neurons under the control of neural activity (Corriveau et al., 1998), regulate retinogeniculate refinement and limit ocular dominance plasticity (Datwani et al., 2009), and co-regulate synapse elimination and learning rules (Lee et al., 2014). Their receptor PirB restricts ocular-dominance plasticity throughout life; in mutant mice lacking functional PirB, cortical ocular-dominance plasticity is more robust at all ages (Syken et al., 2006), and blocking PirB in the adult up-regulates spines and functional synapses, unlocking visual cortical plasticity (Bochner et al., 2014).

The matrix. The organisation of chondroitin-sulfate proteoglycans into perineuronal nets coincides with the end of the critical period, and degrading them with chondroitinase-ABC in the adult rat restores ocular-dominance plasticity (Pizzorusso et al., 2002).

Three systems, three literatures, one function: **the cortex closes.** It closes by pruning what it overbuilt, by installing receptors that restrain further change, and by depositing a matrix that limits access to the neuronal surface.

Everything in the list is, at the far end of life, a mechanism of this disease. C4 becomes C4d and strips spines at LilrB2 (Brott et al., 2025). PirB mediates the amyloid oligomer's attack on plasticity (Kim et al., 2013). The net degrades and the neurons it sheathed lose their protection (de Vries et al., 2024). This is not a coincidence and Part XI is written about it.

6.4 Adulthood: forty years of paying for stability

From the third decade the cell enters the state it will occupy for most of its existence. In adult mouse cortex, the overwhelming majority of spines — approximately 96 per cent — remain stable over a one-month interval, with a half-life greater than thirteen months; in young animals within the critical period the figure is approximately 73 per cent, and most of the change is elimination (Grutzendler et al., 2002).

This is the phase in which the load argument does its real work, and it is the phase the disease literature almost never describes, because nothing happens in it. Nothing happening is the point.

Stability is not the absence of activity; it is the balance of two large opposed activities. The actin in a stable spine is treadmilling continuously. A structure whose molecules turn over while the structure persists is a structure being paid for. Over forty years, the integral of that payment is the dominant term in the cell's structural economy — larger, in total, than the building of the arbor and larger than its eventual dismantling.

The residue is being held, not left alone. Holding a spine at a stable volume means holding the severing–stabilising balance within a narrow band, which means the kinase and phosphatase arms are both running. A stable spine is not a spine where nothing writes to serine 3; it is a spine where the writers and erasers are in equilibrium.

Every excursion from that equilibrium is a purchase. Potentiation requires the spine to remodel: a transient increase in severing to liberate barbed ends, followed by polymerisation and a return to a new, larger stable balance. The remodelling costs energy, and the return requires an instruction to stop. This is the adult use of the migration mechanism, and it is the thing that fails.

So the healthy fifty-year-old pyramidal neuron is best described not as a stable structure but as **a structure held at a set-point by continuous expenditure, capable of leaving that set-point on instruction and returning to it.** Range is the capacity to leave and return. Both halves cost.

6.5 The drift

Normal ageing does something to this system, and what it does is more specific and more interesting than "loss."

The first correction is that ageing is not principally a matter of losing neurons. Age-related cognitive impairment is linked not to forebrain neuron loss but to specific and relatively subtle synaptic alterations in hippocampus and prefrontal cortex (Morrison and Baxter, 2012; Morrison and Hof, 1997). The cells are there. Something about their synapses has changed.

The second correction is more pointed. Chronic in vivo imaging of spines and boutons in mouse somatosensory cortex over a year found that density is **stable** between mature and old animals — but that the *dynamics* are not. In old mice, new spines and boutons are **two to three times more likely to be stabilised over thirty days**, while the long-term retention of already-stable spines over months is **lower**. Spines in old animals are smaller on average, though still capable of making synaptic connections regardless of size (Mostany et al., 2013).

Read that result through the framework of this paper and it is startling. The aged cortex **stabilises new structures more readily and holds old ones less well**. That is not a quantitative decline in a single parameter. It is a *loss of discrimination*: the system that decides what to keep and what to release has become less selective in both directions at once. New spines that should

have been transient are retained; old spines that had earned permanence are released.

That is the behavioural signature of a control system drifting from *signal* toward *setting* — precisely the transition Part V proposes, arriving decades before any disease, in a healthy brain, on the same cell class, measured by an independent method for an unrelated purpose.

It would be an overstatement to say that Mostany and colleagues measured the loss of range. They measured spine dynamics, not phospho-cofilin, and the connection between the two is an inference. But the inference is not idle: it predicts that per-synapse dispersion at serine 3 should already be rising in normal ageing, before disease, and that the ageing brain should therefore sit between the young and the demented on the paper's primary endpoint rather than clustering with the young. Part XII states this as a secondary prediction, and it is the cheapest way to embarrass the argument, because if aged controls are indistinguishable from young controls the drift reading is wrong.

6.6 The ligand tide

Against a cell whose control is already drifting, the late decades bring an increase in what arrives at its surface. Four items, each independently established, and their significance is that they are all *chronic* rather than episodic.

Complement rises with age. Both C4 and C4d increase with age, and more so in Alzheimer's disease; C4d binds LilrB2 and PirB with nanomolar affinity, colocalises with LilrB2 at human excitatory synapses, and infused into wild-type mouse cortex is sufficient to reduce spine density — an effect completely prevented by PirB knockout (Brott et al., 2025).

Amyloid-β oligomers accumulate at post-synaptic densities. Oligomeric amyloid associates with post-synaptic densities and correlates with excitatory synapse loss near senile plaques, with a roughly 60 per cent deficit inside the halo recovering toward control by about 50 micrometres (Koffie et al., 2009). The distribution is local and graded, not uniform.

Activity feeds the ligand. Synaptic activity regulates interstitial fluid amyloid-β levels in vivo, and endocytosis is required for activity-dependent release (Cirrito et al., 2005; Cirrito et al., 2008). A cortex whose inhibition is failing produces more of the ligand that occupies the receptor that drives the residue. Part VIII follows that loop.

Tau arrives in the spine. Tau mislocalisation to dendritic spines mediates synaptic dysfunction independently of neurodegeneration (Hoover et al., 2010), and dendritic tau mediates amyloid-β toxicity in mouse models (Ittner et al., 2010). Synapses containing tau oligomers are eliminated by microglia and astrocytes in human Alzheimer disease (Taddei et al., 2023).

The common property is duration. None of these is a signal. Each is a condition. A receptor engaged by a condition rather than by an event does not transmit information; it sets a level.

6.7 The withdrawal

Simultaneously, and this is the half of the story the ligand literature does not tell, the instruction thins out.

Reelin signalling fails while reelin protein rises. In frontal cortex, reelin messenger RNA and both soluble and guanidine-extractable reelin protein *increase* with advancing Braak stage, while ApoER2 expression does not change — and yet reelin-dependent phosphorylation of Disabled-1 is *reduced*. Reelin and amyloid-β oligomers co-immunoprecipitate from human brain extracts and appear in the same size-exclusion fractions; amyloid treatment increases reelin expression but the secreted reelin is trapped with the aggregates; and amyloid reduces reelin's capacity to induce ApoER2 internalisation and processing. Soluble ApoER2 fragments in cerebrospinal fluid correlate with reelin levels in controls but not in Alzheimer's disease, where the fragments diminish (Cuchillo-Ibáñez et al., 2016).

More ligand, less signal. This is the exact biochemical signature the argument requires: the instruction is not absent from the tissue, it is unable to be delivered.

Reelin-expressing cells are lost where the disease starts. In the entorhinal cortex of human amyloid precursor protein transgenic mice there are significantly fewer reelin-expressing pyramidal cells — the major glutamatergic reelin-expressing population — with reduced reelin in the hippocampus that receives their projection; qualitatively similar reductions of reelin-expressing entorhinal pyramidal neurons were found in human Alzheimer brains. Notably, the number of reelin-expressing GABAergic interneurons was **not** altered in either structure (Chin et al., 2007). Entorhinal layer II reelin-immunoreactive neurons selectively accumulate intracellular amyloid early in the disease (Kobro-Flatmoen et al., 2016).

The receptor is not returned. ApoE4 selectively impairs the recycling of ApoE receptors, reducing glutamate receptor function and synaptic plasticity (Chen et al., 2010). Endosomal recycling machinery is reduced in the Alzheimer brain and the vulnerable region depends on a distinct retromer core dedicated to it (Small et al., 2005; Simoes et al., 2021).

6.8 The end

The cell arrives at its ninth decade with: a control system already less discriminating than it was at forty; a surface chronically occupied by ligands that encode nothing; an instruction pathway in which the ligand is present but trapped, the receptor is present but not recycled, and the source cells are diminished; a metabolic reserve consumed by eighty years of holding a structure it cannot afford to remodel; and an emergency response to that shortfall which is itself a dephosphorylation event that removes actin from the dynamic pool.

At that point the residue is pinned. Which way it is pinned depends on the local receptor complement, which varies from synapse to synapse across the same tissue — which is why bulk measurements disagree, and why the prediction of this paper is about variance.

And then the spines go. Synapse loss is the strongest structural correlate of cognitive impairment in the disease (Terry et al., 1991); it is present in frontal cortex biopsies and tracks cognitive severity (DeKosky and Scheff, 1990); it is measurable in CA1 at the stage of mild cognitive impairment, before dementia (Scheff et al., 2007); and it is now visible in the living brain as widespread reduction in synaptic vesicle glycoprotein 2A binding (Mecca et al., 2020).

6.9 What the structure costs, in the numbers we actually have

The load argument would be stronger with an energy budget for a single spine, and no such budget exists at the precision one would want. It is worth setting out what *is* known, because the gap is instructive and because a reader is entitled to see the weakest quantitative link stated rather than implied.

What is established. Just under half the brain's energy is spent on processes not directly attributable to signalling (Engl and Attwell, 2015; Harris et al., 2012). Within the signalling half, the dominant costs are post-synaptic receptor currents and the sodium–potassium pumping that reverses them. Within the non-signalling half, the candidate large items are lipid synthesis, mitochondrial proton leak, protein synthesis, microtubule treadmilling and actin treadmilling.

What is contested. The cost of actin treadmilling specifically. The published range spans from under 1 per cent of the brain's global energy budget to roughly one-half of neuronal energy use (Engl and Attwell, 2015), the high end coming from the cofilin literature's own estimate (Bernstein et al., 2006). A two-order-of-magnitude uncertainty in a quantity this central is unusual and it should be treated as a live research question rather than as noise.

What the argument requires. Not the high figure. Three weaker statements suffice, and each is safe across the whole published range.

Actin turnover is obligatory. A spine cannot hold its volume without it. This is a structural fact, not an energetic one.

Its cost scales with dynamic actin mass. More spines, or larger spines, cost more. This follows from the mechanism regardless of the constant.

Its cost is large enough that the cell responds to energy failure by reducing it. This is the empirical content of the rod result: neurons under ATP-synthesis inhibition sequester actin into rods, and doing so measurably slows the decline of mitochondrial potential and ATP for about an hour (Bernstein et al., 2006). A cell does not evolve an emergency mechanism to economise on a trivial expense.

Where the gap bites. The framework wants a per-cell **reserve** — the fraction of ATP production not committed to obligatory maintenance — and that quantity is not measured in identified cortical neurons in vivo by any current method. Every statement in this paper about a cell being "near its ceiling" is therefore an inference from arbor extent, firing rate and the qualitative fact of the rod response. Part II.8 says so; the ledger grades Precondition B accordingly; and Experiment 9 is the closest available substitute, testing whether excursion amplitude is budget-limited in a preparation where the budget can be manipulated rather than measured.

It is worth being clear about what a failure here would mean. If actin turnover proves to sit at the bottom of the published range, and if excursion amplitude proves independent of energy availability, then the metabolic precondition is decoration and the load framework rests on three columns instead of four. That would weaken Proposition 3 without touching Propositions 1 or 2.

6.10 A closing note on this biography

The final observation of this part is a negative one and it should be said plainly. **Nothing in the biography above requires a plaque.** The drift is measurable in normal ageing. The complement rise is an ageing phenomenon that disease accelerates. The recycling failure is genotype-dependent. The instruction withdrawal begins with cells dying elsewhere. Amyloid enters the account as the largest single contributor to the ligand tide and as the trigger for three of the four routes to the residue — which is a serious role, and a smaller one than it is usually given.

PART VII — The Writer: the somatostatin interneuron

7.1 The cell that writes a signal it cannot read

In the adult cortex, reelin is not made by the cells that use it. It is made by a defined minority of GABAergic interneurons and secreted **extrasynaptically**, into the perineuronal matrix, where it modulates the surrounding tissue non-synaptically. The cells that do this are bitufted, horizontal and Martinotti cells; the reliable markers of the reelin-positive population are neuropeptide Y and somatostatin; a small number of calbindin cells qualify; and the parvalbumin-expressing basket and chandelier cells, in the authors' explicit formulation, are often immunopositive to parvalbumin but **never** to reelin. Disabled-1, the adaptor that transduces the entire reelin signal, is expressed predominantly in pyramidal neurons (Pesold et al., 1999).

Three cells, three roles, and the field's habit of speaking about "reelin in the cortex" as if it were one system obscures all three:

- the **writer** — a somatostatin- or NPY-expressing interneuron, which secretes reelin into the matrix and does not itself read it;
- the **bed** — the extracellular matrix, which holds the reelin and presents it, and whose chemistry Part VIII shows to be the obligate co-receptor;
- the **reader** — the pyramidal neuron, which carries ApoER2, Disabled-1, and serine 3.

The disease literature has studied the reader intensively, the bed occasionally, and the writer almost not at all in this connection. That is the gap this part addresses.

7.2 Where it writes, and why the address matters

The somatostatin-expressing interneuron of cortex is, in the main, a **dendrite-targeting** cell providing feedback inhibition. It is driven by the pyramidal cells it inhibits, and it inhibits them at their distal dendrites — the Martinotti cell in particular ascends to layer 1 and arborises across the apical tufts of pyramidal neurons whose somata lie in the layers below.

That is an unusually precise coincidence of address, and it deserves to be stated as such because it is not usually noticed. **The cell that secretes reelin into the matrix is the same cell whose axon occupies the compartment where the reader's most distal, most plastic,**

most long-range-driven spines are. It is not merely that somatostatin interneurons happen to be reelin-positive; it is that their output — inhibitory, peptidergic, and matrix-borne — converges on the part of the pyramidal dendrite where the cortex integrates its top-down input.

A second architectural fact reinforces it. Somatostatin cells are roughly thirty per cent of the cortical inhibitory population and are the feedback element of the local circuit: their firing is a function of the pyramidal activity they are meant to limit. In control-theoretic terms the somatostatin interneuron *is* the local governor. A governor that also writes the extracellular protein which sets the pyramidal cell's actin brake is doing two jobs with one cell, and both jobs bear on the same dendrite.

7.3 The oldest non-cholinergic finding in the disease

In 1980, four years after the cholinergic deficit was described, somatostatin-like immunoreactivity was reported reduced in the cerebral cortex of Alzheimer cases relative to neurologically normal individuals. The finding was explicitly framed as the first of a systematic survey of neuropeptides, undertaken because the transmitter systems examined to that point — dopamine, noradrenaline, serotonin, GABA — had shown no consistent involvement while acetylcholine had (Davies et al., 1980).

Four and a half decades later that finding has not been overturned. Somatostatin and neuropeptide Y in cerebrospinal fluid correlate with amyloid and tau measures in elderly patients with mild cognitive impairment (Duron et al., 2018), and the peptide has an established mechanistic link to the disease independent of the cell that makes it: somatostatin regulates brain amyloid-β42 through modulation of neprilysin-mediated proteolytic degradation (Saito et al., 2005).

What the finding has never had is a place. It sits in the literature as an orphan fact — reliably reproduced, mechanistically suggestive, and connected to nothing in particular. The proposal of this part is that its place is the one the reelin literature was always implying and never stated: **the somatostatin cell is the adult source of the signal that governs the pyramidal cell's serine 3, and the earliest robust neurochemical evidence in this disease is evidence that the source is failing.**

7.4 The keystone inference

Stated formally, so it can be attacked:

> **If** *the only characterised physiological system delivering placed, transient serine-3 phosphorylation to pyramidal spines is reelin through ApoER2 (Chai et al., 2009; Beffert et al., 2005; Niu et al., 2008),* **and** *the adult cortical source of that reelin is a defined subset of somatostatin- and NPY-expressing interneurons secreting into the matrix (Pesold et al., 1999),* **and** *somatostatinergic markers are reduced in the Alzheimer cortex (Davies et al., 1980),* **then** *the pyramidal neuron in the Alzheimer cortex is losing its instruction while retaining its*

machinery — and a residue that retains its writers and erasers but loses its instruction is precisely a residue that becomes a setting rather than a signal.

Status: inference, M5. The premises are individually graded M4, M4 and M2 respectively. The conjunction has not been tested, and the specific step that has never been performed is stated in the next section.

Two features make this inference worth the risk of stating it.

It supplies a cellular mechanism for a claim that was otherwise unmoored. Part V argues that the residue loses its range. Range requires a source of placed signal. Without a named source whose failure is independently documented, "loss of placed signal" is a restatement of the conclusion. With one, it becomes a claim about a cell that can be counted, and whose reelin output can be deleted on purpose.

It joins two of the disease's most durable observations that have never been joined. The somatostatin deficit (1980) and the reelin resilience allele (2023) belong to different decades, different subfields and different levels of description. On this reading they are the two ends of one arm: the loss of the writer, and the gain of function in what the writer writes. The RELN-COLBOS variant, carried by a man who remained cognitively intact into his late sixties despite a PSEN1-E280A mutation and very high amyloid burden, activates Disabled-1 more strongly and reduces tau phosphorylation in a knockin mouse (Lopera et al., 2023). A stronger signal from a diminished source is exactly the compensation the argument predicts should help.

7.5 The experiment nobody has done

The inference above turns on a step that is entirely tractable and, so far as the published record shows, has never been performed: **conditionally delete reelin from somatostatin-expressing interneurons in the adult mouse and measure serine-3 phosphorylation dynamics at pyramidal spines.**

The tools are standard. The prediction is specific and unusual: the manipulation should **not** produce a large change in baseline phospho-cofilin. It should reduce the *excursion* — the size of the transient serine-3 response to a plasticity-inducing stimulus — while leaving the resting level within the normal range, and it should widen the synapse-to-synapse distribution.

That prediction is worth dwelling on because it is the point at which the paper's central idea becomes experimentally distinguishable from every conventional alternative. Every conventional account predicts that removing a phosphorylating input lowers the mean. **This account predicts that removing the placed source degrades the dynamics while the mean holds**, because the tonic arms — Rho-kinase on one side, Slingshot on the other — remain and will settle the residue wherever their balance puts it. If the deletion simply lowers mean phospho-cofilin and

nothing else, the "signal versus setting" framing adds nothing, and the reelin arm should be treated as one more tonic input.

7.6 The entorhinal complication, stated rather than buried

An honest account must record a result that cuts against the simplest version of the writer claim.

In the entorhinal cortex of hAPP mice and of human Alzheimer brains, the reelin-expressing population that is depleted is **pyramidal** — the major glutamatergic reelin-expressing population — while the number of reelin-expressing **GABAergic interneurons** was not altered in either the entorhinal cortex or the hippocampus (Chin et al., 2007).

Three points, and they should be weighed rather than resolved by assertion.

This is a regional result, not a global one. The entorhinal cortex is anomalous in having a large glutamatergic reelin source; neocortex is not. The claim of this part is about neocortex and hippocampus, where the interneuronal source predominates (Pesold et al., 1999). The Chin result therefore constrains where the writer claim applies rather than refuting it.

Cell number is not secretion. The measurement that failed to change was the *number* of reelin-expressing interneurons. The somatostatin literature reports reduced peptide immunoreactivity, which is a statement about content, not count. A cell that survives and secretes less is invisible to a count and decisive for the argument. Reconciling the two requires measuring interneuronal reelin *output*, which nobody has done.

The result independently supports the paper's larger structure even as it complicates this part. Whichever cell type is losing reelin in the entorhinal cortex, reelin supply to the reader is falling there, in humans, early, at the site where the disease's cortical pathology begins. The writer claim specifies which cell; the withdrawal claim does not depend on getting that specification right.

7.7 What would refute the writer claim

Stated before the argument was assembled and unchanged by it.

1. **Conditional deletion of reelin from somatostatin interneurons in the adult mouse leaves serine-3 dynamics at pyramidal spines intact.** The writer is then not the source of governance and the keystone fails.
2. **Somatostatin-cell reelin output is shown to be preserved in the Alzheimer cortex** while somatostatin peptide falls. The two markers then dissociate and the peptide deficit says nothing about reelin supply.

3. **Reelin is shown to act tonically rather than as a placed transient signal at adult synapses** — for instance if reelin's effect on serine 3 in mature neurons is found to be spatially uniform and slow. The mechanism then has no privileged status over Rho-kinase tone.

4. **The somatostatin deficit is shown to be entirely a terminal-stage phenomenon**, absent in mild cognitive impairment and early disease. The writer would then be failing after the reader, and could not be part of the reader's failure.

Item 1 is the decisive one and it is a two-year mouse experiment with off-the-shelf reagents.

PART VIII — The Shield and the Shielded

8.1 Exposure is a load, and the matrix is what regulates it

Of the five loads in Part II, exposure is the one with no obvious analogue in the ordinary discussion of neuronal vulnerability, and it is the one on which the therapeutic argument of this paper mostly rests.

The reasoning is simple. Every ligand that writes to serine 3 is extracellular: reelin, amyloid-β oligomers, C4d, the myelin-associated inhibitors that also signal through PirB (Atwal et al., 2008). Whether such a ligand acts as a signal or as a setting is not a property of the ligand. It is a property of the *access regime* — how much of the neuronal membrane the ligand can reach, how quickly it can be cleared from the immediate vicinity, and how freely the receptors it binds can move in the plane of the membrane.

The regulator of that regime is the extracellular matrix, and its condensed form is the perineuronal net.

8.2 The parvalbumin interneuron and its net

The parvalbumin-positive fast-spiking interneuron is, on the first four loads of Part II, the most extreme cell in the cortex. It fires at rates an order of magnitude above a pyramidal neuron, admits and extrudes calcium at correspondingly high rates, and has among the highest metabolic demands of any cortical cell. On the load framework alone it should be the first thing to fail.

It is not. Something protects it, and the something is visible in a light micrograph: it is sheathed in a condensed lattice of chondroitin-sulfate proteoglycans, hyaluronan and tenascin-R that assembles at the close of the critical period and persists for life.

The net's function in plasticity is established: its organisation coincides with the end of the critical period, is delayed by dark rearing, and its enzymatic degradation with chondroitinase-ABC in the adult rat restores ocular-dominance plasticity (Pizzorusso et al., 2002). Its function in this disease is separately established: neurons associated with aggrecan-based perineuronal nets are protected against tau pathology in subcortical regions of the Alzheimer brain (Morawski et al., 2010), and in human frontal cortex, resilience to Alzheimer's disease associates with alterations in perineuronal nets, with excitatory neurons bearing a net showing low amounts of phospho-tau (de

Vries et al., 2024).

The net restricts lateral mobility of surface components and limits access to the neuronal membrane. On the framework of this paper that is exactly the property required: **a net-bearing neuron has a regulated exposure regime, and a netless neuron does not.** The parvalbumin cell survives its enormous loads because the fifth is held low.

This is not offered as a full explanation of parvalbumin-cell fate, which is complex and which includes its own dysfunction in models (Verret et al., 2012; Palop and Mucke, 2016). It is offered as the reason the load framework does not collapse on its most obvious counterexample, and as evidence that the fifth column is doing real work.

8.3 The load transfer

The parvalbumin cell enters the argument of this paper in one way only, and it is worth being exact about it because the alternative is a false unification. **The perineuronal net is a modifier of ligand exposure upstream of the residue, not the site of the lesion.** The parvalbumin cell is aspiny or sparsely spiny; the mechanics of Part III are the mechanics of a spine; no claim is made that this cell fails by the mechanism described here, and it probably does not.

What it does is transfer load. The sequence has four steps and each is independently supported:

1. **Net degradation removes the protection of the interneuron it sheathes**, and net alterations are documented in the Alzheimer cortex (de Vries et al., 2024).
2. **The resulting failure of inhibition raises the activity of the pyramidal population.** Inhibitory interneuron deficit links altered network activity to cognitive dysfunction in an Alzheimer model, and network hyperexcitability is a recognised feature of the disease (Verret et al., 2012; Palop and Mucke, 2016).
3. **Heightened pyramidal activity increases local production of amyloid-β.** Synaptic activity regulates interstitial fluid amyloid-β levels in vivo, and endocytosis is required for the activity-dependent release (Cirrito et al., 2005; Cirrito et al., 2008).
4. **More amyloid means more occupancy of LilrB2 on pyramidal dendrites** — alongside the C4d that rises with age and disease at the same receptor (Brott et al., 2025) — which drives the residue.

The matrix is therefore upstream of the ligand, and the ligand is upstream of the residue. This is a load transfer in the literal sense: the failure of one cell's shielding is discharged as an increase in another cell's exposure.

There is a further term that closes the loop, and it is the reason the framework predicts acceleration rather than steady decline. Electrical load writes to the residue directly — calcium

activates calcineurin, which activates Slingshot-1, which dephosphorylates serine 3 (Wang et al., 2005). A pyramidal population that is firing more is therefore driving its own cofilin toward the active state by two independent routes at once: through the ligand it releases, and through the calcium it admits. **Disinhibition is not merely permissive of the ligand attack; it is itself a write to the switch.** Graded M5, inference from established components; the test is whether pharmacological suppression of pyramidal activity in a model reduces per-synapse dispersion at serine 3.

8.4 The net-bearing pyramidal minority

A minority of excitatory neurons bear a perineuronal net of their own, and in human frontal cortex those neurons are conspicuous: **excitatory neurons bearing a net show low amounts of phospho-tau** (de Vries et al., 2024). The same study reports that aggrecan falls in both Alzheimer's disease and in resilient cases, while the *Wisteria floribunda* agglutinin-reactive sugars fall only in the resilient — a dissociation between the net's core protein and its glycan that is itself worth pursuing and is not resolved here.

This population is the only one in the cortex in which spines, a net, and reelin-responsiveness are properties of a single cell. It is therefore the natural test bed for the exposure argument, and it yields the paper's cheapest prediction:

> **Net-bearing pyramidal neurons retain dynamic range at serine 3 while their netless neighbours lose it.**

This costs one antibody. *Wisteria floribunda* agglutinin is the marker used in the human resilience work; adding it as a channel to the imaging experiment of Part XIII tests the prediction in the same tissue and the same run. If net-bearing excitatory neurons show the *same* per-synapse dispersion as netless ones, the matrix plays no protective role at this residue and the exposure arm of the load framework should be dropped.

8.5 Two sugars, and the cost of conflating them

The perineuronal net is frequently invoked in discussions of reelin, and the two are routinely spoken of as one matrix. They are not, and the distinction has a specific and expensive consequence.

The net is a chondroitin-sulfate structure. Its lecticans — aggrecan, brevican, neurocan, versican — carry chondroitin-sulfate chains, and it is chondroitinase-ABC that removes it and reopens plasticity (Pizzorusso et al., 2002; Fawcett et al., 2022).

Reelin's requirement is for heparan sulfate. N-sulfated heparan sulfate is an obligate co-receptor for reelin-induced ApoER2 dimerisation. Full-length reelin binds heparan sulfate with a dissociation constant of 17 ± 5 nanomolar — tightened to 10 ± 2 nanomolar by the

resilience-associated COLBOS variant — N-sulfation is the critical glycan determinant, and heparinase treatment or knockout of the N-sulfation enzyme NDST1 strips reelin from the cell surface. Decisively, heparinase or free heparin in the medium reduces reelin-induced ApoER2 dimerisation, while N-desulfated heparin does not (Pan et al., 2025).

Heparan-sulfate proteoglycans are components of the diffuse extracellular matrix surrounding all central nervous tissue, not constituents of the condensed net (Fawcett et al., 2022).

Three consequences follow, and each is testable with a pair of enzymes.

Reelin does not require a perineuronal net in order to signal. The alternatives to the net for staging it are the ordinary heparan-sulfate proteoglycans of the neuronal surface — glypicans, syndecans, agrin — including those on the responding pyramidal dendrite itself, which needs no external depot at all. The prediction is that reelin signalling survives chondroitinase and fails under heparinase. If chondroitinase abolishes the signal, the sugar chemistry is wrong and the matrix arm needs rebuilding.

The net's protection and reelin's staging are separable functions of one compartment. A neuron can be shielded without being staged, and staged without being shielded. This matters for interpreting the net-bearing pyramidal minority: if those cells are protected, the framework says it is because their *exposure* is regulated, not because their reelin supply is better.

And there is a drug conflict, verified at the level of chemistry, that nobody has stated. Heparan-sulfate proteoglycans are the route by which pathological tau enters neurons; heparan-sulfate proteoglycans mediate the internalisation and propagation of specific proteopathic seeds (Holmes et al., 2013). That finding has made heparin mimetics an attractive strategy for blocking transcellular tau propagation. The two facts collide: **a heparin mimetic given to stop tau spreading would, by the identical chemistry, silence the reelin signal that restrains tau phosphorylation** — because free heparin in the medium is one of the three manipulations shown to block reelin-induced ApoER2 dimerisation, and the block is N-sulfation-dependent (Pan et al., 2025; Hiesberger et al., 1999).

This is not a speculative interaction. It is the same polymer, the same sulfation requirement, and two opposing consequences, and it argues that the sulfated matrix must be **modulated rather than blocked**. Part XIV states what that means for development programmes.

PART IX — The Extreme Case: the coerulean neuron

9.1 The cell that maximises every load

The noradrenergic neuron of the locus coeruleus is the limiting case of the framework in Part II, and it is included here for that reason and not because the argument depends on it.

Score it on the five loads.

Geometric: extreme. A nucleus of some tens of thousands of cells on each side supplies noradrenaline to essentially the entire forebrain, cerebellum and cord. The individual arbor is correspondingly vast and is largely unmyelinated — the expensive configuration, both to maintain and to protect.

Electrical: high and unremitting. The cell is a tonic pacemaker. It does not have quiet periods in the sense that a cortical pyramidal neuron does; it fires continuously across the waking day and modulates rather than ceases.

Metabolic: extreme. The product of arbor extent and firing rate, compounded by the cost of synthesising, packaging and recovering a catecholamine whose oxidation chemistry is itself a liability.

Proteostatic: high. Delivery and retrieval across an arbor of that extent is the hardest logistics problem any neuron in the brain solves.

Exposure: extreme. The nucleus has no perineuronal net.

By the definition of Part II, constructed without reference to the disease, this is the most load-bearing cell in the human brain. And it is, on the independent evidence of neuropathology, the first structure in the brain to accumulate abnormal tau: pretangle material is present in the locus coeruleus of a substantial fraction of people well before the fourth decade, and it appears there before it appears in the transentorhinal cortex (Braak and Del Tredici, 2011; Braak et al., 2011). Quantitative work confirms the accretion of hyperphosphorylated tau in the locus coeruleus and dorsal raphe as the earliest measurable step (Ehrenberg et al., 2017).

The framework and the pathology agree on the extreme case. That agreement is the strongest single piece of support the load framework has, and it is worth stating precisely what it does and

does not establish.

9.2 What the agreement licenses, and what it does not

It licenses: the claim that the ordinal ranking produced by the five loads is not arbitrary, because at its extreme it predicts the cell the pathology independently identifies as first. A framework that had ranked, say, the cerebellar granule cell first would have been refuted by this test. This one was not.

It does not license: the claim that the locus coeruleus is where the disease begins, still less that everything downstream is a consequence of its failure. The inference from *fails earliest* to *causes the rest* is the classic error of the chain-form account, and this paper declines it explicitly.

Three reasons for the refusal, and they are not rhetorical.

Earliest is not the same as causal. Tau in the locus coeruleus by the fourth decade is present in a very large fraction of the population, and the large majority of that population will never develop Alzheimer's dementia. A lesion that nearly everyone has and few people are harmed by cannot, without a great deal of additional argument, be the cause of the disease in the few.

Extreme load explains early failure without any disease-specific mechanism. If the load framework is right, the most heavily loaded cell should fail first *whatever* the disease is, and indeed the locus coeruleus is affected early in more than one neurodegenerative process. Early coerulean failure is therefore weak evidence about the specific mechanism of Alzheimer's disease and strong evidence about the general principle that load predicts order.

The residue argument does not require it. Everything in Parts III through VIII is stated without reference to the brainstem. The reelin writer, the complement ligand, the integrin route, the Rho tone, the matrix chemistry and the energy accounting are all cortical and all independently established. If the locus coeruleus were spared entirely, the argument of this paper would be unaffected.

9.3 Three routes from the nucleus to the residue

Having demoted it, one must still say what the nucleus does, because there are real mechanistic connections from noradrenergic dysregulation to cofilin serine 3. Three, graded, from most to least secure.

Route 1 — through activity and calcium. Noradrenaline is a gain-setting neuromodulator; a cortex whose noradrenergic input is deranged has deranged pyramidal excitability. Raised pyramidal activity admits calcium, calcium activates calcineurin, calcineurin activates Slingshot-1, and Slingshot-1 dephosphorylates serine 3 (Wang et al., 2005). Raised activity also raises interstitial amyloid-β (Cirrito et al., 2005). **Grade: M5, an inference from**

established components. Each step is solid; the conjunction is not measured.

Route 2 — through Rho tone. Rho-associated kinase is the writer arm of the residue, and ROCK1 and ROCK2 are elevated in the Alzheimer brain; pharmacological inhibition of ROCK2 suppresses amyloid-β production in a mouse model (Herskowitz et al., 2013; Zheng et al., 2025). Adrenergic signalling is among the physiological inputs to Rho-family tone. **Grade: M5, weaker than Route 1, because the specific link from cortical noradrenergic state to cortical ROCK activity has not been measured.**

Route 3 — through matrix proteolysis. Sustained β-adrenergic signalling raises expression of matrix metalloproteinase-9 in peripheral cells through cyclic-AMP, protein kinase A, AP-1 and NF-κB; a gelatinase active in the cortical matrix would degrade the lectican core proteins of the net and shed the ectodomains of the heparan-sulfate proteoglycans that stage reelin, and the low-density-lipoprotein receptor family to which ApoER2 belongs is susceptible to ectodomain shedding. **Grade: M5, and the weakest of the three.** The demonstration that noradrenaline induces this protease has been made in peripheral cells rather than in brain parenchyma in situ; that MMP-9 sheds ApoER2 specifically has been shown for the receptor's kin and not for ApoER2 itself; and direct proteolysis of reelin by this enzyme is not established. It is included because it is the only route that would connect noradrenergic state to reelin *staging* rather than to reelin-independent tone, and because it makes a testable prediction — that cortical MMP-9 activity should track dysregulated rather than physiological noradrenergic signalling.

There is a paradox of sign in Route 3 that must be stated rather than smoothed over. Tonic noradrenaline is broadly anti-inflammatory in the brain, and it is classically the *loss* of coerulean noradrenaline that disinhibits neuroinflammation. The resolution offered is that noradrenaline's sign depends on concentration, receptor subtype and temporal pattern: steady moderate tone holds glia quiescent, while the deranged bursting-and-trough pattern of a failing nucleus, superimposed on receptor up-regulation, engages the β-adrenergic arm that potentiates protease induction. The claim attaches specifically to the deranged phase. Correctly stated, the objection becomes one of the route's predictions rather than a refutation of it — but it is a prediction, not a finding, and the ledger marks it so.

9.4 The demotion, stated in one paragraph

The locus coeruleus is the extreme case that validates the load ranking and the source of three mechanistically plausible but individually weak routes to the residue. It is not, in this account, the origin of Alzheimer's disease, and the account would survive its complete removal. Readers who have encountered arguments in which the brainstem does the causal work should note that the difference here is deliberate: **the nucleus is retained for what it proves about load and demoted for what it cannot prove about time.**

9.5 A note on the other early-failing cells

For completeness, and because the census in Part II should not be read as complete, three further populations satisfy the load definition and are affected early.

Entorhinal layer II projection neurons. Classically the origin of cortical tangle pathology; molecularly resolved to a specific excitatory subpopulation (Leng et al., 2021); the site at which reelin-immunoreactive neurons selectively accumulate intracellular amyloid (Kobro-Flatmoen et al., 2016) and at which reelin-expressing pyramidal cells are depleted in models and in human disease (Chin et al., 2007). Stereological work has argued that viable neurons persist in layer II and CA1 in Alzheimer disease longer than the tangle counts suggest (Hof et al., 2003), which is consistent with a synaptic rather than a somatic lesion as the proximate cause of impairment.

Cholinergic basal forebrain neurons. Enormously extended arbors, netless, and the subject of the oldest neurochemical finding in the disease. They are not developed here because they are not spiny targets in the relevant sense and the cofilin argument does not transfer to them cleanly.

Wake-promoting subcortical populations more generally. Profound degeneration of wake-promoting neurons is documented in Alzheimer's disease (Oh et al., 2019), and these are, without exception, extended-arbor, netless, tonically active cells — which is what the framework predicts and is a further, weak, consistency check.

None of these is required by the argument. All are consistent with it. The reader should treat their inclusion as an invitation to test the framework's ranking against populations it was not built on, which is the only way a descriptive framework earns the right to be called predictive.

PART X — The Bridges

10.1 How to read a bridge

A **bridge**, in the sense used here, is a claim of the form: *two bodies of work, developed independently and usually in mutual ignorance, terminate at the same molecular object.* It is not a claim that one causes the other, and it is not a claim about their order in time.

Bridges have three properties that links in a chain do not.

They can be graded individually. Each bridge has its own evidence and its own maturity grade, and the failure of one does not propagate.

They are symmetrical. A bridge from matrix biology to the residue is equally a bridge from the residue to matrix biology, and can be crossed in either direction — which is why the same bridge sometimes yields a prediction about a drug and sometimes a prediction about a cell.

They accumulate without multiplying error. A chain of seven links each 80 per cent likely is a 21 per cent argument. Seven bridges each 80 per cent likely, converging on one residue, is a much stronger argument than any of them, because they are not conditional on one another.

Seven bridges follow. Each is stated, graded, and given a condition that would demolish it.

10.2 The matrix bridge

The claim. The extracellular sulfated matrix is not context for the residue; it is a required component of the signalling system that governs it, and its two polymers do two different jobs.

The evidence. N-sulfated heparan sulfate is an obligate co-receptor for reelin-induced ApoER2 dimerisation: reelin binds it at 17 ± 5 nanomolar (10 ± 2 nanomolar for the COLBOS variant), N-sulfation is the critical determinant, and heparinase, NDST1 knockout, or free heparin all block dimerisation while N-desulfated heparin does not (Pan et al., 2025). Heparan-sulfate proteoglycans are components of the diffuse matrix surrounding all central nervous tissue; the perineuronal net is a condensed chondroitin-sulfate structure (Fawcett et al., 2022). The net's assembly closes the critical period and its enzymatic removal reopens it (Pizzorusso et al., 2002). Net-bearing neurons carry less tau pathology in subcortical regions and less phospho-tau in human frontal cortex (Morawski et al., 2010; de Vries et al., 2024).

What it buys. A mechanism by which exposure — Load 5 — is regulated, and a chemically specific separation between shielding (chondroitin) and staging (heparan). It converts "the matrix is protective" from a slogan into two dissociable claims that a pair of enzymes can tell apart.

Grade: M5 for the chemistry, M2 for the human net–tau association. Demolished if: chondroitinase abolishes reelin-induced serine-3 phosphorylation, which would mean the two polymers are not functionally separable and the chemistry has been misread.

10.3 The receptor bridge

The claim. One residue is written by at least four receptor systems with at least four distinct ligands, and the identity of the receptor engaged determines the direction of the write.

The evidence. ApoER2 with Disabled-1, ligand reelin: phosphorylation (Chai et al., 2009). LilrB2 and PirB, ligands amyloid-β oligomers, C4d, and myelin inhibitors: dephosphorylation (Kim et al., 2013; Brott et al., 2025; Atwal et al., 2008). β1-integrin through Slingshot-1, ligand amyloid-β42 oligomers: dephosphorylation (Woo et al., 2015a). Rho-associated kinase through LIM kinase: phosphorylation (Rush et al., 2018; Zheng et al., 2025).

What it buys. The mechanism of the disagreement in Part IV. Two synapses in the same cortex with different receptor complements are being driven in opposite directions at the same residue by the same ligand pool. That is not noise. It is the predicted behaviour of the system.

It also buys a specific and underappreciated point about LilrB2. It is an immune receptor, its ligands include a complement fragment that rises with normal ageing, and its physiological function is to *restrain* plasticity throughout life (Syken et al., 2006; Bochner et al., 2014; Djurisic et al., 2019). This is a brake being applied harder, not a novel toxic mechanism.

Grade: M2 to M4 by component; the convergence itself is Observation. Demolished if: ApoER2/Dab1 signalling and LilrB2 signalling, co-manipulated in the same neurons, do not interact at serine 3. The convergence would then be coincidental and the arms should be treated separately.

10.4 The energy bridge

The claim. Actin filament turnover is a material, obligatory energy cost of maintaining a spine; cofilin governs it; and the neuron's emergency response to energy shortfall is itself a cofilin event. Metabolic load and structural control are therefore one system.

The evidence. Just under half the brain's energy is spent on non-signalling processes; among these, the cost of actin treadmilling is contested, with published estimates spanning from under 1 per cent of the global budget to roughly half of neuronal energy use (Engl and Attwell, 2015; Harris et al., 2012). Neurons stressed by inhibition of ATP synthesis rapidly form cofilin–actin rods; for

roughly an hour after formation, neurites bearing rods lose mitochondrial membrane potential and ATP **more slowly** than neurites without them, because rod actin is far less dynamic and its sequestration slows turnover and the associated hydrolysis (Bernstein et al., 2006). Persistent rods disrupt microtubules and degenerate the distal neurite (Minamide et al., 2000). Amyloid-driven cofilin activation through Slingshot-1 translocates cofilin to mitochondria, and Alzheimer brain mitochondria contain increased activated and oxidised cofilin (Woo et al., 2015a).

What it buys. A mechanism for Precondition B. A cell at its energy ceiling cannot fund the excursion, and its response to shortfall drives the residue toward the severing end of the range. It also reframes the rod: not simply a lesion, but an adaptive shutdown of the cell's most expensive structural process that becomes a lesion when it persists.

Grade: M4 for the rod energetics; the magnitude of the actin energy cost is explicitly contested and the argument is stated to survive at the low end. Demolished if: the actin energy cost is resolved to the bottom of the published range *and* rod formation is shown to confer no energetic benefit — in which case the metabolic precondition has no mechanism.

10.5 The proteostasis bridge

The claim. Whether a signal can be received depends on whether its receptor was returned to the surface, and receptor return is a machine that fails in this disease and fails harder in the highest-risk genotype.

The evidence. ApoE4 selectively impairs the recycling of ApoE receptors, thereby reducing glutamate receptor function and synaptic plasticity (Chen et al., 2010) — and ApoER2, the reelin receptor, is an ApoE receptor. Retromer components are reduced in the Alzheimer brain (Small et al., 2005), and the vulnerable region depends on a distinct retromer core dedicated to endosomal recycling (Simoes et al., 2021). A nervous-system-specific plasma-membrane proteasome degrades nascent protein in an activity-dependent manner and is required for learning-induced plasticity (Ramachandran and Margolis, 2017; Ramachandran et al., 2018); proteasome downregulation and dysfunction are early events in Alzheimer proteostasis failure (Jiang et al., 2025). Ephexin5, a RhoA guanine-nucleotide exchange factor, restrains excitatory synapse formation and is degraded by EphB signalling to permit it; reducing its expression ameliorates Alzheimer-like impairment in mice (Margolis et al., 2010; Sell et al., 2017; Hamilton et al., 2017).

What it buys. A mechanism for Precondition C, and — through Ephexin5 — a second, independent connection from the Rho arm of the residue to a synapse-restricting protein whose disposal is itself proteostatic. It also supplies the paper's strongest link to genetic risk: **the ε4 allele's effect on ApoE-receptor recycling is, on this reading, a direct attack on the reelin arm of serine-3 governance.**

Grade: M3 for the ApoE4 recycling result; M2 for the human retromer reduction; the connection to serine 3 specifically is M5 and untested. Demolished if: surface ApoER2 density in ε4 carriers is shown to be preserved at cortical excitatory synapses, or if reelin-induced serine-3 phosphorylation is shown to be insensitive to surface ApoER2 abundance across the physiological range.

10.6 The immune bridge

The claim. The molecules that eliminate synapses in the developing brain are the molecules that eliminate them in the ageing and diseased brain, and the disease is the developmental program running without its developmental constraints.

The evidence. The classical complement cascade mediates central synapse elimination in development (Stevens et al., 2007). C4 localises to human synapses and mediates developmental synapse elimination in mice, with allelic variation in expression conferring psychiatric risk (Sekar et al., 2016). Complement and microglia mediate early synapse loss in Alzheimer mouse models (Hong et al., 2016), C1q-dependent elimination operates in tauopathy and its blockade rescues tau-induced synapse loss (Dejanovic et al., 2018), and tau-oligomer-containing synapses are eliminated by microglia and astrocytes in human Alzheimer disease (Taddei et al., 2023). Class I MHC molecules are activity-regulated in central neurons (Corriveau et al., 1998), regulate developmental refinement (Datwani et al., 2009; Adelson et al., 2016), and co-regulate elimination and learning rules (Lee et al., 2014). Their receptor PirB restricts plasticity across the lifespan (Syken et al., 2006), and blocking it in the adult restores it (Bochner et al., 2014; Djurisic et al., 2019). And C4d, the ligand that rises with age, acts at that same receptor to strip spines, an effect abolished by PirB knockout (Brott et al., 2025).

What it buys. The synthesis of Part XI, and a reframing of the therapeutic target. If the executing machinery is developmental, then the disease is not primarily a matter of a toxin to be removed but of a program to be re-restrained.

Grade: M3 for the developmental biology; M2 for the human Alzheimer complement observations; the identity claim — same program, re-opened — is this paper's Observation and Inference. Demolished if: the complement and MHC-I arms of developmental pruning are shown to act on a molecularly distinct pathway from the one that eliminates synapses in the aged brain.

10.7 The tau bridge

The claim. The residue is not downstream of tau pathology; it is upstream of one of tau's most consequential mislocalisations, and it therefore connects the synaptic and the proteinopathic architectures of this disease at a single molecule.

The evidence. Activated cofilin displaces tau from microtubules, destabilising tau-induced microtubule assembly, missorting tau and promoting tauopathy (Kang and Woo, 2019). Tau mislocalisation to dendritic spines mediates synaptic dysfunction independently of neurodegeneration (Hoover et al., 2010), and dendritic tau mediates amyloid-β toxicity (Ittner et al., 2010). In the other direction, reelin acting through ApoER2 and VLDLR modulates tau phosphorylation (Hiesberger et al., 1999), and the RELN-COLBOS resilience variant reduces tau phosphorylation in a knockin mouse (Lopera et al., 2023). Cofilin–actin rods themselves disrupt microtubules (Minamide et al., 2000). Neuroproteasomes regulate endogenous tau paired-helical-filament formation in an APOE-genotype- and age-dependent manner (Paradise et al., 2026).

What it buys. A single molecule at which the amyloid, matrix, synaptic and tau literatures all have standing. It also explains, without any additional assumption, why a disease whose functional lesion is synaptic should generate a tauopathy: **displacing tau from the microtubule is a cofilin function, and driving cofilin to the active state is what three of the four receptor routes do.**

Grade: M5 for the cofilin–tau displacement mechanism, drawn from review synthesis rather than a single primary demonstration in neurons; M3–M4 for the tau-in-spine results. Demolished if: cofilin activation is shown not to affect tau's microtubule association at concentrations reachable in a spine.

10.8 The neuromodulatory bridge

The claim. Noradrenergic dysregulation reaches the residue by three routes — activity and calcium, Rho tone, and matrix proteolysis — none of which is individually strong.

This bridge is stated in full in Part IX and is not repeated. It is listed here for completeness and to make one point about the architecture of the argument: **it is the weakest of the seven, and the paper's conclusions do not change if it is removed entirely.** In a chain-form account this material would carry the causal weight of the whole story. In a convergence account it is one bridge, graded M5 throughout, with the honest observation that its most specific step — the induction of a matrix protease by noradrenaline — has been demonstrated in peripheral cells and not in brain parenchyma.

Grade: M5 throughout. Demolished if: cortical Rho-kinase activity and matrix protease activity are shown to be independent of noradrenergic state.

10.9 What the bridges jointly establish

Set the seven side by side and the picture that emerges is not a sequence but a **settlement**.

bridge	what crosses it	grade	required by the central claim?
Matrix	the exposure regime, and reelin's obligate co-receptor	M5 / M2	for Precondition D
Receptor	four ligands, four receptors, two directions	M2–M4	**yes — this is the disagreement**
Energy	the cost of turnover; the rod as shutdown	M4	for Precondition B
Proteostasis	receptor return; the ε4 attack on ApoER2	M2–M3	for Precondition C
Immune	the developmental program, re-opened	M2–M3	for Part XI's synthesis
Tau	cofilin displaces tau from microtubules	M5	no — but it is what joins the two architectures
Neuromodulatory	activity, Rho tone, matrix proteolysis	M5	**no — removable without loss**

Three observations about that table.

Only one bridge is load-bearing for the central claim. The receptor bridge is the disagreement of Part IV, and without it there is nothing to reconcile. Everything else supports the *explanation* of the disagreement rather than its existence.

The bridges that support the preconditions are graded lower than the bridges that support the phenomenon. This is as it should be, and it is stated rather than hidden: the phenomenon is better evidenced than the paper's explanation of it. A reader who accepts Part IV and rejects Parts V through IX still has a live problem on their hands, and the field still has to solve it.

The bridge that a chain-form account would have made the spine of the story is the one this paper marks as removable. That is the methodological point of Part I, demonstrated rather than asserted.

PART XI — The Program Re-opened

11.1 The regularity

Assemble the executing molecules of Alzheimer's synaptic failure and ask, of each, what it was doing before the disease. The answer is the same every time, and it is the deepest regularity in the material assembled here.

molecule	job in the developing brain	job in the ageing and diseased brain
Reelin → ApoER2 → Dab1 → pSer3	arrests the migrating neuron by anchoring its leading process (Chai 2009)	stabilises the adult spine; withdrawn as the source cells and the signalling fail (Niu 2008; Cuchillo-Ibáñez 2016)
Complement C4	prunes synapses in postnatal development (Sekar 2016; Stevens 2007)	returns as C4d, binds LilrB2, strips spines (Brott 2025)
MHC class I and PirB / LilrB2	closes the critical period; limits plasticity for life (Corriveau 1998; Syken 2006; Lee 2014)	receptor for amyloid-β oligomers; mediates plasticity loss and spine elimination (Kim 2013)
Perineuronal net (chondroitin sulfate)	closes the critical period; its removal reopens it (Pizzorusso 2002)	degrades in disease; net-bearing neurons resist tau (Morawski 2010; de Vries 2024)
Cofilin at serine 3	severs the actin of the growth cone; executes migration and pruning	severs the actin of the spine; forms the rods of the diseased neurite (Minamide 2000)
Microglia	engulf the synapses complement has tagged	engulf the synapses complement and tau have tagged (Hong 2016; Taddei 2023)
Ephexin5	restrains excitatory synapse formation until EphB degrades it (Margolis 2010)	its reduction ameliorates Alzheimer-like impairment (Sell 2017)

Seven rows, seven independent literatures, one pattern. **Nothing in the right-hand column is new.** There is no molecule in the execution of this disease that the brain did not already possess and use, in youth, for the opposite purpose: to build a circuit by subtracting from an excess.

11.2 What this pattern is not

It is worth stating three readings of the pattern that would be overinterpretations, because each is tempting.

It is not a claim that the disease is a reactivation of a developmental program in the transcriptional sense. No claim is made here about developmental gene expression programs being re-induced. The claim is about molecular identity of the effectors, which is a weaker and better-supported statement.

It is not a claim that development and disease are the same process at different ages. They differ in the most important respect: in development the subtraction is *instructed* — complement tags the synapses that lose the competition, activity selects what is kept, and the whole apparatus is under the control of signals that carry information. In the aged cortex the same effectors act under chronic conditions that carry none.

It is not an argument that pruning is bad. The developmental program is not a latent pathology waiting to be triggered. It is the mechanism by which cortex acquires its structure, and the adult retains it for good reasons — synapse elimination and learning rules are co-regulated by the same molecules (Lee et al., 2014), which means the restraint system is part of how the adult learns.

11.3 What it is

The pattern supports one claim, and it is the paper's synthesis:

> **Alzheimer's synaptic failure is the developmental subtraction program running with its instructions removed.** *The effectors are unchanged. What is lost is the governance — the placed signal that told the machinery which synapse and when, the energy that let a cell follow that instruction and return, the receptor return that made the instruction receivable, and the shielding that kept the effector's ligands episodic rather than ambient.*

This is why the paper is about a residue and not about a toxin, and it is why the therapeutic reading in Part XIV is inverted relative to the field's default.

Consider what the alternative framings imply. If the disease is a toxin, remove the toxin. If the disease is an over-active enzyme, inhibit the enzyme. But if the disease is a **correctly functioning machine running without its instructions**, then both moves are wrong in a specific way: removing one of several ligands leaves the machine running on the others, and inhibiting the effector removes the machine's capacity to do its legitimate job along with its illegitimate one. What the framing implies instead is that the target is the *instruction* — and the instruction, at this residue, has a name, a source cell, a receptor, a matrix requirement and a human resilience allele.

11.4 The resilience allele, read through the synthesis

The RELN-COLBOS variant is worth re-reading in this light, because on the conventional framing it is puzzling and on this one it is not.

A man carrying the PSEN1-E280A mutation, with a very high amyloid burden, remained cognitively intact into his late sixties. He carried a heterozygous gain-of-function variant in *RELN*; the variant activates Disabled-1 more strongly, and a knockin mouse shows reduced tau phosphorylation (Lopera et al., 2023). The variant binds heparan sulfate more tightly than wild-type reelin — 10 ± 2 nanomolar against 17 ± 5 (Pan et al., 2025), which is to say it is better *staged* on the matrix, not merely more potent.

On the conventional framing this is a protective factor of unclear mechanism operating against an overwhelming amyloid burden. On the framing here it is precisely what the synthesis predicts: **the resilient case is the one in which the instruction was strengthened while the toxin was left alone.** The amyloid burden was not reduced. The ligand tide was not stemmed. What differed was the quality of the placed signal at the residue.

That reading is an inference, and the experiment that tests it is stated in Part XIII: the knockin animal should show improved *dynamic range* at serine 3 rather than a tonically raised phospho-cofilin level. If it shows the latter, the reading is wrong and the variant is acting as a stronger tonic input rather than a better signal.

11.5 Why the load framework is required

One might ask whether the synthesis needs Parts II and VI through IX at all. If the disease is the developmental program without instructions, why does it matter which cells carry the most load?

Because the framing alone does not explain *selection*. Every cortical neuron possesses the same effectors; every cortical neuron ages. If the loss of instruction were uniform, the disease would be uniform, and it is conspicuously not. The load framework supplies the selection principle: **the cells that lose governance first are the cells for which governance was always most expensive.**

That is a claim with an ordinal structure, and Part II shows it produces the right ordering at its extreme. It is also a claim that could be wrong in an instructive way. If per-synapse dispersion at serine 3 turns out to be uniform across cortical cell classes and layers — if the heavily loaded neurons are no worse governed than the lightly loaded ones — then the framing survives and the load framework does not, and the interesting question becomes what *else* selects the vulnerable population.

11.6 Why this account produces a tauopathy

An account whose central lesion is synaptic and whose central molecule is an actin regulator owes an explanation of why Alzheimer's disease is, in the neuropathologist's hands, a tauopathy. Most synaptic accounts of this disease either treat tau as a parallel process or reach for it through amyloid. This one has a more direct route, and it runs through the same residue.

Activated cofilin displaces tau from microtubules. The mechanism is set out in the cofilin literature's own synthesis: cofilin acts as a bridge between actin and microtubule dynamics by displacing tau from microtubules, thereby destabilising tau-induced microtubule assembly, missorting tau, and promoting tauopathy (Kang and Woo, 2019). Cofilin–actin rods independently disrupt microtubules in the neurites where they form (Minamide et al., 2000).

Set that beside the four receptor routes of Part IV. Three of them — LilrB2/PirB, β1-integrin/SSH1, and the rod-forming route — drive cofilin toward the active, dephosphorylated state. **Each of those three is therefore, by the mechanism above, a route to tau displacement and missorting**, without requiring any additional assumption about tau kinases.

Three further connections close the loop, and each is independently established.

Tau, once missorted, goes to the spine and does damage there. Tau mislocalisation to dendritic spines mediates synaptic dysfunction independently of neurodegeneration (Hoover et al., 2010), and dendritic tau mediates amyloid-β toxicity in mouse models (Ittner et al., 2010). The compartment the residue governs is the compartment displaced tau enters.

The reelin arm restrains tau directly. Reelin binding to VLDLR and ApoER2 induces Disabled-1 tyrosine phosphorylation and modulates tau phosphorylation (Hiesberger et al., 1999); the resilience-associated *RELN* variant reduces tau phosphorylation in a knockin mouse (Lopera et al., 2023). So the arm that supplies the *stop* at serine 3 is the same arm that brakes tau — one ligand, two consequences, both lost together when the arm fails.

And the disposal machinery couples the two. Neuroproteasomes regulate endogenous tau paired-helical-filament formation in an APOE-genotype- and age-dependent manner (Paradise et al., 2026), and synapses containing tau oligomers are eliminated by microglia and astrocytes in human Alzheimer disease (Taddei et al., 2023) — which returns tau to the same immune-mediated elimination machinery that Part X's immune bridge describes.

What this buys, and what it does not. It buys a single molecular object at which the synaptic and proteinopathic architectures of this disease both have standing, and it buys it without positing a new pathway. It does not buy a theory of tangle formation, of tau strain propagation, or of the topographic staging of tau across the brain, and no such claim is made.

Grade: M5. The cofilin-to-tau-displacement step is drawn from a review synthesis rather than a single primary demonstration in cortical neurons, and it is the weakest link in the chain above. Its falsifier is stated in Part X.7: if cofilin activation does not affect tau's microtubule association at concentrations reachable in a spine, this section is wrong and the two architectures must be joined some other way.

11.7 The image, stated once

A cortex builds itself by overproduction and instructed subtraction. It closes that phase by installing brakes: receptors that restrain plasticity, a matrix that restricts access, a complement system that stands down. It then spends sixty or seventy years holding what it built, paying continuously for the holding, and using the same subtraction machinery in small, instructed doses to learn.

The disease is what happens when the holding becomes unaffordable, the instructions stop arriving, the brakes are engaged by ligands that mean nothing, and the machinery that once sculpted the circuit is left running against a structure nobody is defending.

Not a new process. **An old one, off its leash, in a cell too tired to hold it.**

PART XII — Predictions

A theory that predicts nothing new is a summary. This part states eleven predictions, each of which is a measurement nobody has made, and each of which can come out the wrong way. They are ordered by how much they cost to test.

12.1 The primary prediction

> **P1. In Alzheimer cortex, the synapse-to-synapse dispersion of phospho-serine-3 cofilin is elevated relative to age-matched control, including in regions and at stages where the mean is unchanged.**

This is the prediction on which the paper stands. Three features make it useful.

It is not implied by any of the five accounts in Part IV. Each of them predicts a mean shift and is silent about variance. A result showing elevated dispersion at an unchanged mean is not compatible with any existing reading of the residue.

It fails cleanly. If dispersion is unchanged, or moves only where the mean moves, Proposition 2 is wrong and one of the conventional resolutions in Part IV is right. There is no version of this paper that survives a null result here.

It is measurable now, in accessible human tissue, by a method that has already been applied to these exact synapses. Array tomography resolves individual excitatory synapses in human cortex and quantifies protein content synapse by synapse (Micheva and Smith, 2007; Kay et al., 2013); it was used to place C4d and LilrB2 at human excitatory synapses (Brott et al., 2025). The experiment is a re-analysis of the kind of material that already exists, with an antibody the field already uses.

12.2 Predictions that sharpen the primary one

> **P2. The dispersion is bimodal rather than merely broad**, *with a LilrB2-associated dephosphorylated mode and a Rho-kinase-associated phosphorylated mode.*

Co-staining for LilrB2 assigns each synapse to a mode. If the modes do not separate on receptor content, the receptor explanation in Part V fails even if the dispersion result holds — which would be an interesting outcome, because it would mean the residue is ungoverned for some reason other than differential receptor occupancy.

P3. Dispersion scales with local amyloid proximity rather than with global burden.

Rods are most prominent in neurites contacting amyloid deposits (Minamide et al., 2000), and synapse loss is graded by distance from the oligomeric halo, recovering toward control by roughly 50 micrometres (Koffie et al., 2009). The same geometry should govern the dispersion. A result in which dispersion tracks total plaque load but not local distance would suggest a diffusible or systemic driver instead.

P4. Dispersion is lower on pyramidal neurons that still carry a perineuronal net.

One extra channel — *Wisteria floribunda* agglutinin — in the same imaging run. If net-bearing excitatory neurons show the same dispersion as netless ones, the exposure arm of the load framework should be dropped (Part VIII).

12.3 The ageing prediction

P5. Aged cognitively normal cortex sits between young and demented cortex on the dispersion measure, not with the young.

This follows from Part VI. Spine dynamics in the aged rodent cortex already show a loss of discrimination — new spines over-stabilised, old spines under-retained (Mostany et al., 2013) — and complement rises with normal age (Brott et al., 2025). If the drift toward "setting" begins in health, the dispersion should begin to rise in health.

This is the cheapest way to embarrass the argument. If aged controls are indistinguishable from young controls, the reading of the ageing spine-dynamics data offered in Part VI is wrong.

12.4 Predictions about range rather than level

P6. The excursion of serine-3 phosphorylation in response to a plasticity-inducing stimulus is reduced in Alzheimer tissue and models even where baseline is normal.

This is the most direct statement of the central claim and the hardest measurement in the list, because it requires a live preparation. It is stated because it is the prediction that distinguishes "loss of range" from every account framed in terms of levels.

P7. Conditional deletion of reelin from somatostatin-expressing interneurons in the adult mouse reduces the excursion at pyramidal spines without a large change in baseline phospho-cofilin, and widens the per-synapse distribution.

The writer experiment of Part VII. Note the unusual shape: every conventional account predicts a lowered mean; this one predicts degraded dynamics at a preserved mean.

> **P8. The RELN-COLBOS knockin animal shows improved dynamic range at serine 3 rather than a tonically elevated phospho-cofilin level.**

If the resilience allele works by raising tone, the "signal versus setting" framing has no purchase on the one human genetic anchor the argument possesses (Lopera et al., 2023; Pan et al., 2025).

12.5 Predictions about therapy

> **P9. Rho-kinase inhibition worsens synapses that are already on the dephosphorylated side of the range.**

Fasudil rescues the phosphorylation arm (Rush et al., 2018). On a non-monotonic node, the same drug must push dephosphorylated synapses further from the healthy band. The test is a stratified preparation — tissue or animals sorted by LilrB2 engagement — before further clinical development of the class (Zheng et al., 2025).

> **P10. A heparin mimetic that blocks tau uptake will reduce reelin-induced ApoER2 dimerisation and serine-3 phosphorylation in the same preparation, in an N-sulfation-dependent manner.**

This is the drug conflict of Part VIII, and it is close to a foregone conclusion given that free heparin already blocks the dimerisation while N-desulfated heparin does not (Pan et al., 2025). It is stated as a prediction because nobody has run both readouts in one experiment, and because the magnitude matters: a mimetic that blocks tau uptake at a concentration ten-fold below the one that silences reelin has a therapeutic window, and one that does not, does not.

12.6 A prediction about cells rather than synapses

> **P11. Per-synapse dispersion at serine 3 is ordered across cortical cell classes in the same order as the load ranking of Part II** — *highest on netless deep-layer and entorhinal projection neurons, intermediate on netless superficial pyramidal neurons, lowest on net-bearing cells.*

This is the load framework's own falsifier, and it is the prediction the author expects to be most vulnerable. Load is a plausible organising principle and it may simply not be the one that governs this measurement. If dispersion is flat across cell classes, Proposition 3 fails while Propositions 1 and 2 stand, and the paper reduces to a residue argument with a very long preamble.

12.7 What the predictions have in common

Every prediction above is about a **distribution**, a **dynamic range**, or an **interaction** — never about a mean level.

That is not stylistic. It is the operational content of the claim that the lesion is a loss of governance rather than a displacement. A theory whose distinctive predictions are all about means is a theory about levels; a theory whose distinctive predictions are all about dispersion and excursion is a theory about control. If a future measurement shows that the mean is, after all, the informative quantity at this residue — that it moves reliably, in one direction, across cohorts and preparations — then this paper is unnecessary, the field's existing disagreement was a technical artefact, and the correct response is to say so.

PART XIII — The Experiments, in the Order They Should Be Run

Ten experiments. Each is stated so that a specific result would count against the argument, and they are ordered by what they settle per unit of effort rather than by conceptual tidiness. The first three are, between them, the whole of the paper's empirical exposure on its central claim; the last is the only one that tests the load framework itself.

13.1 Experiment 1 — Per-synapse dispersion in human cortex

Design. Array tomography of Alzheimer and age-matched control cortex on ribbons of serial ultrathin sections, staining phospho-serine-3 cofilin, total cofilin, a synaptic marker, LilrB2, and *Wisteria floribunda* agglutinin. Report the **distribution**, not the mean, and report it separately for net-bearing and netless pyramidal neurons.

Endpoint. For each reconstructed synapse i, the ratio $r(i)$ = integrated phospho-Ser3 signal / integrated total-cofilin signal within the post-synaptic mask. The primary per-case statistic is a robust, scale-invariant dispersion measure of r — the quartile coefficient of dispersion, $(Q3 - Q1)/(Q3 + Q1)$. The unit of analysis is the case, not the synapse.

Why the ratio. It removes synapse-to-synapse variation in cofilin *abundance*, which is large and is not the quantity of interest; and it makes the statistic robust to any process that scales the phospho signal by a common factor across a block.

Refutes if: dispersion is unchanged where the mean is unchanged.

13.2 Experiment 2 — The post-mortem control that Experiment 1 requires

This is not optional and it is stated second rather than in a limitations paragraph, because it governs whether Experiment 1 can be interpreted at all.

The problem. Regulatory serine phosphorylation collapses after circulatory arrest. Biopsy-derived adult human tau is phosphorylated at most of the sites considered abnormal in paired helical filaments, and the hypophosphorylation of autopsy-derived adult tau is attributable to rapid post-mortem dephosphorylation (Matsuo et al., 1994). Any claim about the *absolute level* of a phospho-epitope in autopsy cortex is a claim about what survived the agonal and post-mortem interval.

Why the endpoint survives it, partly. If post-mortem dephosphorylation acts as a common multiplier across a block, a relative dispersion statistic is unchanged while the mean collapses. This inverts the usual objection: the *mean* is the fragile quantity here, and the mean is what the disputing literature has been comparing. Dispersion is computed within a single block that shared one brain, one agonal course, one fixation, one section and one staining run; and Alzheimer and control ribbons can be mounted on the same coverslip and stained in one run, so batch is not confounded with diagnosis (Kay et al., 2013).

What is not answered. If post-mortem dephosphorylation proceeds at *different rates at different synapses* — and local phosphatase content is exactly the sort of thing that varies between synapses — the interval could **manufacture** dispersion rather than preserve it.

The control. A post-mortem ladder in mouse: cortex fixed at 0 (transcardial perfusion), 30 minutes, 2 hours, 6 hours and 12 hours after death, processed identically, dispersion measured at each point. If dispersion is flat across the ladder, the human arm is interpretable. If it rises steeply, the human arm must be read only against tightly interval-matched controls, or abandoned.

Refutes the feasibility of Experiment 1 if: dispersion rises steeply with post-mortem interval in the ladder.

13.3 Experiment 3 — Antibody qualification, and the measurement floor

Design. Commercial phospho-antibodies are validated for Western blot and cryosection immunofluorescence, neither of which predicts performance after aldehyde fixation and resin embedding. Run candidate phospho-Ser3 clones and a total-cofilin antibody through a heterologous cell-based screen that simulates the fixation and embedding chemistry of array tomography (Micheva et al., 2023).

The definitive specificity control, which array tomography makes trivial. Take consecutive ultrathin sections from one block. Treat one ribbon with lambda phosphatase, leave the adjacent ribbon untreated, mount both on the same coverslip, stain and image in one run. Consecutive ultrathin sections are, to a very good approximation, the same tissue. A phospho-specific signal must be abolished on the treated ribbon while total cofilin is unaffected. No other human-tissue method offers a negative control this well matched.

Refutes the programme if: no antibody survives qualification, in which case the whole approach must wait for a better reagent and this should be said publicly rather than worked around.

13.4 Experiment 4 — Co-manipulation at one residue in one system

Manipulate ApoER2/Dab1 signalling and LilrB2 signalling in the same neurons and measure serine-3 phosphorylation and its dynamics. The two pathways have never been run in the same preparation.

Refutes the receptor bridge if: they do not interact at the residue, in which case the convergence proposed here is coincidental and the two arms should be developed separately.

13.5 Experiment 5 — The writer deletion

Conditionally delete *Reln* from somatostatin-expressing interneurons in the adult mouse. Measure, at pyramidal spines: baseline phospho-Ser3, the excursion following a plasticity-inducing stimulus, and the per-synapse distribution.

The prediction is unusual and that is the point: reduced excursion and widened distribution at a *preserved* baseline. Every conventional account predicts a lowered baseline.

Refutes the keystone inference of Part VII if: dynamics are intact, or if the only effect is a lowered mean.

13.6 Experiment 6 — Separate the two sugars

Treat with chondroitinase, which removes the perineuronal net and leaves heparan sulfate; separately with heparinase, which does the reverse. Measure reelin-induced serine-3 phosphorylation after each.

Predicts that reelin signalling survives chondroitinase and fails under heparinase — that is, that reelin does not require a perineuronal net at all.

Refutes the matrix bridge if: chondroitinase abolishes the signal, in which case the net is doing something the sugar chemistry does not predict and the separation of shielding from staging is wrong.

13.7 Experiment 7 — The stratified Rho-kinase test

Stratify tissue or animals by LilrB2 engagement, then apply fasudil or a selective ROCK2 inhibitor and measure per-synapse phospho-Ser3 distribution and synaptic function in each stratum.

Predicts benefit in the stratum whose residue sits on the phosphorylated side and harm in the stratum on the dephosphorylated side.

Refutes the therapeutic corollary if: Rho-kinase inhibition is beneficial across the range, in which case the non-monotonic argument has no clinical consequence even if the biochemistry holds.

13.8 Experiment 8 — The heparin-mimetic window

In one preparation, measure (i) tau uptake and (ii) reelin-induced ApoER2 dimerisation and serine-3 phosphorylation, across a dose range of a candidate heparin mimetic, with an N-desulfated control compound.

Predicts that both readouts move together and in an N-sulfation-dependent manner. The therapeutically decisive question is the *separation* of the two dose–response curves.

Refutes the drug-conflict claim if: tau uptake is blocked at concentrations that leave reelin signalling intact, in which case the conflict is real in principle and irrelevant in practice, which is the best available outcome and should be established before it is assumed.

13.9 Experiment 9 — The energy test

In neurons with graded ATP availability, measure the *excursion* of serine-3 phosphorylation in response to a plasticity stimulus, and the threshold for rod formation, as a function of dynamic actin mass.

Predicts that excursion amplitude falls and rod threshold drops as the ratio of maintained actin mass to available ATP rises — that is, that governance is budget-limited.

Refutes Precondition B if: excursion amplitude is independent of energy availability across the physiological range.

13.10 Experiment 10 — The load index

The quantitative test of Proposition 3, and the only experiment in this list that tests the framework rather than the residue.

Design. In one tissue set, compute for each of a pre-specified list of cortical cell classes: (i) a composite load index from the measurable four of the five loads — geometric from reconstruction, electrical from time-averaged activity or calcium load, proteostatic from recycling and degradation transcriptomic signatures, exposure from lectin histochemistry — and (ii) per-synapse dispersion at serine 3, from the same array-tomography pipeline as Experiment 1. Correlate the two across classes.

The omission, stated up front. Metabolic load cannot be measured per identified cell in vivo by any current method, and the index is therefore built from four columns rather than five. That omission should be reported rather than hidden, and it means the experiment tests a weakened version of the framework.

Predicts a positive rank correlation between the load index and dispersion, with the net-bearing classes falling below the line predicted by their other three loads.

Refutes Proposition 3 if: dispersion is flat across cell classes, or is ordered by something other than load — in which case the framework describes nothing measurable at this residue and the interesting question becomes what does select the vulnerable population.

Why it is listed last. It is the most expensive, it depends on Experiment 1 having succeeded, and it is the test the author expects the framework is most likely to fail. It is included because a framework that never states its own quantitative test is not a framework.

13.11 What the sequence costs

Experiments 3 and 2 are prerequisites and are cheap: one antibody-qualification campaign and one mouse post-mortem ladder. Experiment 1 is a re-staining round on material that already exists in more than one laboratory. Those three, together, settle whether the paper's central claim is true, and none of them requires a new cohort.

Experiments 4, 6 and 9 are culture-scale and can run in parallel with anyone's ongoing work. Experiment 5 is a two-year mouse study with standard reagents. Experiments 7 and 8 are the ones with direct consequences for programmes currently in development, and they should be run before those programmes advance, not after.

The honest summary is that **the whole of this paper's empirical exposure sits in three experiments that cost less than one phase-1 trial**, and that the field is currently choosing the direction of a therapy at this residue by which laboratory it happens to believe.

PART XIV — What This Forbids, and What It Implies for Treatment

14.1 What it forbids

A claim that forbids nothing is not worth grading. This one forbids five things, and three of them are being done now.

It forbids bulk phospho-cofilin as an endpoint. Any trial or biomarker programme that reports mean phospho-serine-3 cofilin from homogenate is measuring a quantity this account predicts is uninformative and potentially sign-inverted relative to the synapses that matter. This is a live practice, and it is the cheapest thing on the list to stop.

It forbids monotonic Rho-kinase therapy as a population treatment. If the residue is pinned in different directions at different synapses, a drug that pushes it one way will help one stratum and harm another. The expected clinical signature is a failed trial with a real positive subgroup — the pattern that has repeatedly been misread in this field as a dosing or staging problem. Rho-kinase inhibitors are in active development for this indication, and a review of the class notes, without apparent alarm, that activation **or** inhibition of these kinases changes dendritic and synaptic structures (Zheng et al., 2025). The sign of the intended therapy is unresolved at the residue through which it acts.

It forbids the inference from "cofilin is dysregulated" to "inhibit cofilin." The reflex at a dysregulated node is to block it. At a non-monotonic effector, blockade moves every synapse in one direction along a curve that turns over, and a subset must cross the turnover point. That is not a theoretical worry; it is the most economical explanation available for a class of intervention that produces convincing rescue in one preparation and no population benefit.

It forbids the assumption that the two literatures can be merged by staging alone. A temporal sequence predicts that variance and mean move together as disease advances. This account predicts that they dissociate. That is a discriminable difference and Experiment 1 discriminates it.

It forbids blocking the sulfated matrix wholesale. A heparin mimetic given to stop transcellular tau propagation would, by the identical chemistry, silence the reelin signal that restrains tau phosphorylation (Pan et al., 2025; Holmes et al., 2013; Hiesberger et al., 1999). The matrix must be modulated, not blocked, and the dose separation must be measured rather than

assumed.

14.2 The counterintuitive target

The therapeutic reading of a non-monotonic node is unusual and worth stating carefully, because the obvious move is the wrong one.

The target implied by this account is not the kinase and not the phosphatase. **It is the restoration of signal-driven control** — the capacity of the residue to be moved transiently and locally and then released. Three things follow, and each is a different kind of intervention from what the field is currently building.

Strengthen the instruction rather than remove the ligand. The physiological system that delivers placed, transient serine-3 phosphorylation is reelin through ApoER2, and it has a human genetic anchor of unusual quality: a gain-of-function *RELN* variant in a man who resisted an autosomal-dominant Alzheimer mutation into his late sixties, with a very high amyloid burden and limited entorhinal tangle burden, the variant activating Disabled-1 more strongly and reducing tau phosphorylation in a knockin mouse (Lopera et al., 2023). The resilient case is the one in which the signal was improved and the toxin was not touched. Whether the protection runs through the cofilin residue is not known and is Experiment 8's neighbour in Part XII (prediction P8).

Restore the receptor rather than the ligand. If Precondition C is real, then the reelin arm can fail with normal reelin and normal receptors simply because the receptors are not returned to the surface — which is exactly what the ε4 allele does to ApoE receptors (Chen et al., 2010), and what retromer deficiency does more broadly (Small et al., 2005; Simoes et al., 2021). Retromer stabilisation is already a pharmacological programme. On this reading it has a synaptic rationale it is not currently claiming: it is a way of putting the reelin receptor back on the membrane.

Protect the writer. If Part VII is right, the somatostatin-expressing interneuron is the adult source of the only characterised placed signal at this residue, and its early failure is a governance failure rather than an inhibitory one. That reframes the therapeutic interest of that cell class entirely: the reason to protect it is not only that it inhibits, but that it *writes*. This is the least developed of the three and the most speculative.

14.3 What follows for amyloid-directed therapy

The account has a specific and non-hostile reading of amyloid removal, and it should be stated because the paper's structure invites misreading in both directions.

Amyloid-β oligomers are a major contributor to the ligand tide of Part VI, and they are the trigger for three of the four receptor routes to the residue. Removing them should therefore reduce ligand occupancy at LilrB2 and β1-integrin and reduce the drive on the residue. That is a real

mechanism and it predicts real benefit.

It also predicts the *shape* of that benefit, and the shape is modest for two reasons that follow directly from the framework. **The ligand tide has other contributors** — C4d rises with age and acts at the same receptor and is sufficient on its own to strip spines (Brott et al., 2025) — so removing one ligand leaves the receptor engaged by another. And **removing ligand does not restore the four preconditions**: it does not resupply the placed signal, refund the energy budget, return the receptors to the surface, or rebuild the shielding. On this account, anti-amyloid therapy relieves one term in the equation and leaves the governance failure intact, which is a reasonable description of the magnitude of effect the class has demonstrated.

The constructive implication is combination, not abandonment: **ligand reduction plus governance restoration** is the pairing this framework recommends, and it predicts a super-additive interaction because the two act on different terms.

14.4 The endpoint must change before the drug does

No programme at this node should proceed on a bulk phospho-cofilin readout. The per-synapse distribution is the quantity of interest, the method to measure it exists and has been applied to these exact synapses, and until it is measured the sign of the intervention is being chosen by which laboratory one happens to believe.

This is the single most actionable statement in the paper and it costs an antibody-qualification campaign and a re-staining round.

14.5 A note on what a positive subgroup means

One consequence deserves separate emphasis because it bears on how trials in this area are read.

If a residue is pinned in different directions at different synapses, and if the proportions differ between patients — by genotype, by regional burden, by stage, by the state of the matrix — then a monotonic drug acting at that residue produces, necessarily, **a mixture of benefit and harm within the same trial population**, with the net effect determined by the mixture. The observable signature is a null or marginal primary endpoint with a real, reproducible positive subgroup that resists explanation by dose or stage.

That pattern is usually read as underpowering, as heterogeneity of diagnosis, or as a failure to treat early enough. On this account it can also be the signature of a correct drug at a non-monotonic target, given to an unstratified population. The two readings imply opposite next steps — one says run it earlier and bigger, the other says stratify by the direction of the lesion before running it again — and the way to tell them apart is Experiment 7.

PART XV — Scope: what portion of the disease this claims

The largest single failure mode of theories in this field is claiming everything and forbidding nothing. This part states the boundary in arithmetic rather than in hedging language, and it does so before the refutation conditions rather than after them, so that a reader can see how much is being claimed before seeing how it might be wrong.

15.1 What is claimed

This account applies to the **execution** of synaptic loss — the mechanical step by which an upstream cause becomes a lost spine — in the population of excitatory pyramidal neurons of neocortex and hippocampus.

Synapse loss is the strongest structural correlate of cognitive impairment in Alzheimer's disease, with neocortical synapse density reaching a multivariate correlation of 0.96 against the Dementia Rating Scale while plaque density contributed 26 per cent of that model's strength (Terry et al., 1991). It is present in frontal cortex biopsies and tracks cognitive severity (DeKosky and Scheff, 1990), it is measurable in CA1 at the stage of mild cognitive impairment (Scheff et al., 2007), and it is now detectable in the living brain (Mecca et al., 2020).

To the extent that dementia in Alzheimer's disease is caused by synapse loss on excitatory cortical neurons, this account claims the final common step. That is a large claim about a well-supported fraction of the disease.

15.2 What is not claimed, and these are large

Initiation. Nothing here says what starts Alzheimer's disease. The account is compatible with amyloid-first, tau-first, vascular-first, inflammatory-first and metabolic-first origins. That compatibility is a weakness as much as a strength, and it is stated as such: an account consistent with every upstream story constrains none of them.

Selective vulnerability, at the level of cause. Part II offers a framework — load — within which vulnerability can be described and ranked, and Part IX shows the framework produces the right ordering at its extreme. Description is not demonstration. That the earliest-failing cells are the most heavily loaded is a correlation across a small number of cell classes, and the framework's own falsifier is stated as prediction P11.

Neuronal death. The account is about spines, not somata. Neuronal death, the loss of whole compartments, and white-matter injury contribute to dementia by routes this account does not describe. There is a reasonable case that the synaptic lesion precedes and partly explains the somatic one — viable neurons persist in entorhinal layer II and CA1 in Alzheimer disease longer than tangle counts suggest (Hof et al., 2003) — but that case is not made here.

Inhibitory interneuron failure. No claim is made that the parvalbumin interneuron fails by the mechanism described here. It probably does not: it is aspiny or sparsely spiny, and the mechanics of Part III are the mechanics of a spine. Its role in the argument is as an exposure regulator upstream of the residue (Part VIII), and that role is itself graded M5.

Dementia not attributable to Alzheimer pathology. A substantial minority of dementia in community-based series is not explained by the pathologies conventionally measured, and a substantial minority of people meeting neuropathological criteria for Alzheimer's disease are not demented at death (Boyle et al., 2019). Primary age-related tauopathy (Crary et al., 2014) and limbic-predominant age-related TDP-43 encephalopathy (Nelson et al., 2019) account for a further share of late-life cognitive impairment. This account speaks to none of those groups.

15.3 The arithmetic, stated plainly

It is worth putting numbers on the restriction, approximately and with their uncertainty acknowledged, because the alternative is a scope section that reads as modest and functions as unlimited.

Attributable-risk analyses of community-based cohorts assign a minority of the population risk of dementia to Alzheimer pathology alone, with vascular disease, TDP-43 pathology, hippocampal sclerosis and Lewy pathology accounting for substantial further fractions, and with a large residual unexplained by any measured pathology (Boyle et al., 2019). Within the fraction that is Alzheimer's disease, this account claims the synaptic-execution arm, which the correlational evidence suggests is the dominant contributor to cognitive severity but is not the whole of it.

The honest summary is that this is a claim about the last step of one route, that route being the one most of the field's structural evidence points at, with no claim over who enters the route or why. A reader who wants a theory of the disease's cause will not find one here and should not be told they have.

15.4 Where the account is most likely to be extended, and where it should not be

Two extensions are natural and one is not, and it is worth marking them so that the paper's boundary is not eroded by well-meaning generalisation.

Natural: other spiny principal neurons in other degenerative conditions. The residue argument is not Alzheimer-specific. Any condition that produces chronic, ungoverned ligand occupancy at receptors that write to serine 3, in a spiny neuron under high maintenance load, should produce the same signature. If the dispersion measure is elevated in Alzheimer cortex and not in, say, a non-degenerative control condition with equivalent post-mortem handling, that is informative; if it is elevated in every neurodegenerative condition examined, the measure is a marker of neuronal distress rather than of this disease, which is worth knowing and would substantially reduce its diagnostic interest while leaving its mechanistic interest intact.

Natural: development and psychiatric conditions. The immune bridge of Part X connects to a body of work on developmental pruning in which allelic variation in complement expression confers psychiatric risk (Sekar et al., 2016). The framework here predicts that conditions of *excessive* developmental pruning and conditions of *ungoverned* adult pruning share effectors and differ in governance. That is a real prediction and it is outside this paper's scope.

Not natural: the extension to non-spiny cells. The temptation to extend the argument to the locus coeruleus neuron's own failure, or to the cholinergic basal forebrain, should be resisted. Those cells are load-bearing by the definition of Part II and they fail early, but the mechanism described here requires a spine, and the paper does not have one for them. Part IX retains the coerulean neuron as an *extreme case of load* and explicitly does not claim that its own degeneration proceeds through the mechanism described for the pyramidal spine.

PART XVI — What Would Refute This

Written before the argument was assembled and unchanged by it. The list is ordered by which proposition each item removes.

16.1 Removing Proposition 1 — the disagreement

1. **A published paper is found that states and resolves the directional conflict at cofilin serine 3 in Alzheimer's disease.** The paper's motivation then largely disappears, and its remaining content is a lifetime framework in search of a problem.
2. **One of the human measurements is retracted or fails replication.** The conflict then has a conventional resolution and Parts V through XI are unnecessary.

16.2 Removing Proposition 2 — the reconciliation

3. **Per-synapse variance of phospho-serine-3 cofilin is not elevated in Alzheimer cortex where the mean is unchanged.** This is the primary falsifier and it is cheap. There is no version of this paper that survives it.
4. **The disputing measurements are shown to differ by stage alone**, with variance tracking the mean throughout. The conventional resolution then suffices and this proposal is unnecessary.
5. **Cofilin's turnover from severing to bundling is shown not to occur at concentrations reachable in a spine.** The reconciliation then rests on an in-vitro property with no cellular reality.
6. **The two receptor systems do not interact at serine 3** when co-manipulated in one preparation. The convergence is then coincidental and the arms should be treated as separate mechanisms that happen to touch the same protein.

Any one of items 3, 5 or 6 removes the paper's central proposal. Item 4 removes its necessity.

16.3 Removing Proposition 3 — the load argument

7. **Per-synapse dispersion is flat across cortical cell classes and layers**, showing no ordering by load. The load framework is then descriptive of nothing measurable at this residue, and the paper reduces to a residue argument with a very long preamble.

8. **Conditional deletion of reelin from somatostatin interneurons leaves serine-3 dynamics intact at pyramidal spines.** Precondition A has no mechanism and the keystone inference of Part VII fails.

9. **Serine-3 excursion amplitude is shown to be independent of energy availability across the physiological range.** Precondition B has no mechanism.

10. **Surface ApoER2 density is shown to be preserved at cortical excitatory synapses in ε4 carriers**, or reelin-induced serine-3 phosphorylation is shown to be insensitive to surface ApoER2 abundance. Precondition C has no mechanism.

11. **Net-bearing pyramidal neurons show the same per-synapse dispersion as netless ones.** Precondition D has no mechanism, and the exposure column of Part II should be dropped from the load definition.

Items 8 through 11 are individually survivable — the framework can lose one precondition and retain the other three — but the loss of any two should be read as a failure of the framework rather than a repair opportunity. That threshold is stated here, in advance, precisely so that it cannot be renegotiated after a result.

16.4 Removing the synthesis of Part XI

12. **The complement and MHC-I arms of developmental pruning are shown to act through a molecularly distinct pathway from the one that eliminates synapses in the aged brain.** The "program re-opened" reading is then a coincidence of nomenclature rather than of mechanism.

13. **Developmental synapse elimination is shown not to require cofilin-mediated actin severing.** The residue would then be incidental to the developmental program and the parallel would be superficial.

16.5 Results that would weaken without refuting

Three outcomes would not falsify the account but would substantially reduce its interest, and they are listed because a theory should say what a disappointing result looks like as well as what a fatal one looks like.

Dispersion is elevated in every neurodegenerative condition examined. The measure would then be a marker of neuronal distress rather than of this disease. The mechanistic claim survives; the diagnostic and stratification value largely does not.

Rho-kinase inhibition proves beneficial across the range, including in synapses with reduced phospho-cofilin. The biochemistry could still be right and the therapeutic corollary of Part XIV would be wrong, which is the part of the paper with the most immediate consequence.

A heparin mimetic blocks tau uptake at concentrations well below those that silence reelin signalling. The drug conflict of Part VIII would be real in principle and irrelevant in practice — the best available outcome, and one that should be established by measurement rather than assumed by anybody developing such a compound.

16.6 The standard the author holds this to

Two entries in the ledger of Part XVII are this paper's own and neither has been tested. The most useful response to this document is not agreement but Experiment 1, and the second most useful is a demonstration that Experiment 1 cannot be interpreted because of Experiment 2.

If the primary prediction fails, the correct conclusion is that the residue moves as a mean, that the field's disagreement was a technical artefact, and that the therapeutic programme currently being run on an unresolved sign is licensed after all. That would also be worth knowing, and it would be worth knowing at the cost of one antibody-qualification campaign rather than at the cost of a trial.

PART XVII — Ledger of Claims

Every load-bearing claim in this paper, with its source and grade. **M1** human population-scale or multi-cohort; **M2** human single cohort or tissue series; **M3** animal replicated; **M4** animal or culture, single laboratory; **M5** in vitro or inference. **Observation** marks a claim about the content of the literature. **Prediction** marks a measurement not yet made.

17.1 The residue and its biochemistry

claim	source	grade	note
Synapse loss is the strongest structural correlate of cognitive impairment	Terry 1991	**M2**	$r = 0.96$ multivariate; plaques 26% of model strength
Synapse loss is present in biopsy and tracks severity	DeKosky and Scheff 1990	**M2**	frontal cortex biopsies, not autopsy
Synapse loss is present at MCI in CA1	Scheff 2007	**M2**	precedes dementia
Synaptic loss is measurable in the living brain	Mecca 2020	**M2**	SV2A PET, widespread
Cofilin is switched at serine 3; phosphorylation prevents F-actin binding	Chai 2009; Bamburg 2021	**M5**	established biochemistry
Cofilin severs at low cofilin:actin ratio and stabilises at high ratio	Bamburg and Bernstein 2016	**M5**	the non-monotonicity Part V turns on
Calcium → calcineurin → Slingshot-1 → serine-3 dephosphorylation	Wang 2005	**M5**	non-neuronal cells; calcineurin inhibitors and SSH1L knockdown both block
Cofilin–actin rods are present in human AD brain and not normal brain	Minamide 2000; Bamburg and Bernstein 2016	**M2**	most prominent in neurites contacting amyloid
Rod formation requires active (dephosphorylated) cofilin	Minamide 2000	**M4**	overexpression of active cofilin is sufficient
Rods transiently retard loss of mitochondrial potential and ATP	Bernstein 2006	**M4**	~1 hour benefit; actin turnover taken as ~50% of neuronal ATP use

claim	source	grade	note
Actin treadmilling cost is contested across two orders of magnitude	Engl and Attwell 2015	**Review**	<1% of global budget to ~half of neuronal use; argument stated to survive the low end

17.2 The five laboratories

claim	source	grade	note
Reelin/ApoER2/Dab1 phosphorylates serine 3 and stabilises	Chai 2009	**M4**	developmental context; leading process of migrating neurons
Reelin promotes dendritic spine development in hippocampal pyramidal neurons	Niu 2008	**M4**	adult-relevant postsynaptic role
ApoER2 is in the PSD in complex with NMDA receptors; its spliced exon is required for reelin-enhanced LTP and memory	Beffert 2005	**M3/M4**	activity-dependent splicing
Reelin/ApoER2/VLDLR modulates tau phosphorylation	Hiesberger 1999	**M4**	direct receptor binding shown
LilrB2/PirB are Aβ-oligomer receptors at nanomolar affinity	Kim 2013	**M3**	human LilrB2 present in human brain
Aβ→LilrB2 enhances cofilin signalling (dephosphorylation), seen in human AD brain	Kim 2013	**M2**	human observation is of cofilin signalling, not a per-synapse ratio
The dephosphorylation direction reproduces in an independent laboratory	Kawaguchi 2022	**M4**	antagonist rather than knockout
Aβ→β1-integrin→SSH1 activates cofilin; AD brain mitochondria contain increased activated/oxidised cofilin	Woo 2015a	**M4 / M2**	third receptor route, third laboratory
RanBP9 regulates SSH1 levels and mediates cofilin–actin pathology in vivo	Woo 2015b	**M4**	upstream scaffold
Aβ signals to cofilin through both LIMK1 and SSH1	Kang and Woo 2019	**Review**	states the two-armed structure explicitly

claim	source	grade	note
Aβ→ROCK increases phospho-cofilin in human AD cortex; necessary and sufficient for impairment	Rush 2018	**M2**	opposite direction to the three rows above
C4d binds LilrB2/PirB, is elevated in AD, and is sufficient to strip spines	Brott 2025	**M2 / M3**	human colocalisation M2; sufficiency in mouse M3, abolished in PirB-null
ROCK1/2 are elevated in AD and the inhibitor class is in development	Herskowitz 2013; Zheng 2025	**M2 / review**	the review itself states "activation or inhibition" alters synaptic structure

17.3 The lifetime of the reader

claim	source	grade	note
Human prefrontal spine density in childhood exceeds adult by 2–3×	Petanjek 2011	**M2**	layer IIIC and layer V pyramidal neurons
Spine elimination continues beyond adolescence through the third decade	Petanjek 2011	**M2**	newborn to 91 years, large human sample
~96% of adult spines stable over one month; half-life >13 months; ~73% in critical period	Grutzendler 2002	**M4**	mouse, two-photon, layer 5 V1
Ageing impairs cognition without forebrain neuron loss; the lesion is synaptic	Morrison and Baxter 2012; Morrison and Hof 1997	**Review**	the standard correction to "ageing is neuron loss"
In aged cortex, new spines are 2–3× more likely to be stabilised while long-term retention of stable spines falls	Mostany 2013	**M4**	density itself unchanged; spines smaller
Complement C4 and C4d increase with age and more so in AD	Brott 2025	**M2**	the ligand tide's age term
Oligomeric Aβ associates with PSDs; ~60% synapse deficit within the halo, recovering by ~50 µm	Koffie 2009	**M3**	the geometry the dispersion prediction inherits

claim	source	grade	note
Synaptic activity regulates interstitial Aβ; endocytosis required	Cirrito 2005, 2008	**M3**	the activity→ligand term
Tau mislocalises to spines and mediates synaptic dysfunction independently of degeneration	Hoover 2010; Ittner 2010	**M3/M4**	dendritic tau
Tau-oligomer-containing synapses are eliminated by glia in human AD	Taddei 2023	**M2**	array-tomography-scale human observation

17.4 The writer, the shield and the matrix

claim	source	grade	note
Reelin is secreted by NPY/somatostatin interneurons into perineuronal matrix and never by parvalbumin cells; Dab1 is predominantly pyramidal	Pesold 1999	**M4**	rat cortex; writer, bed and reader are three different cells
Somatostatin-like immunoreactivity is reduced in AD cortex	Davies 1980	**M2**	the oldest non-cholinergic finding; never overturned
Somatostatin and NPY in CSF correlate with Aβ and tau in MCI	Duron 2018	**M2**	
Somatostatin regulates brain Aβ42 through proteolytic degradation	Saito 2005	**M4**	independent mechanistic link
Reelin protein rises in AD frontal cortex while Dab1 phosphorylation falls; Aβ traps reelin and blocks ApoER2 processing	Cuchillo-Ibáñez 2016	**M2**	more ligand, less signal
Reelin-expressing pyramidal cells are depleted in entorhinal cortex in hAPP mice and human AD; reelin-expressing interneurons are not	Chin 2007	**M2/M3**	the honest complication of Part VII.6
Entorhinal layer II reelin-immunoreactive neurons selectively accumulate intracellular amyloid early	Kobro-Flatmoen 2016	**M4**	

claim	source	grade	note
PNN organisation coincides with the close of the critical period; chondroitinase reopens it	Pizzorusso 2002	**M3**	
Aggrecan-net neurons are protected against tau in subcortical AD regions	Morawski 2010	**M2**	
Excitatory neurons bearing a PNN carry low phospho-tau in human frontal cortex	de Vries 2024	**M2**	aggrecan falls in both AD and resilient; WFA sugars fall only in the resilient
N-sulfated heparan sulfate is an obligate co-receptor for reelin-induced ApoER2 dimerisation	Pan 2025	**M5**	K_D 17 nM, COLBOS 10 nM; heparinase, NDST1-KO and free heparin block; N-desulfated heparin does not
HSPGs are diffuse-ECM components; the PNN is condensed chondroitin sulfate	Fawcett 2022	**Review**	basis for the claim that reelin needs no net
HSPGs mediate internalisation and propagation of proteopathic tau seeds	Holmes 2013	**M4**	the other use of the same polymer
Interneuron deficit links altered network activity to cognitive dysfunction	Verret 2012; Palop and Mucke 2016	**M4 / review**	the disinhibition term

17.5 Development, proteostasis and genetics

claim	source	grade	note
The classical complement cascade mediates CNS synapse elimination in development	Stevens 2007	**M3**	
C4 localises to human synapses and mediates developmental synapse elimination; allelic C4A expression confers psychiatric risk	Sekar 2016	**M1 / M3**	genetics M1; mouse mechanism M3
Complement and microglia mediate early synapse loss in AD models	Hong 2016	**M3**	
C1q-dependent elimination operates in tauopathy; C1q antibodies rescue	Dejanovic 2018	**M3**	

claim	source	grade	note
MHC-I is activity-regulated in central neurons and limits developmental refinement and plasticity	Corriveau 1998; Datwani 2009; Adelson 2016; Lee 2014	**M3**	
PirB restricts ocular-dominance plasticity at all ages; blocking it in the adult restores plasticity	Syken 2006; Bochner 2014; Djurisic 2019	**M3**	the brake, applied for life
PirB is also a receptor for myelin inhibitors	Atwal 2008	**M4**	a third ligand class at the same receptor
ApoE4 selectively impairs ApoE-receptor recycling, reducing glutamate receptor function and plasticity	Chen 2010	**M3**	the ε4 attack on Precondition C
Retromer components are reduced in AD; the vulnerable region uses a distinct recycling-dedicated core	Small 2005; Simoes 2021	**M2 / M4**	
A neuronal membrane proteasome degrades nascent protein in an activity-dependent way and is required for plasticity	Ramachandran and Margolis 2017; Ramachandran 2018	**M4**	spine-local degradation
Early proteasome downregulation drives proteostasis failure in AD	Jiang 2025	**M2**	
Neuroproteasomes regulate endogenous tau PHF formation, APOE- and age-dependently	Paradise 2026	**M3/M4**	
Ephexin5 restrains excitatory synapse formation; reducing it ameliorates AD-like impairment	Margolis 2010; Sell 2017; Hamilton 2017	**M3/M4**	a second Rho-arm connection
A RELN gain-of-function allele confers resistance to autosomal-dominant AD	Lopera 2023	**M2**	n = 1 index case; stronger Dab1 activation; reduced tau phosphorylation in knockin
Entorhinal layer II vulnerability is molecularly localised to a specific excitatory subpopulation	Leng 2021	**M2**	
LC tau precedes cortical tau and is present in young adults	Braak and Del Tredici 2011; Braak 2011; Ehrenberg 2017	**M2**	the extreme-load consistency check

claim	source	grade	note
A minority of population dementia risk is attributable to AD pathology alone	Boyle 2019	**M1**	the scope constraint of Part XV

17.6 This paper's own claims

These are the entries a reader should attack first. None is a result.

claim	grade	where tested
Five literatures converge on the pyramidal dendritic spine without saying so	**Observation**	verifiable from the five preparations; §3.4
Three laboratories report Aβ removing the phosphate and one reports it adding the phosphate; none addresses the others	**Observation**	verifiable directly from the sources in §4.7
The lesion is loss of range at serine 3, not displacement of the mean	**Inference (M5)**	falsifier P1 / Experiment 1
Governed range has four material preconditions — placed signal, energy, receptor return, exposure baseline	**Inference (M5)**	falsifiers 8–11 in Part XVI
The somatostatin interneuron is the adult writer whose loss ungoverns the reader's residue	**Inference (M5)**	Experiment 5
The five loads are settled at one residue	**Inference (M5)**	prediction P11
Alzheimer's synaptic failure is the developmental subtraction program running without its instructions	**Inference (M5)**	falsifiers 12–13 in Part XVI
Per-synapse variance of pSer3 is elevated in AD where the mean is not	**Prediction**	never measured; method exists
Aged normal cortex sits between young and demented on dispersion	**Prediction**	never measured
Net-bearing pyramidal neurons retain dynamic range	**Prediction**	one extra WFA channel
Deleting reelin from somatostatin cells degrades excursion at a preserved mean	**Prediction**	never performed
A heparin mimetic blocking tau uptake will silence reelin signalling by the same chemistry	**Prediction**	dose separation never measured

Twelve entries, none of them a result, and every one of them attached to a stated measurement. That is the whole of what this paper adds, and it is offered in that spirit.

APPENDIX A — The Case Against This Paper

An argument should be able to state its own strongest opposition better than its opponents will. What follows is the case against, written as forcefully as the author can manage, with a reply where one exists and an admission where one does not.

A.1 "The convergence is trivial, because everything converges on the cytoskeleton"

The objection. Cofilin sits downstream of most receptors that change spine shape. Finding that five receptor systems write to it is like finding that five roads lead to the sea. The convergence is a property of the cell's architecture, not a discovery about the disease, and dressing it up as a claim about Alzheimer's is a category error.

The reply. The objection is correct about the convergence and wrong about what is being claimed. Part I.2 concedes the point explicitly: convergence at a residue is compatible with the residue being entirely passive. What raises it above passivity is not that many arrows point at it but that the node is **non-monotonic**, so that departures in either direction are pathological, and that consequence does not follow from the anatomy of signalling. It follows from the biochemistry of one protein.

What survives the objection intact: the disagreement in Part IV, which is a fact about the literature and not about cell architecture.

A.2 "The variance prediction is unfalsifiable in practice"

The objection. Immunofluorescence intensity distributions are heavy-tailed and are affected by fixation, antibody batch, illumination, section thickness, embedding and post-mortem interval. A dispersion measure will differ between any two blocks for reasons that have nothing to do with biology. When it differs, the paper will claim vindication; when it does not, the paper will blame technique. That is not a falsifiable prediction; it is an unfalsifiable one wearing a statistic.

The reply. This is the most serious objection in the list and it is met by design rather than by argument. The endpoint is a **ratio** of phospho to total signal in the same mask, which cancels any factor that scales both channels; the statistic is **scale-invariant**, which cancels any factor that scales the ratio uniformly across a block; the comparison is **within-case**, so between-case handling differences shift each case's scale factor without altering its internal spread; and diagnosis is **not**

confounded with batch, because Alzheimer and control ribbons can be stained in one run on one coverslip (Kay et al., 2013).

What is not met. If post-mortem dephosphorylation proceeds at different rates at different synapses, the interval could manufacture dispersion. Nothing in the human arm excludes this from inside itself, which is why Experiment 2 — the mouse post-mortem ladder — is stated as a prerequisite rather than a control, and why a steeply rising ladder would render the human arm uninterpretable and should be published as such.

A.3 "Loss of range is unfalsifiable in principle"

The objection. Any result can be absorbed. Mean up? Pinned in the phosphorylated direction. Mean down? Pinned the other way. Mean unchanged? Loss of range. The framework has an explanation for every outcome, which means it forbids none.

The reply. The framework does not predict "any mean." It predicts a specific and unusual *conjunction*: **elevated dispersion with a preserved mean**, with modes that separate on receptor content, in tissue where the ligand distribution is spatially graded. That conjunction is not implied by any of the five accounts in Part IV and is straightforwardly absent in a large class of possible results — for instance, a shifted mean with unchanged dispersion, which would refute it, or unimodal broadening that does not sort on LilrB2, which would refute the mechanism while sparing the phenomenon.

The residual concession. The objection is right that a *sufficiently* determined proponent could rescue the framework from most results. The defence against that is not logical but procedural: Part XVI states the refutation conditions in advance, including the threshold at which the loss of two preconditions should be read as failure rather than as an opportunity for repair.

A.4 "The load framework is a metaphor with a table"

The objection. Five "loads," scored ordinally by the author, on cells the author selected, producing a ranking the author already knew. No load has been measured. No index has been computed. The framework does no work that "these big, busy, unprotected cells fail first" does not already do in plain English.

The reply, in two parts. *Partly conceded.* Section 2.7 concedes the framework has not been quantitatively validated and Section 2.8 states what validation would require, including the one measurement — per-cell metabolic reserve in vivo — that no current method supplies.

Partly resisted. The framework does one thing plain English does not: it separates exposure from the other four and thereby predicts that a maximally busy cell can be spared if it is shielded. That prediction is not a restatement; it picks out the parvalbumin interneuron and the net-bearing

pyramidal minority, and it is refutable by one lectin channel (falsifier 11).

A.5 "The reelin arm is a developmental result being asked to carry an adult claim"

The objection. Chai and colleagues measured cofilin phosphorylation in the leading processes of *migrating* neurons. The paper's central mechanism — placed, transient serine-3 phosphorylation at adult spines — has never been demonstrated at an adult spine. Everything in Parts V and VII stands on an extrapolation across four decades of developmental time and two entirely different cellular contexts.

The reply. Conceded in substance and answered in part. The adult synaptic role of the reelin/ApoER2 axis is independently established — spine development (Niu et al., 2008), post-synaptic-density localisation in complex with NMDA receptors, activity-dependent splicing required for reelin-enhanced potentiation and normal memory (Beffert et al., 2005), and modulation of tau phosphorylation (Hiesberger et al., 1999). What has *not* been demonstrated is the specific conjunction: reelin producing a placed, transient serine-3 excursion at an adult spine.

Status. This is the single largest unmeasured step in the paper, and it is not disguised: the ledger grades the reelin-to-serine-3 claim M4 in a developmental context, and Experiment 5 is designed to supply the adult measurement. A reader who thinks the extrapolation illegitimate is entitled to discount Parts V.2, VII and XI accordingly, while retaining Parts III, IV and XII.

A.6 "The paper wants the somatostatin cell to matter and the evidence is a peptide assay from 1980"

The objection. The keystone inference rests on reduced somatostatin-like immunoreactivity — a measure of *peptide content*, not of cell number, and certainly not of reelin secretion. Meanwhile the one study that counted reelin-expressing interneurons in the relevant tissue found their number **unchanged** (Chin et al., 2007). The keystone is therefore built on a marker that does not measure the quantity the argument needs, against a direct count that points the other way.

The reply. Part VII.6 states the Chin result rather than burying it, and the reply there is that a count is not a secretion rate and that the two literatures measure different things. But the objection lands harder than that reply fully answers, and it should be said plainly: **nobody has measured interneuronal reelin output in the Alzheimer cortex.** Until somebody does, the keystone inference is an inference about a quantity that has never been observed. That is why it is graded M5 and why Experiment 5 is a deletion experiment rather than an observational one — the causal test is available even though the observational one is missing.

A.7 "The therapeutic implications are unusable"

The objection. "Restore governed control" is not a drug target. One cannot make a molecule that restores range. The actionable content of the paper reduces to *do not give a Rho-kinase inhibitor to an unstratified population* and *check whether your heparin mimetic silences reelin* — both of which are cautions, not therapies.

The reply. The two cautions are, on their own, worth the paper: one of them concerns a class in active development and the other concerns a chemistry with a verified conflict. But the objection understates the constructive content. Three concrete programmes follow from Part XIV — strengthen the instruction (the *RELN* arm, with a human resilience allele already characterised), restore the receptor (retromer stabilisation, already a pharmacological programme, here given a synaptic rationale it is not currently claiming), and protect the writer. None is easy. All are more specific than "restore control."

A.8 What the author would change if forced to choose one weakness

If the paper had to be reduced to its most vulnerable joint, it would be **A.5** — the extrapolation of the reelin-to-serine-3 mechanism from migrating neurons to adult spines. It is the step on which Propositions 2 and 3 both lean, it has never been measured, and it would be measurable in a competent laboratory within a year. A reader wanting to demolish this paper efficiently should not begin with the philosophy of convergence claims. They should measure whether reelin produces a local, transient serine-3 excursion at an adult pyramidal spine, and if it does not, everything from Part V onward collapses.

APPENDIX B — Powering the primary experiment

The primary prediction is only useful if the experiment that tests it can be sized. This appendix sets out the design parameters that determine whether it can, and flags the two numbers that have to be estimated from pilot data rather than assumed.

B.1 The unit of analysis

The unit is the **case**, not the synapse. This is not a conservative choice; it is the only correct one. Synapses within a case are not independent — they share a brain, an agonal course, a fixation, a block, a section and a staining run — and treating tens of thousands of synapses as independent observations would produce vanishing p-values from three cases per group and would be meaningless.

Each case contributes one number: the quartile coefficient of dispersion of the per-synapse phospho-to-total ratio, computed across all reconstructed synapses in that case. The comparison is that number, between groups.

B.2 What has to come from a pilot

Two quantities determine the sample size and neither can be assumed.

The within-case sampling precision of the dispersion statistic. How many reconstructed synapses per case are needed before the case-level dispersion estimate stabilises? This is answerable directly from existing array-tomography datasets by subsampling: compute the statistic on random subsets of increasing size and find the plateau. It is a computation, not an experiment, and it should be done before any new tissue is stained.

The between-case variance of the statistic within a diagnostic group. This is the denominator of the effect size and it is the number nobody has. It can only come from a pilot of several control cases carried through the full pipeline.

B.3 What can be stated in advance

Direction is predicted, so a one-sided test is defensible — but the two-sided result should be reported regardless, because a *reduction* in dispersion in disease would be a genuinely surprising and important finding that a one-sided design would discard.

The mean must be reported alongside the dispersion, always. The prediction is about a conjunction, and a dispersion result reported without its accompanying mean is uninterpretable with respect to the claim.

Post-mortem interval must be reported per case and modelled, not merely matched. Given Experiment 2's purpose, interval should enter the analysis as a covariate, and the sensitivity of the result to that covariate should be stated.

Region and layer must be pre-specified. Dispersion is predicted to be graded by local ligand exposure (prediction P3); pooling across regions with different plaque burdens will inflate within-case dispersion for reasons the hypothesis itself predicts, which would be a self-fulfilling design. The primary analysis should be within a pre-specified region and cortical layer, with the geometry analysis run separately.

B.4 The pre-registration that this argues for

The design above is unusually vulnerable to analytic flexibility, because dispersion statistics admit many defensible variants and the choice can be made after seeing the data. The appropriate response is to fix, before unblinding: the dispersion statistic and its formula; the synapse-inclusion criteria; the region and layer; the covariate set; and the primary comparison. Everything else is exploratory and should be labelled so.

This is stated here because a paper that proposes a distributional endpoint has an obligation to say how that endpoint could be abused, and because the author's own prediction would be the first beneficiary of the abuse.

APPENDIX C — The residue's regulators, in one place

A short reference for readers coming to cofilin from another direction. Each entry gives the molecule, its action on serine 3, its principal upstream control, and where in this paper it does work.

Cofilin-1 (n-cofilin). The effector. Binds cooperatively along ADP-actin; severs at low occupancy, stabilises and bundles at high occupancy (Bamburg and Bernstein, 2016). Phosphorylation at serine 3 abolishes filament binding. *Parts III, V.*

LIM kinase 1 and 2 (LIMK1/2). Writers. Phosphorylate serine 3. Activated downstream of Rho-family GTPases. *Parts III.2, IV.5, X.3.*

Rho-associated kinase (ROCK1/2). Upstream of LIMK; the canonical route by which raised RhoA activity stabilises spine actin. Elevated in the Alzheimer brain; the target of an inhibitor class in development for this indication (Herskowitz et al., 2013; Zheng et al., 2025). *Parts IV.5, XII.9, XIV.1.*

Slingshot-1 (SSH1/SSH1L). Eraser. Dephosphorylates serine 3. Activated by calcineurin in response to calcium, and released from inhibitory 14-3-3 binding (Wang et al., 2005; Woo et al., 2015a). The route by which electrical load writes to the residue. *Parts II.2, III.2, IV.4, VIII.3.*

Chronophin. A second serine-3 phosphatase; present in the neuronal repertoire and listed among the regulators downstream of amyloid (Kang and Woo, 2019). Not developed here.

Calcineurin. Calcium-dependent phosphatase; dephosphorylates SSH1L and raises its cofilin-phosphatase activity (Wang et al., 2005). The coupling between activity and the residue.

Reelin. Extracellular ligand. Through ApoER2/VLDLR, Disabled-1, Src-family kinases and PI3-kinase, produces serine-3 phosphorylation, placed and transient (Chai et al., 2009). Requires N-sulfated heparan sulfate as an obligate co-receptor (Pan et al., 2025). *Parts IV.1, V.2, VI.1, VII, VIII.5.*

ApoER2 (LRP8). Reelin receptor; component of the post-synaptic density in complex with NMDA receptors; activity-dependently spliced (Beffert et al., 2005); recycling impaired by ApoE4 (Chen et al., 2010). *Parts V.2, X.5.*

LilrB2 / PirB. Immune receptor; ligands include amyloid-β oligomers (Kim et al., 2013), the complement fragment C4d (Brott et al., 2025), and myelin inhibitors (Atwal et al., 2008).

Engagement produces cofilin activation. Restrains plasticity throughout life (Syken et al., 2006). *Parts IV.2, VI.6, X.3, X.6.*

β1-integrin. Receptor for amyloid-β42 oligomers in low- or intermediate-activation conformers; engagement activates SSH1 and cofilin (Woo et al., 2015a). *Part IV.4.*

Cofilin–actin rods. Bundled 1:1 aggregates requiring dephosphorylated and oxidised cofilin; present in Alzheimer brain and not normal brain; transiently energy-sparing, chronically destructive (Minamide et al., 2000; Bernstein et al., 2006; Bamburg and Bernstein, 2016). *Parts III.6, IV.6, X.4.*

References

Adelson JD, Sapp RW, Brott BK, Lee H, Miyamichi K, Luo L, Cheng S, Djurisic M, Shatz CJ. Developmental sculpting of intracortical circuits by MHC class I H2-Db and H2-Kb. *Cerebral Cortex*. 2016;26(4):1453–1463. PMID 25316337.

Atwal JK, Pinkston-Gosse J, Syken J, Stawicki S, Wu Y, Shatz C, Tessier-Lavigne M. PirB is a functional receptor for myelin inhibitors of axonal regeneration. *Science*. 2008;322(5903):967–970. PMID 18988857.

Bamburg JR, Bernstein BW. Actin dynamics and cofilin-actin rods in Alzheimer disease. *Cytoskeleton (Hoboken)*. 2016;73(9):477–497. PMID 26873625. doi:10.1002/cm.21282

Bamburg JR, Minamide LS, Wiggan O, Tahtamouni LH, Kuhn TB. Cofilin and actin dynamics: multiple modes of regulation and their impacts in neuronal development and degeneration. *Cells*. 2021;10(10):2726. PMID 34685706. doi:10.3390/cells10102726

Beffert U, Weeber EJ, Durudas A, Qiu S, Masiulis I, Sweatt JD, Li WP, Adelmann G, Frotscher M, Hammer RE, Herz J. Modulation of synaptic plasticity and memory by Reelin involves differential splicing of the lipoprotein receptor Apoer2. *Neuron*. 2005;47(4):567–579. PMID 16102539. doi:10.1016/j.neuron.2005.07.007

Bernstein BW, Chen H, Boyle JA, Bamburg JR. Formation of actin-ADF/cofilin rods transiently retards decline of mitochondrial potential and ATP in stressed neurons. *American Journal of Physiology. Cell Physiology*. 2006;291(5):C828–C839. PMID 16738008. doi:10.1152/ajpcell.00066.2006

Bochner DN, Sapp RW, Adelson JD, Zhang S, Lee H, Djurisic M, Syken J, Dan Y, Shatz CJ. Blocking PirB up-regulates spines and functional synapses to unlock visual cortical plasticity and facilitate recovery from amblyopia. *Science Translational Medicine*. 2014;6(258):258ra140. PMID 25320232.

Boyle PA, Yu L, Leurgans SE, Wilson RS, Brookmeyer R, Schneider JA, Bennett DA. Attributable risk of Alzheimer's dementia attributed to age-related neuropathologies. *Annals of Neurology*. 2019;85(1):114–124. PMID 30421454. doi:10.1002/ana.25380

Braak H, Del Tredici K. The pathological process underlying Alzheimer's disease in individuals under thirty. *Acta Neuropathologica*. 2011;121(2):171–181. PMID 21170538. doi:10.1007/s00401-010-0789-4

Braak H, Thal DR, Ghebremedhin E, Del Tredici K. Stages of the pathologic process in Alzheimer disease: age categories from 1 to 100 years. *Journal of Neuropathology and Experimental Neurology*. 2011;70(11):960–969. PMID 22002422. doi:10.1097/NEN.0b013e318232a379

Brott BK, Raissi AJ, Micheva KD, Vielmetter J, Mendes MS, Baccus CJ, Huang J, Shatz CJ. C4d, a high-affinity LilrB2 ligand, is elevated in Alzheimer's disease and mediates synapse pruning. *Proceedings of the National Academy of Sciences USA*. 2025;122(38):e2519253122. PMID 40966293. doi:10.1073/pnas.2519253122

Chai X, Förster E, Zhao S, Bock HH, Frotscher M. Reelin stabilizes the actin cytoskeleton of neuronal processes by inducing n-cofilin phosphorylation at serine3. *Journal of Neuroscience*. 2009;29(1):288–299. PMID 19129405. doi:10.1523/JNEUROSCI.2934-08.2009

Chen Y, Durakoglugil MS, Xian X, Herz J. ApoE4 reduces glutamate receptor function and synaptic plasticity by selectively impairing ApoE receptor recycling. *Proceedings of the National Academy of Sciences USA*. 2010;107(26):12011–12016. PMID 20547867.

Chin J, Massaro CM, Palop JJ, Thwin MT, Yu GQ, Bien-Ly N, Bender A, Mucke L. Reelin depletion in the entorhinal cortex of human amyloid precursor protein transgenic mice and humans with Alzheimer's disease. *Journal of Neuroscience*. 2007;27(11):2727–2733. PMID 17360894. doi:10.1523/JNEUROSCI.3758-06.2007

Cirrito JR, Yamada KA, Finn MB, Sloviter RS, Bales KR, May PC, Schoepp DD, Paul SM, Mennerick S, Holtzman DM. Synaptic activity regulates interstitial fluid amyloid-beta levels in vivo. *Neuron*. 2005;48(6):913–922. PMID 16364896.

Cirrito JR, Kang JE, Lee J, Stewart FR, Verges DK, Silverio LM, Bu G, Mennerick S, Holtzman DM. Endocytosis is required for synaptic activity-dependent release of amyloid-beta in vivo. *Neuron*. 2008;58(1):42–51. PMID 18400162.

Corriveau RA, Huh GS, Shatz CJ. Regulation of class I MHC gene expression in the developing and mature CNS by neural activity. *Neuron*. 1998;21(3):505–520. PMID 9768838.

Crary JF, Trojanowski JQ, Schneider JA, Geda YE, Wharton SB, Alafuzoff I, et al. Primary age-related tauopathy (PART): a common pathology associated with human aging. *Acta Neuropathologica*. 2014;128(6):755–766. PMID 25348064. doi:10.1007/s00401-014-1349-0

Cuchillo-Ibáñez I, Mata-Balaguer T, Balmaceda V, Arranz JJ, Nimpf J, Sáez-Valero J. The β-amyloid peptide compromises Reelin signaling in Alzheimer's disease. *Scientific Reports*. 2016;6:31646. PMID 27531658. doi:10.1038/srep31646

Datwani A, McConnell MJ, Kanold PO, Micheva KD, Busse B, Shamloo M, Smith SJ, Shatz CJ. Classical MHCI molecules regulate retinogeniculate refinement and limit ocular dominance plasticity. *Neuron*. 2009;64(4):463–470. PMID 19945389.

Davies P, Katzman R, Terry RD. Reduced somatostatin-like immunoreactivity in cerebral cortex from cases of Alzheimer disease and Alzheimer senile dementia. *Nature*. 1980;288(5788):279–280. PMID 6107862. doi:10.1038/288279a0

Dejanovic B, Huntley MA, De Mazière A, Meilandt WJ, Wu T, Srinivasan K, Jiang Z, Gandham V, Friedman BA, Ngu H, et al. Changes in the synaptic proteome in tauopathy and rescue of tau-induced synapse loss by C1q antibodies. *Neuron*. 2018;100(6):1322–1336.e7. PMID 30392797.

DeKosky ST, Scheff SW. Synapse loss in frontal cortex biopsies in Alzheimer's disease: correlation with cognitive severity. *Annals of Neurology*. 1990;27(5):457–464. PMID 2360787.

de Vries LE, Bahnerth A, Swaab DF, Verhaagen J, Carulli D. Resilience to Alzheimer's disease associates with alterations in perineuronal nets. *Alzheimer's & Dementia*. 2024;21(2):e14504. PMID 39737731. doi:10.1002/alz.14504

Djurisic M, Brott BK, Saw NL, Shamloo M, Shatz CJ. Activity-dependent modulation of hippocampal synaptic plasticity via PirB and endocannabinoids. *Molecular Psychiatry*. 2019;24(8):1206–1219. PMID 29670176.

Duron E, Vidal JS, Grousselle D, Gabelle A, Lehmann S, Pasquier F, Bombois S, Buée L, Allinquant B, Schraen-Maschke S, Baret C, Rigaud AS, Hanon O, Epelbaum J. Somatostatin and neuropeptide Y in cerebrospinal fluid: correlations with amyloid peptides Aβ1–42 and tau proteins in elderly patients with mild cognitive impairment. *Frontiers in Aging Neuroscience*. 2018;10:297. PMID 30327597.

Ehrenberg AJ, Nguy AK, Theofilas P, Dunlop S, Suemoto CK, Di Lorenzo Alho AT, Leite RP, Diehl Rodriguez R, Mejia MB, Rüb U, et al., Grinberg LT. Quantifying the accretion of hyperphosphorylated tau in the locus coeruleus and dorsal raphe nucleus: the pathological building blocks of early Alzheimer's disease. *Neuropathology and Applied Neurobiology*. 2017;43(5):393–408. PMID 28117917. doi:10.1111/nan.12387

Engl E, Attwell D. Non-signalling energy use in the brain. *Journal of Physiology*. 2015;593(16):3417–3429. PMID 25639777. doi:10.1113/jphysiol.2014.282517

Fawcett JW, Fyhn M, Jendelova P, Kwok JCF, Ruzicka J, Sorg BA. The extracellular matrix and perineuronal nets in memory. *Molecular Psychiatry*. 2022;27(8):3192–3203. PMID 35760878. doi:10.1038/s41380-022-01634-3

Grutzendler J, Kasthuri N, Gan WB. Long-term dendritic spine stability in the adult cortex. *Nature*. 2002;420(6917):812–816. PMID 12490949. doi:10.1038/nature01276

Hamilton AM, Lambert JT, Parajuli LK, Vivas O, Park DK, Stein IS, Jahncke JN, Greenberg ME, Margolis SS, Zito K. A dual role for the RhoGEF Ephexin5 in regulation of dendritic spine outgrowth. *Molecular and Cellular Neuroscience*. 2017;80:66–74. PMID 28185854.

Harris JJ, Jolivet R, Attwell D. Synaptic energy use and supply. *Neuron*. 2012;75(5):762–777. PMID 22958818.

Herskowitz JH, Feng Y, Mattheyses AL, Hales CM, Higginbotham LA, Duong DM, Montine TJ, Troncoso JC, Thambisetty M, Seyfried NT, Levey AI, Lah JJ. Pharmacologic inhibition of ROCK2 suppresses amyloid-β production in an Alzheimer's disease mouse model. *Journal of Neuroscience*. 2013;33(49):19086–19098. PMID 24305806.

Hiesberger T, Trommsdorff M, Howell BW, Goffinet A, Mumby MC, Cooper JA, Herz J. Direct binding of Reelin to VLDL receptor and ApoE receptor 2 induces tyrosine phosphorylation of disabled-1 and modulates tau phosphorylation. *Neuron*. 1999;24(2):481–489. PMID 10571241.

Hof PR, Bussière T, Gold G, Kövari E, Giannakopoulos P, Bouras C, Perl DP, Morrison JH. Stereologic evidence for persistence of viable neurons in layer II of the entorhinal cortex and the CA1 field in Alzheimer disease. *Journal of Neuropathology and Experimental Neurology*. 2003;62(1):55–67. PMID 12528818.

Holmes BB, DeVos SL, Kfoury N, Li M, Jacks R, Yanamandra K, Ouidja MO, Brodsky FM, Marasa J, Bagchi DP, et al., Diamond MI. Heparan sulfate proteoglycans mediate internalization and propagation of specific proteopathic seeds. *Proceedings of the National Academy of Sciences USA*. 2013;110(33):E3138–E3147. PMID 23898162.

Hong S, Beja-Glasser VF, Nfonoyim BM, Frouin A, Li S, Ramakrishnan S, Merry KM, Shi Q, Rosenthal A, Barres BA, Lemere CA, Selkoe DJ, Stevens B. Complement and microglia mediate early synapse loss in Alzheimer mouse models. *Science*. 2016;352(6286):712–716. PMID 27033548.

Hoover BR, Reed MN, Su J, Penrod RD, Kotilinek LA, Grant MK, Pitstick R, Carlson GA, Lanier LM, Yuan LL, Ashe KH, Liao D. Tau mislocalization to dendritic spines mediates synaptic dysfunction independently of neurodegeneration. *Neuron*. 2010;68(6):1067–1081. PMID 21172610.

Ittner LM, Ke YD, Delerue F, Bi M, Gladbach A, van Eersel J, Wölfing H, Chieng BC, Christie MJ, Napier IA, Eckert A, Staufenbiel M, Hardeman E, Götz J. Dendritic function of tau mediates amyloid-β toxicity in Alzheimer's disease mouse models. *Cell*. 2010;142(3):387–397. PMID 20655099.

Jiang S, Srikanth M, Serpe R, Yavari S, Gaur P, Collins GA, Soni R, Menon V, Myeku N. Early proteasome downregulation and dysfunction drive proteostasis failure in Alzheimer's disease. *Brain*. 2025;148(12):4372–4388. PMID 40488453.

Kang DE, Woo JA. Cofilin, a master node regulating cytoskeletal pathogenesis in Alzheimer's disease. *Journal of Alzheimer's Disease*. 2019;72(s1):S131–S144. PMID 31594228. doi:10.3233/JAD-190585

Kawaguchi Y, Matsubayashi J, Kawakami Y, Nishida R, Kurihara Y, Takei K. LOTUS suppresses amyloid β-induced dendritic spine elimination through the blockade of amyloid β binding to PirB. *Molecular Medicine*. 2022;28(1):154. PMID 36510132. doi:10.1186/s10020-022-00581-7

Kay KR, Smith C, Wright AK, Serrano-Pozo A, Pooler AM, Koffie R, Bastin ME, Bak TH, Abrahams S, Kopeikina KJ, McGuone D, Frosch MP, Gillingwater TH, Hyman BT, Spires-Jones TL. Studying synapses in human brain with array tomography and electron microscopy. *Nature Protocols*. 2013;8(7):1366–1380. PMID 23787894. doi:10.1038/nprot.2013.078

Kim T, Vidal GS, Djurisic M, William CM, Birnbaum ME, Garcia KC, Hyman BT, Shatz CJ. Human LilrB2 is a β-amyloid receptor and its murine homolog PirB regulates synaptic plasticity in an Alzheimer's model. *Science*. 2013;341(6152):1399–1404. PMID 24052308. doi:10.1126/science.1242077

Kobro-Flatmoen A, Nagelhus A, Witter MP. Reelin-immunoreactive neurons in entorhinal cortex layer II selectively express intracellular amyloid in early Alzheimer's disease. *Neurobiology of Disease*. 2016;93:172–183. PMID 27195475.

Koffie RM, Meyer-Luehmann M, Hashimoto T, Adams KW, Mielke ML, Garcia-Alloza M, Micheva KD, Smith SJ, Kim ML, Lee VM, Hyman BT, Spires-Jones TL. Oligomeric amyloid β associates with postsynaptic densities and correlates with excitatory synapse loss near senile plaques. *Proceedings of the National Academy of Sciences USA*. 2009;106(10):4012–4017. PMID 19228947. doi:10.1073/pnas.0811698106

Lee H, Brott BK, Kirkby LA, Adelson JD, Cheng S, Feller MB, Datwani A, Shatz CJ. Synapse elimination and learning rules co-regulated by MHC class I H2-Db. *Nature*. 2014;509(7499):195–200. PMID 24695230.

Leng K, Li E, Eser R, Piergies A, Sit R, Tan M, Neff N, Li SH, Rodriguez RD, Suemoto CK, et al., Grinberg LT, Kampmann M. Molecular characterization of selectively vulnerable neurons in Alzheimer's disease. *Nature Neuroscience*. 2021;24(2):276–287. PMID 33432193.

Lopera F, Marino C, Chandrahas AS, O'Hare M, Villalba-Moreno ND, Aguillon D, Baena A, Sanchez JS, Vila-Castelar C, Ramirez Gomez L, et al., Quiroz YT, Arboleda-Velasquez JF. Resilience to autosomal dominant Alzheimer's disease in a Reelin-COLBOS heterozygous man. *Nature Medicine*. 2023;29(5):1243–1252. PMID 37188781. doi:10.1038/s41591-023-02318-3

Margolis SS, Salogiannis J, Lipton DM, Mandel-Brehm C, Wills ZP, Mardinly AR, Hu L, Greer PL, Bikoff JB, Ho HY, Soskis MJ, Sahin M, Greenberg ME. EphB-mediated degradation of the RhoA GEF Ephexin5 relieves a developmental brake on excitatory synapse formation. *Cell*. 2010;143(3):442–455. PMID 21029865.

Matsuo ES, Shin RW, Billingsley ML, Van deVoorde A, O'Connor M, Trojanowski JQ, Lee VM. Biopsy-derived adult human brain tau is phosphorylated at many of the same sites as Alzheimer's disease paired helical filament tau. *Neuron*. 1994;13(4):989–1002. PMID 7946342. doi:10.1016/0896-6273(94)90264-x

Mecca AP, Chen MK, O'Dell RS, Naganawa M, Toyonaga T, Godek TA, Harris JE, Bartlett HH, Zhao W, Nabulsi NB, Vander Wyk BC, Varma P, Arnsten AFT, Huang Y, Carson RE, van Dyck CH. In vivo measurement of widespread synaptic loss in Alzheimer's disease with SV2A PET. *Alzheimer's & Dementia*. 2020;16(7):974–982. PMID 32400950.

Micheva KD, Smith SJ. Array tomography: a new tool for imaging the molecular architecture and ultrastructure of neural circuits. *Neuron*. 2007;55(1):25–36. PMID 17610815. doi:10.1016/j.neuron.2007.06.014

Micheva KD, Gong B, Collman F, Weinberg RJ, Smith SJ, Trimmer JS, Murray KD. Developing a toolbox of antibodies validated for array tomography-based imaging of brain synapses. *eNeuro*. 2023;10(12):ENEURO.0290-23.2023. PMID 37945352. doi:10.1523/ENEURO.0290-23.2023

Minamide LS, Striegl AM, Boyle JA, Meberg PJ, Bamburg JR. Neurodegenerative stimuli induce persistent ADF/cofilin-actin rods that disrupt distal neurite function. *Nature Cell Biology*. 2000;2(9):628–636. PMID 10980704. doi:10.1038/35023579

Morawski M, Brückner G, Jäger C, Seeger G, Arendt T. Neurons associated with aggrecan-based perineuronal nets are protected against tau pathology in subcortical regions in Alzheimer's disease. *Neuroscience*. 2010;169(3):1347–1363. PMID 20497908. doi:10.1016/j.neuroscience.2010.05.022

Morrison JH, Baxter MG. The ageing cortical synapse: hallmarks and implications for cognitive decline. *Nature Reviews Neuroscience*. 2012;13(4):240–250. PMID 22395804. doi:10.1038/nrn3200

Morrison JH, Hof PR. Life and death of neurons in the aging brain. *Science*. 1997;278(5337):412–419. PMID 9334292.

Mostany R, Anstey JE, Crump KL, Maco B, Knott G, Portera-Cailliau C. Altered synaptic dynamics during normal brain aging. *Journal of Neuroscience*. 2013;33(9):4094–4104. PMID 23447617. doi:10.1523/JNEUROSCI.4825-12.2013

Nelson PT, Dickson DW, Trojanowski JQ, Jack CR, Boyle PA, Arfanakis K, Rademakers R, Alafuzoff I, Attems J, Brayne C, et al., Schneider JA. Limbic-predominant age-related TDP-43 encephalopathy (LATE): consensus working group report. *Brain*. 2019;142(6):1503–1527. PMID 31039256. doi:10.1093/brain/awz099

Niu S, Yabut O, D'Arcangelo G. The Reelin signaling pathway promotes dendritic spine development in hippocampal neurons. *Journal of Neuroscience*. 2008;28(41):10339–10348. PMID 18842893. doi:10.1523/JNEUROSCI.1917-08.2008

Oh J, Eser RA, Ehrenberg AJ, Morales D, Petersen C, Kudlacek J, Dunlop SR, Theofilas P, Resende EPF, Cosme C, et al., Grinberg LT. Profound degeneration of wake-promoting neurons in Alzheimer's disease. *Alzheimer's & Dementia*. 2019;15(10):1253–1263. PMID 31416793.

Palop JJ, Mucke L. Network abnormalities and interneuron dysfunction in Alzheimer disease. *Nature Reviews Neuroscience*. 2016;17(12):777–792. PMID 27829687.

Pan L, Song X, Su G, Gandy LA, Fang B, Buttaci M, Gibson J, Xia K, Zhang F, Liu J, Wang L, Temple S, Wang C. N-Sulfated heparan sulfate promotes Reelin signaling as a co-receptor. *Journal of the American Chemical Society*. 2025;147(51):46773–46779. PMID 41359976. doi:10.1021/jacs.5c15573

Paradise V, Konrad-Vicario KD, Nguyen C, Sharif NA, Wang X, Mukim RD, Sabu M, Corjuc BT, Bafia J, Fu J, et al., Ramachandran KV. Neuroproteasomes regulate endogenous tau paired helical filament formation in an APOE genotype- and age-dependent manner. *Nature Neuroscience*. 2026. PMID 42215643. doi:10.1038/s41593-026-02297-x

Pesold C, Liu WS, Guidotti A, Costa E, Caruncho HJ. Cortical bitufted, horizontal, and Martinotti cells preferentially express and secrete reelin into perineuronal nets, nonsynaptically modulating gene expression. *Proceedings of the National Academy of Sciences USA*. 1999;96(6):3217–3222. PMID 10077664. doi:10.1073/pnas.96.6.3217

Petanjek Z, Judaš M, Šimic G, Rasin MR, Uylings HBM, Rakic P, Kostovic I. Extraordinary neoteny of synaptic spines in the human prefrontal cortex. *Proceedings of the National Academy of Sciences USA*. 2011;108(32):13281–13286. PMID 21788513. doi:10.1073/pnas.1105108108

Pizzorusso T, Medini P, Berardi N, Chierzi S, Fawcett JW, Maffei L. Reactivation of ocular dominance plasticity in the adult visual cortex. *Science*. 2002;298(5596):1248–1251. PMID 12424383. doi:10.1126/science.1072699

Ramachandran KV, Margolis SS. A mammalian nervous-system-specific plasma membrane proteasome complex that modulates neuronal function. *Nature Structural & Molecular Biology*. 2017;24(4):419–430. PMID 28287632.

Ramachandran KV, Fu JM, Schaffer TB, Na CH, Delannoy M, Margolis SS. Activity-dependent degradation of the nascentome by the neuronal membrane proteasome. *Molecular Cell*. 2018;71(1):169–177.e6. PMID 29979964.

Rush T, Martinez-Hernandez J, Dollmeyer M, Frandemiche ML, Borel E, Boisseau S, Jacquier-Sarlin M, Buisson A. Synaptotoxicity in Alzheimer's disease involved a dysregulation of actin cytoskeleton dynamics through cofilin 1 phosphorylation. *Journal of Neuroscience*. 2018;38(48):10349–10361. PMID 30341179. doi:10.1523/JNEUROSCI.1409-18.2018

Saito T, Iwata N, Tsubuki S, Takaki Y, Takano J, Huang SM, Suemoto T, Higuchi M, Saido TC. Somatostatin regulates brain amyloid β peptide Aβ42 through modulation of proteolytic degradation. *Nature Medicine*. 2005;11(4):434–439. PMID 15778722.

Scheff SW, Price DA, Schmitt FA, DeKosky ST, Mufson EJ. Synaptic alterations in CA1 in mild Alzheimer disease and mild cognitive impairment. *Neurology*. 2007;68(18):1501–1508. PMID 17470753.

Sekar A, Bialas AR, de Rivera H, Davis A, Hammond TR, Kamitaki N, Tooley K, Presumey J, Baum M, Van Doren V, Genovese G, Rose SA, Handsaker RE, Daly MJ, Carroll MC, Stevens B, McCarroll SA. Schizophrenia risk from complex variation of complement component 4. *Nature*. 2016;530(7589):177–183. PMID 26814963. doi:10.1038/nature16549

Sell GL, Schaffer TB, Margolis SS. Reducing expression of synapse-restricting protein Ephexin5 ameliorates Alzheimer's-like impairment in mice. *Journal of Clinical Investigation*. 2017;127(5):1963–1968. PMID 28346227.

Simoes S, Guo J, Buitrago L, Qureshi YH, Feng X, Kothiya M, Cortes E, Patel V, Kannan S, Kim YH, Chang KT, Hussaini SA, Moreno H, Di Paolo G, Andersen OM, Small SA. Alzheimer's vulnerable brain region relies on a distinct retromer core dedicated to endosomal recycling. *Cell Reports*. 2021;37(13):110182. PMID 34965419.

Small SA, Kent K, Pierce A, Leung C, Kang MS, Okada H, Honig L, Vonsattel JP, Kim TW. Model-guided microarray implicates the retromer complex in Alzheimer's disease. *Annals of Neurology*. 2005;58(6):909–919. PMID 16315276.

Stevens B, Allen NJ, Vazquez LE, Howell GR, Christopherson KS, Nouri N, Micheva KD, Mehalow AK, Huberman AD, Stafford B, Sher A, Litke AM, Lambris JD, Smith SJ, John SW, Barres BA. The classical complement cascade mediates CNS synapse elimination. *Cell*. 2007;131(6):1164–1178. PMID 18083105.

Syken J, GrandPre T, Kanold PO, Shatz CJ. PirB restricts ocular-dominance plasticity in visual cortex. *Science*. 2006;313(5794):1795–1800. PMID 16917027. doi:10.1126/science.1128232

Taddei RN, Perbet R, Mate de Gerando A, Wiedmer AE, Sanchez-Mico M, Connors Stewart T, Gaona A, Melloni A, Amaral AC, Duff K, Frosch MP, Gómez-Isla T. Tau oligomer-containing synapse elimination by microglia and astrocytes in Alzheimer disease. *JAMA Neurology*. 2023;80(11):1209–1221. PMID 37812432.

Terry RD, Masliah E, Salmon DP, Butters N, DeTeresa R, Hill R, Hansen LA, Katzman R. Physical basis of cognitive alterations in Alzheimer's disease: synapse loss is the major correlate of cognitive impairment. *Annals of Neurology*. 1991;30(4):572–580. PMID 1789684. doi:10.1002/ana.410300410

Verret L, Mann EO, Hang GB, Barth AM, Cobos I, Ho K, Devidze N, Masliah E, Kreitzer AC, Mody I, Mucke L, Palop JJ. Inhibitory interneuron deficit links altered network activity and cognitive dysfunction in Alzheimer model. *Cell*. 2012;149(3):708–721. PMID 22541439.

Wang Y, Shibasaki F, Mizuno K. Calcium signal-induced cofilin dephosphorylation is mediated by Slingshot via calcineurin. *Journal of Biological Chemistry*. 2005;280(13):12683–12689. PMID 15671020. doi:10.1074/jbc.M411494200

Woo JA, Zhao X, Khan H, Penn C, Wang X, Joly-Amado A, Weeber E, Morgan D, Kang DE. Slingshot-Cofilin activation mediates mitochondrial and synaptic dysfunction via Aβ ligation to β1-integrin conformers. *Cell Death and Differentiation*. 2015a;22(6):921–934. PMID 25698445. doi:10.1038/cdd.2015.5

Woo JA, Boggess T, Uhlar C, Wang X, Khan H, Cappos G, Joly-Amado A, De Narvaez E, Majid S, Minamide LS, Bamburg JR, Morgan D, Weeber E, Kang DE. RanBP9 at the intersection between cofilin and Aβ pathologies: rescue of neurodegenerative changes by RanBP9 reduction. *Cell Death and Disease*. 2015b;6(3):e1676. PMID 25741591. doi:10.1038/cddis.2015.37

Zheng C, Xia W, Zhang J. Rock inhibitors in Alzheimer's disease. *Frontiers in Aging*. 2025;6:1547883. PMID 40182055. doi:10.3389/fragi.2025.1547883

All references were checked against the source record at the time of writing: title, authorship, year, journal and the direction of the reported effect were each verified rather than recalled. The directional disagreement documented in Part IV is the reason that check was necessary, and it is the reason no claim in this paper about the sign of a phosphorylation is made from memory.