The Honest Cascade

Todd Golde's amyloid-first account of Alzheimer's disease, read for what it adds rather than for what it claims: the organ-failure reframe, the vacancy between trigger and degeneration, and the boundary conditions any rival theory has to clear

The lifelong journey · A673T · Brain organ failure · The four models · The dystrophic neurite · The absent brainstem

Benjamin Aaron Gustafsson

AdultCognitiveDisease.com

August 2026

Eight pages arguing that amyloid initiates Alzheimer's disease, and then a paragraph conceding that nobody knows how it produces the disease, that this is the question that matters most, and that there are four ways it might work. This paper is about what that concession is worth — and about the two hundred and forty-three references that never reach the brainstem.

Contents

A note on sources

Part I The Account

1. The Claim, in Its Author's Terms
2. Nosology First
3. The Lifelong Journey
4. What "Brain Organ Failure" Actually Claims
5. The Genetics as the Load-Bearing Wall

Part II Six Years of Record

6. The Genotype That Verged on Causality
7. The Effect Sizes That Arrived on Schedule
8. The Trial That Read Out Negative, and What It Actually Refuted
9. The Nosology, Adopted

Part III The Named Vacancy

10. The Sentence That Matters
11. Four Models, No Winner
12. The Delay Is the Datum
13. The Scaffold Model, and What Its Author Did Next
14. What a Theory Would Have to Supply to Fill the Vacancy

Part IV The Address

15. The One Feature That Distinguishes This Disease
16. What Is Inside a Dystrophic Neurite
17. Seeding in Trans
18. The Experiment the Vacancy Demands

Part V Where the Account Is Weak

19. The Absent Brainstem
20. The Coverage Fraction, Stated Outward and Not Inward
21. Reserve, Which Does the Work and Is Not Specified
22. The Extrapolation That Is Not Flagged
23. Unfalsifiable at the Level of the Synthesis

Part VI What It Adds

24. A Boundary Condition, Not a Rival
25. The Division of Labour
26. The Reframe That Survives the Debate
27. What the Analogy Licenses, and What It Does Not

Part VII Ledger, Refutation, and Tests

28. A Graded Ledger
29. Four Errors of Record
30. Five Conditions of Refutation
31. Eight Experiments, in Order
32. Limitations of This Evaluation

33. Conclusion — The Honest Cascade
References

Abstract

The amyloid cascade hypothesis is usually evaluated as a verdict — right or wrong, alive or dead — and the evaluation usually turns on whether removing amyloid-β from a symptomatic brain improves cognition. That is the least interesting question that can be asked of it, and it is not the question the most rigorous version of the hypothesis is trying to answer.

This paper takes one such version. In 2020 Todd Golde wrote a synthesis of Alzheimer's disease as a lifelong trajectory — genes acting from conception, reserve accumulating through childhood, amyloid depositing for twenty to thirty silent years, and a symptomatic phase that he declined to call early disease and instead called brain organ failure. The argument was submitted to a prize for a new theory of Alzheimer's disease and published essentially intact two years later. It is amyloid-first, and it is written by someone who has spent three decades working on amyloid. It is also the most self-critical statement of that position anyone has produced, and it contains, stated plainly and in its author's own words, a description of the hole in the middle of his own theory.

This evaluation asks what the account adds. Four things, in order of how much they survive the amyloid debate itself.

First, a reframing: symptomatic Alzheimer's disease is not early disease but late organ failure. That claim is independent of what starts the disease. It holds if amyloid is the trigger and it holds if amyloid is a passenger, and it has consequences for trial design, for the meaning of a diagnosis, and for what the word *early* is allowed to mean.

Second, a precisely specified vacancy. The account states that amyloid is necessary to trigger the disease and is not sufficient to produce it, names the missing link — how amyloid accumulation drives the tau, glial and neurodegenerative phase — sets out four non-exclusive candidate mechanisms, and declines to choose. A theory that specifies the dimensions of its own hole is a theory another investigator can complete. Most theories do not do this.

Third, a boundary condition. Deterministic mutations in *APP*, *PSEN1* and *PSEN2*; the protective A673T variant; the dose-dependent bidirectionality of *APOE*; the near-inevitability of Alzheimer pathology in trisomy 21 and in small chromosome-21 duplication kindreds. These are the strongest human causal data in the field, and they are constraints on every theory of this disease, not only on amyloid-first ones. Read as a constraint rather than as a claim, the account is more useful than its author advertises.

Fourth, an address. The dystrophic neurite — the swollen, vesicle-filled axonal structure surrounding a plaque, which distinguishes Alzheimer's disease from nearly every other neurodegenerative disorder — is named as the underexamined lesion and as the plausible site at which amyloid and tau actually meet.

The account is also weak in four specific places, and this paper states them at their full weight. It never bounds its own explanatory fraction, though it bounds the field's data carefully. It leaves cognitive reserve — a load-bearing element that determines when symptoms appear — explicitly unmechanised. It makes an unflagged inference from *lifelong low amyloid protects* to *late amyloid removal prevents*, which the genetics does not license and on which an entire prevention programme now rests. And in two hundred and forty-three references it does not cite the brainstem: the pretangle tau that appears in the locus coeruleus of people in their twenties and thirties, decades before the first cortical plaque, is the strongest available challenge to its ordering, and it is not engaged.

Six years of subsequent record are used to grade the account rather than to argue with it. Two of its predictions have been confirmed: that *APOE4* homozygosity verges on a deterministic genotype, formalised in 2024, and that removing the trigger after neurodegeneration is established would produce small effects, confirmed twice in 2023. One test of the prevention paradigm has read out negative, and that result is examined carefully, because it refutes a molecule rather than a paradigm and the difference matters.

The paper closes with a graded ledger of twenty-four propositions, four errors of record, five conditions of refutation, and eight experiments in the order in which they should be done.

A note on sources

The primary text evaluated here is the ten-page synthesis Golde submitted to the Oskar Fischer Prize in 2020, together with its four figures and a bibliography of two hundred and forty-three items. That argument was published, substantially unchanged, as *Alzheimer's disease — the journey of a healthy brain into organ failure* (Golde, 2022), and the two documents are treated throughout as one statement of a position. Quotations are from that statement. Where the argument is graded against subsequent evidence, the evidence is cited to its primary source and was checked against the indexed record rather than recalled.

The author's laboratory work through 2026 is treated as a separate object and is used in Part II and Part III, where it bears on whether the programme has moved in the direction its own synthesis pointed.

Part I — The Account

1. The Claim, in Its Author's Terms

It is a discourtesy to evaluate a position in a weakened form, so we begin by stating this one at full strength.

The proposition is that Alzheimer's disease is a complex, insidious, degenerative proteinopathy triggered by the formation of amyloid-β aggregates; that over decades this pathology drives neurofibrillary tangle formation, dysfunction of virtually every cell type in the brain, and eventually overt neurodegeneration; and that by the time any of this becomes clinically visible, the correct description of the organ is not *early disease* but *failure*.

Four moves distinguish this from the standard statement of the amyloid cascade hypothesis, and all four are worth noticing before the account is examined.

The first is that it begins with **nosology** — with the naming and classification of the disease — rather than with mechanism. The opening pages are about what the word *Alzheimer's* denotes, how often it is applied to brains that do not have the pathology, and how often it is withheld from brains that do. Almost no synthesis in this field opens this way. It is the move that makes everything after it gradeable, and Chapter 2 is devoted to it.

The second is that the unit of analysis is **the lifespan**, not the lesion. The account's central figure is a horizontal line from age zero to age eighty on which genetic factors are active throughout, reserve is built in the first decades, lifestyle and comorbidity modify the trajectory continuously, and the pathological sequence — amyloid deposition, then complex cellular dysfunction, then tauopathy and secondary proteinopathies, then neurodegeneration — occupies only the final third. The disease is presented as a biography rather than a mechanism.

The third is that the linear cascade is **explicitly disowned**. The account builds on Hardy and Higgins and on Hardy and Selkoe (Hardy & Higgins, 1992; Hardy & Selkoe, 2002) and then states that the linear cascade those papers proposed is far more complex than originally thought. The word *trigger* does a great deal of work in this text, and it is chosen in preference to *cause*. Amyloid initiates; it does not execute.

The fourth is that the account **names its own principal gap and does not fill it**. In the author's words, the most critical poorly understood aspect of the hypothesis is how amyloid pathology drives the downstream cellular dysfunction that leads to the neurodegenerative phase, and — in the first person, which is rare in a document of this kind — that this central question needs to be answered in order to complete our understanding of the disease. Four candidate mechanisms are offered. None is endorsed. The figure that presents them ends with the sentence *Models are not mutually exclusive*.

That fourth move is the reason this paper exists. A theory can fail in two ways: by being wrong, and by being unable to say what would make it wrong. This one does something rarer than either. It identifies the joint on which its conclusion depends, states that the joint is not established, specifies the alternatives, and leaves the question open in public. Whatever else that is, it is honest, and it converts a claim into a piece of work someone else can pick up.

2. Nosology First

The account opens with a definition problem, and the problem is not academic.

Alzheimer's disease is defined pathologically by plaques and tangles, and it is diagnosed clinically by a progressive amnestic dementia. Those two definitions do not pick out the same people. The account states the mismatch in both directions, which is unusual — most treatments of this issue state it in one.

In the first direction, dementias that are not Alzheimer's disease are routinely classified as Alzheimer's disease. Clinical criteria cannot reliably separate it from the several other conditions that produce a similar syndrome, and studies of mild cognitive impairment frequently enrol people with no significant neuropathology at all.

In the second direction, the controls are contaminated. A substantial fraction of cognitively normal older people carry Alzheimer pathology at autopsy. They are not disease-free; they are pre-symptomatic, or resilient, and either way their inclusion as controls attenuates every comparison made against them.

The consequence the account draws is a rule about how to read the literature: any study that rests solely on a clinical diagnosis of Alzheimer's disease or mild cognitive impairment is confounded in both directions simultaneously, and the confounding is not random. This is applied consistently and to the author's own disadvantage. When lifestyle and comorbidity are discussed, the account notes that the epidemiology measures *dementia*, not *Alzheimer's disease*, and declines to import the associations wholesale. When traumatic brain injury is discussed, the account observes that the more consistent post-mortem association is with chronic traumatic encephalopathy, which is a tauopathy with a different spatiotemporal distribution. When stress, depression and sleep disruption are discussed, the account raises the possibility that the association runs backwards — that these are early manifestations of a pathology already underway rather than causes of it.

Two things follow.

The first is that the account is unusually well positioned to be graded, because it has told the reader which of its own supporting evidence to discount. A synthesis that says *these associations are with dementia and not necessarily with this disease* has pre-emptively removed a category of evidence from its own column.

The second is that the discipline is applied outward and not inward. The account bounds the field's data with real care and never states what fraction of the disease its own ordering explains. That asymmetry is taken up in Chapter 20, and it is the single most consequential omission in the

document.

It is worth recording that the field subsequently moved the way this chapter argues. The 2024 revision of the diagnostic criteria defines Alzheimer's disease biologically, by biomarker evidence of the pathology, rather than by the clinical syndrome (Jack et al., 2024). That is the nosological position the account took in 2020, and it is now the criteria.

3. The Lifelong Journey

The account's organising figure shows a hypothetical individual who becomes symptomatic at eighty, with the contributing factors drawn as bands beneath the lifeline. It repays close reading, because the claims encoded in the drawing are stronger than the claims made in the text.

Genetic factors run the whole width of the figure. This is not the trivial statement that genotype is fixed at conception. It is the claim that Alzheimer risk genes are doing biological work continuously, from development onward, and that their effect on disease risk is the integral of that work rather than an event late in life. The account supports this with the observation that *APOE* genotype associates with hippocampal volume and myelination in infants and children — and then, creditably, immediately limits it: these early differences do not translate into measurable differences in cognitive ability, and there is no strong correlation between Alzheimer risk genes and early-life development or intelligence. The band is drawn across the whole lifespan and the text declines to claim that anything visible happens in the first half of it.

Reserve is built early and spent late. Educational attainment and intelligence are among the few factors reproducibly associated with altered dementia risk. The account takes the standard division — brain reserve as structural tolerance, cognitive reserve as differences in how tasks are performed — and takes a definite position on the mechanism: current data more strongly support the view that these factors do not alter the presence of pathology but alter the ability to retain function despite it. Reserve moves the symptom threshold, not the lesion.

Lifestyle modifies the reading rather than the pathology. Diet, exercise, blood pressure and metabolic health are associated with modest reductions in later-life dementia risk. The account's reading is that these act largely through vascular and general health — that is, through comorbidity and reserve — rather than directly on Alzheimer pathophysiology, and it notes that model-system results suggesting otherwise are inconsistently replicated with modest effect sizes.

The pathological sequence occupies the last three decades and is strictly ordered. Amyloid deposition, then complex cellular dysfunction, then tauopathy and other secondary proteinopathies, then neurodegeneration, with age-related comorbidities rising in the final years.

The figure therefore encodes a three-layer model. There is a *pathology layer* with a fixed internal order. There is a *tolerance layer* — reserve — that determines the mapping from pathology to symptoms without touching the pathology. And there is a *modifier layer* — lifestyle, vascular health, comorbidity — that acts mostly on the tolerance layer and only weakly on the pathology layer.

This is a more sophisticated structure than the cascade it is built on, and it is falsifiable in a way the cascade is not. If reserve turned out to alter amyloid deposition itself, the middle layer would

collapse into the first. If lifestyle interventions were shown to change the pathology rather than the threshold, the third layer would move. Both are testable. Neither has been settled.

4. What "Brain Organ Failure" Actually Claims

The phrase is the account's most quotable contribution and it is not decorative. Four distinct claims are packed into it.

Claim one: the clinic sees the disease late, and the word *early* is being misused. Mild cognitive impairment is the earliest clinically distinguishable stage, and it is not early. By the time it is diagnosed the pathologies have been present for many years, structural change and neurodegeneration are already established, and amyloid is likely to have been present and plateaued for over a decade. The account's phrasing is that even the earliest symptomatic stages represent long-standing pathology and early brain *organ failure*. A patient at the threshold of diagnosis is at roughly the position of a patient with a first episode of decompensated heart failure: it is the first clinical event, and it is a late biological one.

Claim two: the failing thing is an organ, not a cell type. This is the substantive content of the analogy. Late-stage Alzheimer's disease is presented as involving pathophysiological alteration in every cell type in the brain — neurons, astrocytes, microglia, oligodendrocytes, vasculature — with proteomic and transcriptomic studies showing thousands of changes. The account explicitly criticises the neuron-centric framing of the cascade's own experimental programme as one of its flaws.

Claim three: single-target therapy at this stage is a category error. Heart failure is treated with multiple agents acting on different limbs of a failing system, and the account argues the same must be true here: modifying the course of symptomatic disease will require targeting multiple aspects of the pathophysiology, whether by one pleiotropic drug or by combination.

Claim four: restoration, not just arrest, is required. The account states that drugs will likely need some ability to restore or regenerate functional brain circuits to have transformative impact. This is the most demanding claim in the document and the least discussed. It concedes that even perfect arrest of the pathological process in a symptomatic patient leaves a brain that has lost tissue, and that recovering function from that state is a different and harder problem than stopping the disease.

The four claims have different strengths. The first is close to established, and the biomarker literature the account cites carries it. The second is well supported and now conventional. The third is an inference from the first two rather than a result. The fourth is a conjecture, and it is stated as one.

What the phrase does *not* claim is worth stating too. It is not a claim about aetiology. Nothing in *brain organ failure* requires amyloid to be the trigger. The reframing is portable: any account of this disease that agrees the clinical phase is preceded by decades of accumulating pathology can adopt it without adopting anything else in the document. That portability is the reason Chapter 26 treats it as the account's most durable single contribution.

5. The Genetics as the Load-Bearing Wall

Strip the account of its narrative and what remains is a genetic argument, and the genetic argument is the strongest human causal evidence in this field. It has four legs.

Leg one: the deterministic mutations converge on one molecule by three routes. Rare variants in *APP*, *PSEN1* and *PSEN2* cause autosomal dominant Alzheimer's disease with onset typically between forty and sixty. The account groups their biochemical effects into three classes: increased total amyloid-β production; an increased ratio of the longer, more aggregation-prone species, typically Aβ42 and sometimes Aβ43; and alteration of the peptide sequence itself in ways that promote aggregation. Three mechanistically distinct routes, one convergent consequence — the propensity of the peptide to aggregate. Convergence of this kind is the strongest form of genetic argument available, because it is difficult to explain by anything other than the shared endpoint.

Leg two: the protective variant runs the argument backwards. The A673T substitution in *APP*, identified in an Icelandic population, reduces amyloid-β production and confers lifelong protection against both Alzheimer's disease and age-related cognitive decline (Jonsson et al., 2012). This is the single most important observation in the account, and the reason is structural. Loss-of-function and gain-of-function evidence in the same gene, in humans, in opposite directions, is as close to an experiment as human genetics gets. It is very hard to construct an account in which amyloid-β is a bystander and a variant that lowers it protects against the disease for a lifetime.

Leg three: *APOE* is bidirectional and dose-dependent. The ε4 allele raises risk and the ε2 allele lowers it, both in an allele-dose-dependent way. The account observes that the roughly tenfold risk of the ε4/ε4 genotype *verges on causality*, and that the rarer ε2/ε2 genotype confers almost complete protection — the latter established in a five-thousand-person neuropathological series (Reiman et al., 2020). The mechanism is tied to amyloid: ε4 is associated with earlier and heavier deposition in autopsy and in amyloid ligand imaging, ε2 with delay and reduction. Here the account is careful in a way that is easy to miss. It lists the non-amyloid biology of apolipoprotein E — lipid transport, immune function, cardiovascular health, tau pathology in mouse models — and then states that it remains unclear whether these are major contributors to the isoform-dependent risk. It does not claim the amyloid route is the only route. It claims it is the established one.

Leg four: gene dosage settles the direction of causation. Individuals with trisomy 21 carry three copies of *APP*, produce more amyloid-β, deposit it in the first decade of life, and reach Alzheimer-type plaque and tangle burdens in the fourth and fifth. The account then supplies the control the trisomy cannot: within selected kindreds, duplication or triplication of a much smaller

chromosome-21 region containing *APP* alone produces early-onset familial Alzheimer's disease. The small duplications isolate *APP* from the rest of chromosome 21, and the phenotype survives the isolation.

Set against these four legs, the rest of the genetic material in the account is handled with notable restraint. Common risk loci from genome-wide association are acknowledged to have mostly unknown functional bases, with biological imputation described as broadly pointing to immune function, lipid metabolism, tau-binding proteins and precursor-protein metabolism, and *speculative* for most individual associations (Kunkle et al., 2019; and subsequently Bellenguez et al., 2022, which extended the count to seventy-five loci). The rare coding variants in *TREM2*, *PLCG2* and *ABI3* that implicate microglial innate immunity are cited as suggestive that increased microglial activation through these proteins is protective, with the explicit caveat that this is far from settled science (Sims et al., 2017).

And there is an aside in this section that deserves more attention than it has received. The account observes that a variant with a large biological effect and a small effect on disease risk is unlikely to be a good therapeutic target, and that loci with subtle biological effects may make better targets than loci with prominent ones. That is a genuine insight about the relationship between effect size in biology and effect size in epidemiology, it is stated in a single sentence, and it cuts against the way most target-selection is done.

Part II — Six Years of Record

A synthesis written in 2020 can be graded rather than argued with, because the intervening record contains results the author could not have known. This part uses that record. It is not a defence of the account and it is not a prosecution; it is an audit of four specific commitments against what has since been measured.

6. The Genotype That Verged on Causality

In 2020 the account made a claim that was, at the time, a rhetorical flourish: that given the commonness of the disease, the roughly tenfold risk attributable to a homozygous ε4/ε4 genotype *verges on causality*.

In 2024 that flourish was converted into a formal result. A study combining the National Alzheimer's Coordinating Center series with five biomarker cohorts — 3,297 individuals for the pathological analysis and 10,039 for the clinical analysis — reported that almost all *APOE4* homozygotes show Alzheimer pathology; that their biomarkers diverge from ε3/ε3 homozygotes from age fifty-five; that by sixty-five nearly all have abnormal cerebrospinal amyloid and three-quarters have a positive amyloid scan; and that the age of symptom onset is both earlier and more narrowly predictable than in ε3 homozygotes. The authors concluded that ε4/ε4 constitutes a distinct genetically determined form of the disease, with a predictability of onset and a biomarker sequence that mirror autosomal dominant Alzheimer's disease and Down syndrome (Fortea et al., 2024).

Three things should be said about this.

It is a confirmed prediction, and confirmed predictions are worth more than accommodated observations. The account did not merely file ε4 among risk factors; it said the homozygous genotype was approaching determinism. Four years later a large series said so in the technical vocabulary. Nothing in the 2020 document had to be adjusted to absorb the finding.

It strengthens the strongest leg and not the weakest one. The finding is about ordering and penetrance — about *when* the biology starts and how reliably — and not about the mechanism by which amyloid produces degeneration. It moves Chapter 5's third leg from *very strong* to *nearly as strong as the deterministic mutations*, and it moves the vacancy of Part III not at all.

It sharpens a problem the account did not anticipate. The same study reports that at the dementia stage there were no differences in amyloid or tau positron-emission tomography across haplotypes, despite the earlier clinical and biomarker divergence. Genotype governs the timing of arrival at the pathological state and not the state itself. That is a clean result and it is awkward for any simple dose-response reading of amyloid: two groups reach the same imaging endpoint, on different schedules, with different clinical histories. It is comfortable, however, for the three-layer model of Chapter 3, in which genotype sets the clock on the pathology layer and something else governs the mapping to symptoms.

7. The Effect Sizes That Arrived on Schedule

The account's most exposed therapeutic claim was made before either of the trials that tested it read out. It states that it has proven extremely hard to modify symptomatic disease by targeting amyloid deposition, even with therapeutics showing good target engagement; that this failure is *not surprising* given the triggering role of amyloid aggregate accumulation; and that such data inform us that once substantial neurodegenerative changes are present, targeting the trigger is not sufficient to provide much benefit.

That is a prediction with a shape: anti-amyloid therapy in symptomatic patients should produce an effect that is real, small, and disproportionate to the amount of amyloid removed.

Both trials that subsequently succeeded produced exactly that shape.

In CLARITY-AD, 1,795 participants with mild cognitive impairment or mild dementia received lecanemab or placebo for eighteen months. Amyloid burden fell by 59.1 centiloids relative to placebo — near-complete plaque clearance in most treated participants. The primary clinical endpoint moved by 0.45 points on an eighteen-point scale (1.21 versus 1.66; 95% CI −0.67 to −0.23), with infusion reactions in 26.4% and amyloid-related imaging abnormalities with oedema in 12.6% (van Dyck et al., 2023). Donanemab, in TRAILBLAZER-ALZ 2, produced slowing of the same general magnitude in early symptomatic disease, with the effect largest in the low-to-medium tau stratum (Sims et al., 2023).

A rival reading of the same numbers should be acknowledged. It has been argued that plaque must be reduced below roughly twenty centiloids before meaningful benefit appears, that there is a necessary lag between removal and clinical effect, and that the speed of removal is therefore the critical drug property — on which reading the modest effects reflect incomplete and slow clearance rather than the exhaustion of trigger-dependence (Karran & De Strooper, 2022). That reading is testable by longer exposure and faster-clearing agents, and it predicts that effect sizes should grow with time on treatment. The account's reading predicts that they should plateau. This is one of the few places where two live interpretations of the same trials make divergent quantitative predictions, and the extension data will decide it.

The dissociation is nonetheless the point. Essentially all of the plaque was removed and roughly a quarter of the decline was averted. If plaque were the proximate cause of ongoing degeneration in a symptomatic brain, that ratio would be difficult to explain. If plaque is the trigger of a process that has long since become self-sustaining, the ratio is what one would expect — and the account said so in advance.

It is worth being precise about what this does and does not establish. It does not establish that the amyloid-first ordering is correct; a theory in which amyloid is one of several co-drivers predicts partial benefit too. What it establishes is that this particular account was not embarrassed by the trials. Its author's position in 2020 was that these drugs, given at this stage, would do about this much — and the observation that the field's mainstream framing at the time was considerably more optimistic is a fact about the field, not about the drugs.

The tau stratification in the donanemab result adds something the account did not predict but is well positioned to absorb. Benefit was larger where tau pathology was lower. On the account's own logic — trigger first, executor second — the trigger-directed drug should work best in patients whose executor has not yet been fully recruited. The data behaved as if that were true.

8. The Trial That Read Out Negative, and What It Actually Refuted

The account's therapeutic hope is prevention: intervene against amyloid in amyloid-positive but asymptomatic individuals, and expect this to succeed where treatment of symptomatic disease has not. It names the paradigms — the prevention studies then underway in individuals with deposition before symptom onset, and in those at genetic risk without deposition.

The first of these to report was negative.

The A4 study randomised 1,169 amyloid-positive but cognitively unimpaired individuals aged sixty-five to eighty-five to solanezumab or placebo for 240 weeks. The primary cognitive composite declined by 1.43 in the treated group and 1.13 in the placebo group — a difference of −0.30 in the wrong direction, with a confidence interval spanning zero ($P = 0.26$). Solanezumab did not slow cognitive decline in preclinical Alzheimer's disease (Sperling et al., 2023).

An honest reading has to separate two things this result could mean.

Reading one: the prevention paradigm failed. Amyloid was present, the drug was given for four and a half years before symptoms, and nothing happened. If amyloid deposition is the trigger, removing it before the downstream cascade is established should have worked, and it did not.

Reading two: the molecule failed, and the paradigm was not tested. This reading is supported by a number in the same abstract that is easy to skip. Amyloid burden *increased* in both arms over the trial — by 11.6 centiloids on solanezumab and 19.3 on placebo. The drug targets monomeric amyloid-β. It did not clear plaque; it modestly slowed accumulation. Participants finished the trial with more amyloid than they started with.

The second reading is the correct one, and it is not special pleading, for two reasons. The first is that the account itself named insufficient target engagement as one of three canonical reasons for the preceding two decades of trial failure — the others being toxicity and mistimed administration. Solanezumab is the textbook instance of the failure mode the document identified in advance. The second is that the comparison is now available: in symptomatic patients, drugs that do remove plaque produce a measurable clinical effect, and this drug, which does not, produced none. The variable that differs across those trials is target engagement.

What follows is a statement about the state of the evidence that should be uncomfortable for everyone. **The prevention paradigm has not yet been tested at adequate target engagement.** The trials that could test it — plaque-clearing antibodies given to asymptomatic

amyloid-positive individuals, and to individuals at genetic risk before deposition — are running and have not reported. In dominantly inherited disease, the first such attempt found that neither gantenerumab nor solanezumab slowed cognitive decline in the primary analysis, though gantenerumab reduced amyloid and moved downstream biomarkers (Salloway et al., 2021).

The correct entry in the ledger is therefore *untested*, not *refuted*. That entry will be uncomfortable to defend for much longer. A paradigm that has been the field's principal hope for a decade, whose first completed test failed for reasons that can be explained, and whose decisive tests are still pending, occupies a position that is defensible now and will not be defensible indefinitely.

9. The Nosology, Adopted

The account's opening argument was that the disease should be defined by its pathology rather than by its syndrome, and that the biomarker era makes this possible for the first time in living people.

In 2024 the diagnostic criteria were revised along exactly those lines. Alzheimer's disease is now defined biologically, by biomarker evidence of amyloid and tau pathology, with clinical staging layered on top of a biological diagnosis rather than constituting it (Jack et al., 2024). Someone who is amyloid- and tau-positive and cognitively normal has the disease under these criteria. The 2020 document's position — that a substantial fraction of elderly controls are in a preclinical stage and should not be treated as disease-free — became the definition.

The account should be credited for this and then immediately examined on it, because the revision creates a problem the document did not foresee and does not have the resources to solve.

If the disease is defined by its biology, and if a large fraction of people who carry that biology never become demented, then the biological definition has severed the disease from the thing that makes it matter. The account has an answer available — reserve, which determines the mapping from pathology to symptoms — but the answer is only a name (see Chapter 21). Under the new criteria, the question *why do some people with the disease never get ill from it?* is no longer peripheral. It is the central clinical question, and the document treats it in a page and a half.

There is a second-order consequence for the account's own coverage argument. Defining the disease biologically removes the misclassification problem in one direction — non-Alzheimer dementias will no longer be counted as Alzheimer's disease — and enormously enlarges it in the other, by admitting to the diagnosis a large population of people with pathology and no illness. The account's nosological rigour has been vindicated and, in the same motion, handed a harder problem than the one it solved.

Part III — The Named Vacancy

10. The Sentence That Matters

Buried in the middle of the account, after eight pages of argument that amyloid initiates the disease, is a sentence that concedes the argument is incomplete: the lack of definitive insight into how amyloid accumulation triggers the neurodegenerative phase strongly suggests that though amyloid is necessary to trigger the disease, it is not sufficient.

Then, in the first person: the most critical poorly understood aspect of the hypothesis is how amyloid pathology drives the downstream cellular dysfunction that leads to the neurodegenerative phase, and *in my opinion* this central question needs to be answered in order to complete our understanding of the disease and provide a better framework for therapeutic intervention.

This is the most important passage in the document, and it is important for a structural reason rather than a rhetorical one.

Consider what a theory of a disease is for. It has to say what starts the process, what carries it forward, and why the clinical phenotype takes the form it does. The account answers the first question with more human evidence than any competing answer has. It answers the third with the organ-failure reframe. And it declines to answer the second, having stated that the second is the one that matters most.

There is a temptation to treat this as a weakness to be scored — a theory with a hole in the middle. That reading is available and it is superficial. Consider the alternative behaviours a synthesis could exhibit at this joint:

Assert through it. Claim that amyloid oligomers are directly neurotoxic, cite a portion of the several thousand papers reporting toxicity of applied aggregates, and proceed. This is what the field's standard exposition does.

Rename it. Introduce a term — *neuroinflammation*, *proteostatic collapse*, *network dysfunction* — that has the grammatical form of a mechanism and the content of a placeholder.

Absorb it. Broaden the theory until the joint is internal to it and no longer needs to be crossed.

Name it, size it, and leave it open. State that the joint is not established, enumerate the candidate mechanisms, specify what each would require, and decline to choose.

The account does the fourth. The result is that its central gap is a *specification* rather than an absence: a hole with stated dimensions, into which a competing account can be fitted and tested. This is the single most useful thing the document does for anyone who is not persuaded by it, and it is why the correct reading of this account is not as a rival theory but as a problem statement

with a very well-evidenced first clause.

11. Four Models, No Winner

The fourth figure of the account sets out the candidates. They are presented without ranking and with the note that they are not mutually exclusive. Stated at full strength, with what each requires and what would decide it:

Model 1 — Amyloid is a direct neurotoxin. Aggregated amyloid-β, in some conformation, damages neurons directly. This is the field's default and the account is notably cool toward it. Its four objections are stated in the document itself and are worth repeating because they are the best short critique of direct-toxicity models in print: amyloid accumulates for ten to twenty years before neurodegeneration is evident, so why is toxicity not observed during that phase; not all deposits behave alike, with diffuse plaques common in the normal aged brain and unassociated with overt pathology while neuritic plaques carry dystrophy, tau and gliosis; rodent deposition models show poor correlation between amyloid accumulation and neurodegeneration; and most of the toxicity literature rests on exogenously applied aggregates in culture, with in vivo demonstrations of Alzheimer-like neurodegeneration driven by amyloid *far fewer and far less well reproduced*. That is an amyloid researcher writing the strongest available case against the standard amyloid toxicity claim.

Model 2 — Amyloid is an indirect neurotoxin, mediated by glia. Aggregates activate microglia and astrocytes; the activated state does the damage. The genetics is congenial: *TREM2*, *PLCG2*, *ABI3* and a large fraction of the risk loci are microglial (Sims et al., 2017; Kunkle et al., 2019; Bellenguez et al., 2022). The account is careful here in a way that is characteristic — it notes that immune modulation can move amyloid and tau in *opposite* directions with the same stimulus, and concludes that we have only a rudimentary understanding of the temporality, directionality and impact of immune change in this disease. A mechanism whose sign flips depending on which pathology you measure is not yet a mechanism.

Subsequent work from the author's own laboratory illustrates the difficulty precisely. Antagonising interleukin-10 and interleukin-4 signalling in the brain by expressing decoy receptors reduced amyloid burden — robustly when expressed neonatally, and for the interleukin-4 decoy when delivered to the adult hippocampus — with opposing effects of the two decoys on glial proliferation. Neither manipulation altered tau pathology in two tau transgenic models, despite robust expression and clear effects on glia (Koller et al., 2025). Immune tone is demonstrably a lever on amyloid. Whether it is the lever that connects amyloid to tau is, on this evidence, not established.

Model 3 — Amyloid is a scaffold. The deposit is not itself the agent; it is a surface on which biologically active proteins co-accumulate, and the slow accumulation of those proteins eventually overcomes compensatory responses and triggers the degenerative phase. The account

names candidates: apolipoprotein E and clusterin, both genetically associated; heparan sulfate proteoglycans; the signalling molecules midkine and pleiotrophin; α1-antichymotrypsin. It notes that altered expression of plaque-associated proteins modifies amyloid, gliosis, tauopathy and dystrophic neurites in model systems. This is Chapter 13's subject and, in this evaluation's judgement, the strongest of the four.

Model 4 — Amyloid acts through the cerebrovasculature. Deposition alters vessels; barrier disruption, even transient, initiates downstream changes. The account is realistic about the limits: cerebral amyloid angiopathy plainly contributes in some patients, and fulminant disease occurs with little or none of it.

The four models are not equally well specified, and it is worth noticing what kind of claim each is. Model 1 is a claim about a molecule. Model 2 is a claim about a cell type. Model 3 is a claim about a *composite* — the deposit plus its cargo. Model 4 is a claim about a compartment. Only Model 3 makes the identity of the pathogenic entity depend on something other than amyloid-β itself, which is why it is the one that most changes what a therapy would have to do.

12. The Delay Is the Datum

Any candidate for the vacancy must explain a specific quantity, and the field discusses it far less than it discusses toxicity.

Amyloid deposition precedes neurodegeneration by ten to twenty years. This is not a soft impression; it is the consistent finding of the biomarker cascade in autosomal dominant and sporadic cohorts, and it is the axis of the account's own first figure. Something is present, in quantity, for two decades, and then something else begins.

This is the hardest constraint in the disease, and it discriminates between the four models better than any toxicity assay.

Model 1 fails it, on the account's own reasoning. If aggregated amyloid is directly toxic, the latency has to be explained by an auxiliary hypothesis: a threshold, a specific conformer that appears late, a failing compensatory mechanism. Each of these is a way of saying that something *other than the amyloid* determines when damage begins — which concedes the point.

Model 2 fails it in its simple form and passes in a modified one. Glial activation begins early; plaques are surrounded by reactive glia from the outset. For a glial mechanism to explain the delay, the glial response has to *change character* over time, from a tolerated or protective state to a damaging one. That is a real and testable proposition, and the field's single-cell work has been circling it, but the account correctly notes that its temporality and directionality are unresolved.

Model 3 passes it natively. This is the model's distinguishing virtue and the account states it explicitly: gradual, amyloid-dependent accumulation of these proteins could overcome compensatory responses to trigger the degenerative phase, *accounting for the long delay* between the onset of deposition and neurodegeneration in humans. If the pathogenic entity is the deposit plus decades of accreted cargo, then a two-decade latency is not an anomaly requiring an auxiliary hypothesis. It is the accretion time. The clock is built into the mechanism.

Model 4 is neutral on it. Vascular effects could be early or late; nothing in the model sets a timescale.

A latency of two decades is a great deal of information, and a mechanism that predicts it without adjustment should be preferred over mechanisms that must be fitted to it. That is an ordinary rule of inference and it is not applied here as often as it should be.

13. The Scaffold Model, and What Its Author Did Next

The scaffold model was the least fashionable of the four in 2020. It is also the one the author's own laboratory went on to test, and the results are the strongest post-2020 development bearing on the vacancy.

In 2024 a study from that laboratory compared the Alzheimer brain proteome and its network structure with the brain proteomes of amyloid-depositing mice, in order to separate conserved from divergent protein networks. The conserved networks define what the authors call an *amyloid responsome*. Proteins in the most conserved module, designated M42, accumulate in plaques, in cerebral amyloid angiopathy, and in dystrophic neuronal processes. The module's proteins bind amyloid fibrils in vitro. Overexpression of two of them — midkine and pleiotrophin — increases the accumulation of amyloid-β in plaques and in vascular amyloid, establishing that these are not passive passengers but modifiers of the deposition process itself. Midkine and pleiotrophin also co-accumulate with cardiac transthyretin amyloid, which places the phenomenon outside the brain and outside amyloid-β. The paper's closing position is that amyloid-scaffolded accumulation of numerous M42 proteins is a central mechanism mediating downstream pathophysiology (Levites et al., 2024).

Four features of this deserve to be stated plainly.

It is the 2020 model, executed. Midkine and pleiotrophin were named in the submission as candidate plaque-associated bioactive proteins. Four years later they are the module's exemplars and the mechanism is asserted as central. The account's author did not leave the vacancy for someone else; he moved his programme into it.

It converts a list into a network. The observation that plaques contain many proteins is decades old and has been treated as compositional trivia. Defining a *conserved* module — conserved between species, and between brain amyloid and a peripheral amyloidosis — turns the composition into a candidate mechanism, because conservation implies constraint.

It is bidirectional, which is a problem and an opportunity. Overexpressing midkine or pleiotrophin increases amyloid accumulation. That makes them modifiers of deposition. It does not yet make them the executors of degeneration. The scaffold model requires the accreted proteins to *do damage*, and the demonstration to date is that they do *deposition*. This is the gap inside the gap, and it is where the model must next be tested.

It has been extended regionally and the pattern is not uniform. A subsequent analysis of five matrisome proteins across four brain regions spanning the range of Alzheimer

neuropathological change found that midkine in plaques increased consistently with severity everywhere, while SPOCK3, COL25A1, EGFL8 and syndecan-4 accumulated markedly only in occipital cortex and hippocampus, sparsely in striatum, and not at all in cerebellum — with different members associating with neuritic versus diffuse deposits, and some overlapping with tau pathology and peri-plaque dystrophic processes (Tsering et al., 2025).

That last result is the most interesting one in this literature, because it bears directly on the objection the account raised against Model 1. The unanswered question there was: what is the difference between a diffuse plaque, which sits harmlessly in aged cortex, and a neuritic plaque, which is surrounded by dystrophy and tau? The scaffold model offers a candidate answer — *the cargo differs* — and the regional data show cargo that is not uniformly distributed and that sorts by deposit type. That is the shape of an explanation for the diffuse-versus-neuritic distinction, and no other model in the figure has one.

14. What a Theory Would Have to Supply to Fill the Vacancy

The value of a well-specified hole is that the specification doubles as a set of acceptance criteria. Anyone proposing to fill this one — whether with a scaffold mechanism, a glial mechanism, or something not in the figure at all — has to satisfy the following. The list is derived from the account's own objections and from the constraints in Chapters 5 and 12.

Criterion 1 — Latency. The mechanism must explain a ten-to-twenty-year gap between deposition and degeneration without an auxiliary threshold assumption. If it needs a threshold, it must say what accumulates to reach it.

Criterion 2 — Deposit-type selectivity. It must distinguish diffuse from neuritic deposits. Any account on which all amyloid is equivalently pathogenic contradicts the neuropathology of the normal aged brain.

Criterion 3 — Genetic compatibility. It must be consistent with three copies of *APP* being sufficient to produce the disease, with A673T being sufficient to prevent it, and with the *APOE* dose-response in both directions. This is a constraint, not an option; the genetics is the best-evidenced part of the field and any mechanism must survive it.

Criterion 4 — Species divergence. It must explain why rodent amyloid deposition models show poor correlation between deposition and neurodegeneration. A mechanism present in both species predicts neurodegeneration in mice, and mice do not deliver it. Either the mechanism is absent in rodents, or something protective is present in them, and the theory must say which.

Criterion 5 — Tau recruitment. It must specify how the process reaches tau, because tau is what tracks cognitive decline. A mechanism that damages neurons without engaging tau describes a different disease.

Criterion 6 — Therapeutic sign. It must predict the observed trial results: near-total plaque clearance in symptomatic patients yielding real but small clinical benefit, with benefit larger at lower tau burden.

The scaffold model, at present, passes 1, 2 and 6, is compatible with 3, is silent on 4, and has a promising but unproven route to 5 through the dystrophic neurite, which is the subject of Part IV. No other candidate does better, and it is worth noticing how much of that performance comes from the model having been designed against the account's own objections rather than against the field's.

Part IV — The Address

15. The One Feature That Distinguishes This Disease

Late in the account, in a section arguing for a less neuron-centric view, there is a paragraph that reads like an aside and is not one:

One area that deserves much more attention is the pathology referred to as dystrophic neurites. Dystrophic neurites surrounding amyloid plaques are the one feature that distinguishes Alzheimer's disease from almost all other neurodegenerative disorders except familial British and Danish dementia.

Take that claim seriously for a moment, because it is a strong one and it is not often made.

Neuronal loss is common to every neurodegenerative disease. Gliosis is common to all of them. Protein aggregation is common to all of them — that is the account's own argument for reading this disease as one of a family of proteinopathies. Tangles occur in progressive supranuclear palsy, corticobasal degeneration, chronic traumatic encephalopathy and several other tauopathies. Extracellular amyloid deposits occur in the cognitively normal aged brain in quantity.

What does not occur elsewhere is the specific composite lesion: an extracellular amyloid deposit encircled by swollen, vesicle-engorged axonal and dendritic processes. And the exception proves the rule in the strict sense — the two diseases that do produce it, familial British and Danish dementia, are the two diseases that produce a *different* amyloid from a *different* gene and are otherwise close phenocopies of Alzheimer's disease, clinically and pathologically. Both arise from mutations in *ITM2B*, which the account writes as *ITMB2* (see Chapter 29), and both deposit a non-amyloid-β peptide that behaves, in the tissue, the way amyloid-β does.

If a lesion appears in exactly the set of diseases characterised by an extracellular amyloid deposit of a secreted peptide in the brain parenchyma, and in no other, then it is a candidate for being the *disease-specific* consequence of that deposit. Everything else — the loss, the gliosis, the tangles — is the tissue's general repertoire. This is the specific part.

That is a strong argument and the account makes it in three sentences and moves on.

16. What Is Inside a Dystrophic Neurite

The lesion's contents are the reason it should be the focus rather than a curiosity.

Dystrophic neurites are swollen axonal structures filled with dysfunctional vesicles, many carrying endosomal and lysosomal markers. They are not a smear of debris; they have internal organisation, and the organisation is sequential — different layers form in a defined order in human brain, with distinct molecular composition (Sharoar et al., 2019), and a component of the swelling is dysfunctional tubular endoplasmic reticulum. They can be produced experimentally without any amyloid at all: inhibiting lysosomal proteolysis disrupts the axonal transport of degradative organelles and generates an Alzheimer-like axonal dystrophy (Lee, Sato & Nixon, 2011).

That last result matters more than the account allows. It means the dystrophic neurite is not simply a mechanical consequence of contact with a plaque. It is what an axon looks like when its degradative traffic fails, and a plaque is one way — perhaps the principal way in this disease — of making that traffic fail. The lesion is a *transport and disposal* lesion that happens to be located at the plaque.

Three consequences follow, and the account draws none of them.

The dystrophic neurite is where the account's four models converge. The endolysosomal contents implicate the disposal machinery. Contact with the deposit implicates the plaque. Peri-plaque glial activity implicates microglia. The accumulated matrisome and M42 proteins are found in dystrophic processes (Levites et al., 2024; Tsering et al., 2025). Whatever is happening in this structure is happening at the intersection of every candidate in Figure 4.

It has the right timescale. Dystrophic neurites are a feature of neuritic plaques and not of diffuse ones. The diffuse-to-neuritic transition is slow. If the lesion tracks that transition, it is on the two-decade clock the mechanism has to explain.

It is a compartment, not a cell. A dystrophic neurite is part of a neuron whose soma may be some distance away and may be intact. That reframes what is being lost. In the early phase, the disease may not be killing neurons at all; it may be disabling processes, one plaque-adjacent segment at a time, in cells that remain alive and countable. Neuronal counts would understate this by construction.

17. Seeding in Trans

The reason to care about all this is the account's fifth criterion: any mechanism has to reach tau.

The dystrophic neurite is the best available candidate for the place where it happens, and the key experiment is one the account cites without dwelling on it.

Human Alzheimer brain-derived pathological tau was injected into amyloid-plaque-bearing mice that do not overexpress tau. The injections recapitulated three Alzheimer-relevant tau lesions, and the crucial one is *neuritic plaque tau* — tau aggregation in the dystrophic processes around plaques. The conclusion was that amyloid plaques enhance the seeding of brain-derived tau by facilitating tau aggregation within neuritic plaques (He et al., 2018).

Read carefully, this is a solution to a problem that has bedevilled the field for thirty years: how a peptide outside the cell influences a protein inside it. The answer in this experiment is that it does not have to. It has to create a local compartment in which tau, arriving at low seed concentration, aggregates more readily than it would elsewhere. The plaque is not a signal transmitted to the neuron; it is a *place* — a site where axonal transport has stalled, degradative organelles have accumulated, cargo is not being cleared, and the local concentration of a seedable protein rises past the point at which it templates.

The account's own speculation about extracellular amino-terminal tau fragments points in a compatible direction. It notes that amyloid accumulation in mice and humans is associated with extracellular release of tau truncated before the microtubule-binding domain, that this species is a biomarker of amyloid and not of tau pathology, and asks — speculatively, and flagged as such — whether the cleavage that produces the secreted fragment leaves a truncated intracellular species more prone to aggregation. That is a hypothesis about a proteolytic event with two products, one of which is a marker and the other a substrate. It is testable and, as far as the record shows, it has not been directly tested in the seven years since it was written down.

18. The Experiment the Vacancy Demands

Putting Parts III and IV together yields a specific proposal, which this evaluation offers as the single most informative experiment available on the account's own terms.

The scaffold model says the pathogenic entity is the deposit plus its accreted cargo. The dystrophic-neurite work says the deposit's effect on tau is local and compartmental. The regional matrisome data say the cargo differs by deposit type and by region. These three combine into a testable proposition:

The difference between a diffuse plaque and a neuritic plaque is the presence of specific accreted proteins, and those proteins are what convert a deposit into a tau-seeding compartment.

The proposition has the property that good hypotheses have: it is straightforwardly refutable. If diffuse and neuritic plaques in the same brain, from the same region, matched for size and age, carry the same cargo, the proposition fails and the scaffold model loses its answer to the deposit-selectivity criterion. If neuritic plaques are distinguished by cargo, the next question — whether the cargo is cause or consequence of the dystrophy — is answerable by the manipulations already demonstrated for midkine and pleiotrophin.

Chapter 31 places this in a sequence with seven other experiments. It is placed first.

Part V — Where the Account Is Weak

The four weaknesses in this part are not the ones the account is usually attacked for. It is usually attacked for being an amyloid position, which is not a weakness but a thesis, and one better evidenced than most of its rivals. The weaknesses that matter are structural, and three of the four are omissions rather than errors.

19. The Absent Brainstem

Two hundred and forty-three references. No entry for the locus coeruleus, none for the brainstem, none on noradrenergic systems, and none from the body of work that established when the first abnormal tau appears in the human brain and where.

The omitted result is this. In an examination of 2,332 non-selected brains from individuals aged one to one hundred, using AT8 immunocytochemistry for abnormal tau and separate staining for amyloid-β, fifty-eight cases carried subcortical tau predominantly in the locus coeruleus with no abnormal cortical tau at all. Pretangle material restricted to subcortical sites was seen chiefly at younger ages. Amyloid plaques were present in 44.2% of the whole series and generally developed in the fifth decade. The authors' conclusion was that tauopathy associated with sporadic Alzheimer's disease may begin earlier than previously thought, and possibly in the lower brainstem rather than the transentorhinal region — and, stated flatly in the same abstract, that the first plaques occurred in the neocortex *after* the onset of tauopathy in the brainstem (Braak, Thal, Del Tredici et al., 2011).

This is the most direct available challenge to the account's ordering, it is human autopsy material at large scale, it was published nine years before the account was written, and it is not addressed.

The omission is the more conspicuous because the same laboratory's cortical staging paper *is* cited, as reference 186. The account's engagement with this body of work stops at 1991.

Now, in fairness — and this is the part a hostile reading would omit — the challenge is not fatal, and the reasons it is not fatal are interesting.

Pretangle tau is not Alzheimer's disease. Abnormal tau in the locus coeruleus is close to universal in middle age. Most people who have it never develop dementia. If the disease is defined as the syndrome, or as the plaque-and-tangle pathology, then subcortical pretangle material is a precondition rather than the disease, and its priority in time does not establish its priority in causation.

The amyloid interaction is real and it is in the account's favour. The in vivo work is the more informative here. Combining MRI measures of locus coeruleus integrity with amyloid and tau imaging in 174 individuals, and validating against locus coeruleus measures in 1,524 and 2,145 autopsy cases, one study found that lower locus coeruleus integrity was associated with elevated entorhinal tau among cognitively unimpaired people — and, critically, that the association with tau spread *beyond* the medial temporal lobe and with memory decline held *in the context of elevated amyloid*-β (Jacobs et al., 2021). Read plainly: early brainstem tau by itself stays where it is; it becomes a cortical disease in the presence of amyloid.

That reading is compatible with the account, and it is arguably a strengthening of it. The trigger, on this reading, is not the trigger of the *first tau*; it is the trigger of the *conversion* of a common, indolent, age-related subcortical tauopathy into a spreading cortical one. That is a more specific and more interesting claim than the one the account actually makes, and the account could have made it.

The criticism, therefore, is not that the account is refuted by the brainstem data. It is that the account **does not know they exist**, and a synthesis whose central commitment is the *ordering* of events in time is obliged to engage the largest human dataset on the ordering of events in time. The strongest version of the position requires that engagement, and the position is presented without it.

20. The Coverage Fraction, Stated Outward and Not Inward

Chapter 2 credited the account for its handling of misclassification. That credit stands. This chapter states the other half.

The account bounds *the field's data* with real care: it names the fraction of dementia attributable to this disease as roughly seventy per cent over the age of sixty, it separates Alzheimer's disease from dementia, and it handles misclassification in both directions. It also observes that pure Alzheimer's disease, defined by plaques and tangles alone, is rare over eighty, and that TDP-43 inclusions, hippocampal sclerosis, Lewy pathology and vascular change are variably present.

Having said all this, it never asks the corresponding question about itself: *of the cognitive decline that occurs in people who do have this disease, what fraction does the amyloid-first ordering explain?*

The answer is available and the account does not cite it. In a clinicopathological series of 1,079 individuals with repeated cognitive assessment and autopsy, ninety-four per cent had at least one neuropathology, seventy-eight per cent had two or more, and thirty-five per cent had four or more. Alzheimer pathology was the most frequent, at sixty-five per cent, and occurred in isolation in nine. More than two hundred and thirty distinct neuropathological combinations were observed, each in under six per cent of the cohort. Alzheimer pathology accounted, on average, for about half of the observed cognitive loss — with the person-specific proportion ranging from twenty-two per cent to one hundred (Boyle et al., 2018). The subsequent formalisation of limbic-predominant age-related TDP-43 encephalopathy as a distinct and common entity sharpened the same point (Nelson et al., 2019).

Two consequences.

The first is arithmetic. If the disease is seventy per cent of late-life dementia, and if within an affected individual it carries on average half the cognitive loss, then the account's ordering — even taken as entirely correct — is a complete explanation of something well short of the clinical problem it opens by describing. Nothing in the document acknowledges this, and the document opens with prevalence and cost figures for the whole of dementia.

The gap may be larger than the neuropathology alone suggests. In a plasma proteomic analysis across four cohorts totalling 2,139 participants, with a subset carrying paired brain data, the known neuropathologies accounted for only about half of the proteins associated with cognitive function — and many of the plasma proteins associated with those neuropathologies were not strongly correlated with their levels in brain, pointing to peripheral contributions to the clinical endophenotype (Afshar et al., 2025). The author is a co-author of that study. Its implication for

his own synthesis is that the unexplained fraction of cognitive decline is not merely the sum of the other named pathologies; some of it is not accounted for by anything currently on the list.

The second is methodological, and it is the more serious. A synthesis that carefully bounds the evidence it inherits and does not bound its own conclusions has applied its rigour asymmetrically. This is not dishonesty; it is the ordinary direction in which scepticism runs, and almost every synthesis in this field does the same thing. But the account's distinguishing virtue is precisely that it does *not* do the ordinary thing with evidence, which makes the one place it does conspicuous.

The remedy is a single sentence the account could have written and did not: *the ordering described here applies to the amyloid-initiated pathology, which accounts for roughly half the cognitive decline in the people who have it, and other processes account for the remainder.* That sentence costs the position nothing. Its absence is what allows a reader to finish the document with the impression that a theory of amyloid is a theory of dementia.

21. Reserve, Which Does the Work and Is Not Specified

Follow the load path through the account's own model and it terminates somewhere unexpected.

The pathology layer determines what is in the brain. The tolerance layer — reserve — determines when the person becomes ill. Therefore, for any given individual, *reserve determines the age of onset of the disease as experienced*. The account says as much: individuals with high reserve may tolerate much higher levels of pathology and neurodegeneration before showing symptoms; those with lower reserve show symptoms sooner. It also says that these early-life factors do not alter the presence of pathology but the ability to retain cognitive function despite it.

And then: the biological factors that underlie reserve are speculative.

A beam is holding up the clinical half of the model and the account says it does not know what the beam is made of. It offers the standard bipartite division — brain reserve as structural, cognitive reserve as strategic — but the division is taxonomic, not mechanistic, and the account is candid that it is.

This is a more serious omission after 2024 than it was in 2020, for the reason given in Chapter 9. Under a biological definition of the disease, a large population has the disease and is well, and the only thing standing between the pathology and the illness is the unspecified beam. The field's central clinical question has moved onto the one part of the model that is a placeholder.

What makes the omission harder to excuse is that the account cites, without noticing what it is, one of the two best pieces of human evidence about the mechanism.

The Christchurch case is cited as a genetic protection datum: an individual from the world's largest autosomal dominant kindred, homozygous for the *APOE3* Christchurch variant, who did not develop mild cognitive impairment until her seventies — three decades after the expected age of onset for her mutation — and who had unusually high brain amyloid levels with limited tau and limited neurodegeneration (Arboleda-Velasquez et al., 2019). A second such case, a man carrying a variant in *RELN*, was subsequently characterised alongside her, and the two were compared directly for common features of resistance (Lopera et al., 2023).

Read those cases against the account's own structure and something jumps out. These individuals were not protected by having *less amyloid*. They had a great deal of amyloid. They were protected at the joint — between the trigger and the executor — which is exactly the vacancy of Part III. Nature has run the experiment the account says has not been done: it has taken the trigger to full strength and interrupted the transmission to tau, twice, in human beings, with the resulting

phenotype documented for decades.

The account files these under genetics, where they demonstrate that protective variants exist. They belong under mechanism, where they demonstrate that the trigger-to-executor step is *interruptible* and localise the interruption to something in the apolipoprotein E and reelin signalling neighbourhood. That is the single largest missed inference in the document, and it is missed because the material is placed in the wrong chapter.

22. The Extrapolation That Is Not Flagged

The account is careful about inference almost everywhere. There is one place where it is not, and the whole therapeutic programme rests on it.

The A673T variant establishes that a lifetime of modestly reduced amyloid-β production confers protection against Alzheimer's disease and against age-related cognitive decline (Jonsson et al., 2012). The account treats this as the fundamental observation supporting the position — correctly — and then moves, without comment, to the therapeutic conclusion that robust, safe interventions targeting amyloid should prove efficacious in prevention trials.

These are different propositions.

The variant demonstrates that *never accumulating much amyloid* protects. A prevention trial in amyloid-positive individuals tests whether *removing amyloid that has already accumulated* protects. Between them lies everything the account itself has argued: that the deposit has been present for one to two decades by the time it is detectable; that co-accumulating proteins may be the operative agents; that the process may become self-sustaining. If the scaffold model of Chapter 13 is right, then two decades of amyloid positivity have already deposited a cargo whose removal is not addressed by removing amyloid-β.

This is a general feature of genetic evidence and it is well understood in other fields: a lifelong difference in exposure produces a larger effect than a late intervention on the same exposure, and the size of the gap between them is not derivable from the genetics. Estimating the effect of a late intervention from a lifelong genetic variant systematically overstates it.

None of this makes the prevention programme wrong. It makes the inference from A673T to prevention *unlicensed by the datum cited to support it*, and the account does not flag the step. Given how much of the document's care is spent flagging weaker inferences than this one — the reverse-causation caveat on depression and sleep, the discount on lifestyle epidemiology, the refusal to extrapolate from mouse immune modulation — the omission stands out. It is the one place where the argument's conclusion runs ahead of its evidence without the author saying so.

The honest formulation is: prevention is the most plausible remaining application of an amyloid-directed therapy, the genetics makes it worth attempting, and the genetics does not predict its magnitude.

23. Unfalsifiable at the Level of the Synthesis

The account's individual claims are refutable. Chapter 30 lists five ways to refute them. The synthesis as a whole is a different matter, and the property that makes it comprehensive is the property that makes it hard to kill.

Three features do the work.

Non-exclusive mechanisms. Four candidate models for the central joint, explicitly not mutually exclusive. This is epistemically honest — the truth may well be all four — and it means no single negative result at the joint can falsify the synthesis. Refute direct toxicity, and three models remain. Refute the glial route, and the scaffold stands.

A modifier layer that absorbs anomalies. Reserve, comorbidity, vascular health and lifestyle sit between the pathology and the phenotype. Any observation of pathology without symptoms is absorbed by reserve. Any observation of symptoms disproportionate to pathology is absorbed by comorbidity. Both absorptions are probably correct. Both also mean that the pathology-to-phenotype mapping cannot generate a decisive anomaly.

A trigger that is necessary but not sufficient. *Necessary and not sufficient* is the most defensible position available and the least vulnerable. Findings that amyloid alone does not produce the disease confirm the insufficiency. Findings that it precedes the disease confirm the necessity. The only observation that could straightforwardly refute the claim is a well-characterised case of the disease — with its full pathology — arising without amyloid, which the definitional structure of the field makes nearly impossible to produce, because a dementia without amyloid is by construction classified as something else.

That last point deserves to be sat with. Under the 2024 biological criteria, the proposition *Alzheimer's disease is triggered by amyloid* is close to being true by definition, because amyloid positivity is now part of what the term denotes. The account argued for biological definition and was right to. One consequence is that its own central claim has become harder to test.

A bounded alternative already exists in the literature and the account does not engage it. The probabilistic model of the disease retains amyloid as the principal driver in autosomal dominant cases while treating sporadic late-onset disease as arising from a variable combination of amyloid-dependent and amyloid-independent contributions of differing weight in different individuals (Frisoni et al., 2022). That model makes the explanatory fraction an explicit parameter rather than an unstated implication. Whatever its merits, it demonstrates that the bounded version of the account's position is writable.

The remedy, then, is not to abandon the synthesis. It is to state a bounded version: which fraction, which population, which route. The bounded version is testable and the account has all the materials to write it.

Part VI — What It Adds

24. A Boundary Condition, Not a Rival

The usual way to read this account is as one entry in a competition of theories: amyloid-first versus tau-first, versus inflammation-first, versus metabolic, versus vascular, versus infectious. Read that way, it is a strong entry and it is still only an entry, and a reader who does not accept its conclusion has no further use for it.

There is a more productive reading. The genetic core of Chapter 5 is not a claim that belongs to the amyloid position. It is a set of *facts about human beings* that any account of this disease must accommodate, whatever it thinks causes it. Stated as constraints rather than as arguments:

C1. Three copies of *APP* are sufficient to produce Alzheimer-type plaque and tangle pathology, in the first decade for deposition and the fourth to fifth for full pathology, and this holds in kindreds carrying duplication of a small chromosome-21 region containing *APP* alone.

C2. A single substitution in *APP* that lowers amyloid-β production confers lifelong protection against both the disease and age-related cognitive decline.

C3. Deterministic mutations in three genes, acting by three distinct biochemical routes, converge on the aggregation propensity of one peptide, and produce a disease whose natural history closely resembles the sporadic form.

C4. *APOE* alters risk bidirectionally and by allele dose, with ε4/ε4 approaching a deterministic genotype with near-full penetrance of the biology by sixty-five, and ε2/ε2 conferring near-complete protection.

C5. Interruption is possible at the joint: individuals with the deterministic mutation and heavy amyloid burden have deferred onset by three decades with limited tau and limited neurodegeneration.

Any theory of this disease must be consistent with all five. A theory in which amyloid-β is epiphenomenal must explain C1 and C2. A theory in which the primary lesion is metabolic, inflammatory or vascular must explain why gene dosage of one secreted peptide is sufficient to produce the syndrome. A theory in which tau is primary must explain C2 and C5 — must explain, in particular, why a variant that lowers amyloid protects for a lifetime, and why the two documented human resisters had heavy amyloid and little tau.

Read this way, the account's contribution is *not conditional on its conclusion*. Someone who thinks amyloid is a passenger still has to clear C1 through C5, and the account is where they are assembled and stated cleanly.

There is a symmetry worth naming. The account's weakest section, Part V's coverage problem, is the mirror of its strongest: it assembles the constraints and then does not notice that constraints bound a *portion* of the disease. Both the strength and the weakness come from the same source — the genetics is so good that it is easy to mistake it for the whole.

25. The Division of Labour

Combine Part III with Part VI and a specific working arrangement falls out, which this evaluation regards as the most useful way to hold the account.

The account owns the trigger. No competing proposal has evidence of this quality for what initiates the process. The genetics is human, it is convergent, it is bidirectional, and it includes a dosage series. Arguments about whether amyloid *initiates* should be conducted against this evidence and mostly are not.

The vacancy belongs to whoever can fill it. The step from trigger to degeneration is not owned by anyone, and the account says so. This is where the tau, glial, metabolic, endolysosomal, matrix and vascular programmes are actually working, whether or not they describe themselves that way. Much of what is presented in this field as a rival to the amyloid hypothesis is, structurally, a candidate for a position the amyloid hypothesis has advertised as open.

The two are separately gradeable. A programme can be right about the executor and wrong about the trigger, and vice versa. Treating the disease as a single contested question forces investigators into positions on the trigger that their evidence does not support, in order to justify work on the executor that stands on its own.

The author's own recent work is an example of the division being observed in practice. A study of the two haplotypes of the unfolded-protein-response kinase PERK found their canonical stress-response functions nearly indistinguishable, but identified a subset of translation events permitted under the risk-associated haplotype and not the other — among them the transcription factor DLX1, which is itself genetically linked to progressive supranuclear palsy risk, whose solubility shifts in tauopathy brain tissue and whose fly homologue, when silenced, reduces tau-induced toxicity in vivo (Lessard et al., 2026). Nothing in that work is about amyloid. It is a candidate mechanism for how tau becomes toxic — the second half of the vacancy, in a disease where the first half does not apply at all. Work on the executor does not require a position on the trigger, and this is what it looks like when the two are kept apart.

This division also clarifies a persistent confusion about what the trials showed. Near-total plaque clearance producing a quarter of the decline averted is a result about the *relative contribution of trigger and executor at the time of treatment*. It is not evidence against the trigger. It is a measurement of how much of the ongoing damage the trigger was still responsible for at that moment — and the answer, in a symptomatic brain, is: some, and not most. Everyone should have expected that, and the account did.

26. The Reframe That Survives the Debate

Of everything in the document, one thing will still be standing whatever happens to the amyloid hypothesis.

Symptomatic Alzheimer's disease is not early disease. It is organ failure.

The reason it survives is that it does not depend on the aetiology. It requires only that the pathology precede the syndrome by a long interval, which is now established by imaging and fluid biomarkers in multiple independent cohorts and is not seriously disputed by any party to the causal argument. Whether the pathology is amyloid, tau, a metabolic lesion or a clearance failure, the clinic meets the patient decades downstream of the initiating event.

Four consequences follow, and they are consequential in proportion to how uncomfortable they are.

The word *early* is being used incorrectly, across the entire field. *Early Alzheimer's disease* denotes mild cognitive impairment and mild dementia — the enrolment criteria for the trials of Chapter 7. Biologically these are late. Every use of the word in a trial title is a claim about the clinical scale that reads, to anyone who accepts this reframe, as a claim about the biological one. The account is direct about this and the field has not absorbed it.

A disease-modifying trial in symptomatic patients is testing a compressed hypothesis. It tests trigger removal, in an organ that has already failed, over eighteen to twenty-four months, against a functional endpoint. Three of those four elements are unfavourable by the account's own reasoning. A negative result in such a trial is close to uninformative about the biology; a positive one, as obtained, is informative mainly about how much trigger-dependence remains late.

Prevention is not one strategy among several; it is the only one addressed to the disease as described. Everything else is addressed to organ failure. This is the reframe's sharpest implication and the account states it.

Organ failure has its own therapeutic literature, and it is not aetiological. Heart failure is not treated by addressing the myocardial infarction that caused it. It is treated by unloading, by neurohormonal blockade, by rhythm control — by managing the failing system. The account's explicit call for multi-target therapy in symptomatic disease follows directly, and so does its most demanding claim: that restoration or regeneration of circuits will be needed for transformative effect, because arrest of a process is not recovery of function.

The last of these is where the analogy is doing the most work and getting the least examination, and Chapter 27 takes it up.

27. What the Analogy Licenses, and What It Does Not

Analogies in medicine are load-bearing whether or not their users intend them to be, and this one is used repeatedly. It is worth separating what it earns from what it assumes.

What it earns.

Staging is separate from aetiology. Heart failure is classified by functional stage and treated by stage; the cause is a separate axis. Applying this to Alzheimer's disease is a genuine importation and the field has partially adopted it in the biological-plus-clinical-staging structure of the 2024 criteria.

Multiple concurrent targets are normal, not a counsel of despair. No one expects a single agent to treat heart failure. The account's argument that symptomatic Alzheimer's disease will need combinations is not pessimism; it is the standard expectation for an organ-failure syndrome.

Late intervention on the initiating event is not expected to help much. Nobody proposes to treat established heart failure by reperfusing an infarct from ten years ago. Stated in those terms, the small effect sizes of Chapter 7 stop being a scandal and become the expected result of a category error the field made.

What it assumes.

That the failing organ has functional reserve to recruit. Heart failure therapy works substantially by improving the performance of surviving myocardium. The equivalent claim for the brain is that surviving circuits can be made to perform better. This is plausible and it is an assumption, and the account's regeneration clause is an acknowledgement that the assumption may not hold.

That the failure modes are enumerable. Heart failure therapy targets a small number of well-characterised compensatory pathways. The account's own description of late-stage disease — thousands of altered proteins and transcripts across every cell type — is an argument that the brain's failure is not similarly enumerable. The analogy supplies the strategy and simultaneously undercuts the tractability of executing it.

That the endpoint is measurable on a useful timescale. Heart failure trials have hard endpoints. Cognitive trials have composites. This is not a defect of the analogy so much as an obstacle to acting on it.

The net assessment is that the reframe is sound and the therapeutic programme it implies is under-specified. *Treat the failing organ with combinations and restore circuits* is a research direction rather than a plan, and the account presents it as the latter. That said, an under-specified plan aimed at the right stage of the disease is a better position than a well-specified plan aimed at the wrong

one, which is where the field spent two decades.

Part VII — Ledger, Refutation, and Tests

28. A Graded Ledger

The table grades twenty-four propositions drawn from the account, from its supporting evidence, and from the six years of record since it was written. Grades run:

Established — demonstrated in human material, replicated by independent groups, not seriously contested. **Well supported** — strong evidence including independent replication, with a stated limitation. **Supported, model-restricted** — good evidence confined to animal or cell systems, without human confirmation of the quantitative claim. **Contested** — reported and contradicted in the primary literature, unresolved. **Untested** — the decisive experiment has not been done. **Beyond the evidence** — asserted, or implied, at a strength the cited evidence does not carry.

#	Proposition	Grade	Principal evidence	What would change the grade
1	Deterministic *APP*/*PSEN1*/*PSEN2* mutations converge on amyloid-β aggregation propensity by three routes	Established	Human genetics with biochemical characterisation across many kindreds	A deterministic mutation acting on the downstream cascade by a route independent of amyloid
2	A673T confers lifelong protection through reduced amyloid-β production	Established	Jonsson 2012, population-scale	Failure to replicate the protective association in an independent population
3	*APOE* alters risk bidirectionally and by allele dose, through amyloid deposition	Established	Autopsy and amyloid imaging series; Reiman 2020 for ε2/ε2	— not in dispute
4	*APOE4* homozygosity approaches a deterministic genotype	Established	Fortea 2024; 3,297 pathological and 10,039 clinical cases	— strengthened since the account was written
5	*APP* gene dosage is sufficient to produce the pathology	Established	Trisomy 21; small chromosome-21 duplication kindreds	— the small duplications are the control that closes this
6	Amyloid deposition precedes symptom onset by twenty to thirty years	Established	Biomarker cascade in dominantly inherited and sporadic cohorts	— not in dispute

#	Proposition	Grade	Principal evidence	What would change the grade
7	Cortical tau imaging signal and volumetric neurodegeneration begin roughly together, after amyloid	Well supported	Longitudinal biomarker cohorts	Cohorts showing tau PET regularly preceding amyloid positivity
8	Abnormal tau appears in the locus coeruleus before the first cortical plaque	Established — **and not addressed by the account**	Braak, Thal & Del Tredici 2011; 2,332 brains	— the finding is not in doubt; its interpretation is
9	Symptomatic disease represents long-standing pathology, correctly described as organ failure	Established	Biomarker ordering plus post-mortem proteomics and transcriptomics	— this is the account's most durable claim
10	Reserve alters the pathology-to-symptom threshold without altering the pathology	Well supported	Educational and cognitive-reserve epidemiology; clinicopathological discordance	Demonstration that reserve proxies alter deposition itself
11	Lifestyle factors act mainly through vascular health and reserve, not on amyloid pathophysiology	Supported, contested	Epidemiology measuring dementia rather than the disease	A prevention trial altering amyloid biomarkers through lifestyle
12	Amyloid-β aggregation is necessary to trigger the disease	Well supported, and increasingly definitional	Constraints 1–5 of Chapter 24	A fully characterised case of the syndrome and pathology without amyloid — which the 2024 criteria make hard to constitute
13	Amyloid-β is not sufficient to produce the disease	Established	The account's own four objections; diffuse plaques in normal aged cortex	— conceded by the author
14	The mechanism linking amyloid accumulation to tau and neurodegeneration is unresolved	Established as a statement about the field	Four non-exclusive models offered; none endorsed	Any of the experiments in Chapter 31
15	Aggregated amyloid-β is a direct neurotoxin in vivo	**Contested**	Extensive in vitro literature; in vivo demonstrations sparse and poorly reproduced, per the account	A reproducible in vivo demonstration at physiological exposure

#	Proposition	Grade	Principal evidence	What would change the grade
16	Glial activation mediates amyloid's downstream toxicity	Supported, direction unresolved	Microglial risk genetics; immune modulation moving amyloid and tau in opposite directions	Temporal resolution of when the glial response changes sign
17	Amyloid-scaffolded accumulation of co-depositing proteins mediates downstream pathophysiology	Supported, model-restricted — **strengthened since 2020**	Levites 2024 (conserved M42 responsome); Tsering 2025 (regional, deposit-type specificity)	Demonstration that the accreted proteins damage neurons, not only that they modify deposition
18	Midkine and pleiotrophin modify amyloid deposition	Well supported	Overexpression increases plaque and vascular amyloid; fibril binding in vitro; co-accumulation with cardiac transthyretin amyloid	— robust; the open question is downstream, not here
19	Dystrophic neurites distinguish this disease from other neurodegenerative disorders, except familial British and Danish dementia	Well supported	Comparative neuropathology	A comparable peri-deposit dystrophy characterised in an unrelated proteinopathy
20	Plaques promote tau seeding locally, within peri-plaque dystrophic processes	Well supported, model-restricted	He 2018: human-brain-derived tau into plaque-bearing, non-tau-overexpressing mice	Human demonstration; or failure to replicate the neuritic-plaque-tau result
21	Trigger-directed therapy after neurodegeneration is established yields limited benefit	Established — **predicted, then confirmed**	CLARITY-AD and TRAILBLAZER-ALZ 2: near-complete plaque removal, roughly a quarter of decline averted	— the prediction and the result agree
22	Amyloid-directed intervention in the preclinical window will prevent or substantially delay the disease	**Untested**	A4 was negative, but amyloid *rose* in both arms; the drug did not clear plaque	The running trials of plaque-clearing agents in preclinical and at-risk populations
23	The amyloid-first ordering accounts for the bulk of late-life cognitive decline	**Beyond the evidence**	Never claimed outright; implied by framing. Boyle 2018: ~50% of decline on average, 22–100% person-specific	An explicit statement of the explanatory fraction, which would remove the row

#	Proposition	Grade	Principal evidence	What would change the grade
24	Circuit restoration or regeneration will be required for transformative benefit in symptomatic disease	Conjecture, stated as such	None; an inference from the organ-failure reframe	— appropriately labelled by the author

Four rows deserve comment.

Row 8 is the account's largest omission and not its largest error. The brainstem finding is established and it is unaddressed. As Chapter 19 argued, engaging it would probably have strengthened the position rather than weakened it, by converting a claim about the first tau into a claim about the conversion of an indolent subcortical tauopathy into a spreading cortical one.

Row 17 is where the account's own programme went, and the movement is to its credit. A model that was one of four unranked candidates in 2020 is now supported by a conserved cross-species proteomic module, with two of its members shown to modify deposition. The grade is held at *model-restricted* because the demonstrated function is modification of deposition rather than production of damage.

Row 22 is the position's principal exposure. It is *untested*, not *refuted*, and Chapter 8 gives the reason. It will not remain a comfortable entry. Two of the three running trials that could settle it are due to report within the period over which this assessment would need to be revised.

Row 23 is the row this evaluation would most like to see removed, and it can be removed by a sentence. The account never states that the ordering explains all of late-life cognitive decline. It also never states that it does not, while opening with the prevalence and cost of all dementia. One clause stating the explanatory fraction would cost the argument nothing and would close the gap.

29. Four Errors of Record

These were checked against the primary sources. Three are substantive enough to record; the fourth is typographic and is listed for completeness.

E1 — The Christchurch case is assigned to the wrong mutation. The text describes the *APOE3* Christchurch carrier as a carrier of the familial-Alzheimer-linked *PSEN1* **E280G** variant. The primary report describes the individual as a member of the world's largest autosomal dominant Alzheimer's disease kindred (Arboleda-Velasquez et al., 2019) — the Colombian **E280A** kindred. E280G is not that mutation. The bibliography entry itself is correct; the error is in the text. It does not affect the argument, and it is the kind of error that propagates when a paper is cited from a secondary source.

E2 — The gene symbol for familial British and Danish dementia is inverted. The text gives *ITMB2 (BRI2)*; the symbol is **ITM2B**. Cosmetic in isolation, but this gene carries the account's best analogical argument for the amyloid-to-degeneration step — a non-amyloid-β peptide producing a close phenocopy of the disease — so the reference should be findable.

E3 — Thal staging is named in the text and Thal's paper is not in the bibliography. The text refers to Braak and Thal staging of tangle and amyloid pathology and cites two items: the NIA-AA neuropathological assessment guidelines, which operationalise Thal phasing, and Braak & Braak 1991. Thal's primary paper (Thal et al., 2002) does not appear among the 243 references. The claim is supported by the guideline citation and the attribution is loose.

E4 — A malformed in-text citation range. The claim that mechanistic insight into tau-induced neurodegeneration is lacking carries the range "241,-243". The three bibliography items exist and are on topic; only the punctuation is broken.

Two observations about the register in which these should be read. First, an error rate of three substantive slips across 243 references and ten pages is low for a synthesis of this scope, and every load-bearing genetic citation checked in this evaluation resolved correctly. Second, E1 is the interesting one, because the case it misidentifies is — as Chapter 21 argued — the most important single piece of evidence in the document about the joint the document says is unknown. The account cites it, misattributes its kindred, and files it in the wrong section.

30. Five Conditions of Refutation

The account does not state what would refute it. These five are derived from propositions it commits to in writing, and each is answerable by an experiment that could be specified today.

R1 — The ordering. *If a large post-mortem or longitudinal biomarker series shows that cortical tau pathology, neurodegeneration or cellular dysfunction regularly and demonstrably precedes amyloid deposition in individuals who go on to develop the disease, the ordering claim fails.* The qualifier *cortical* is doing necessary work: the brainstem pretangle finding of Chapter 19 already satisfies a weaker version of this condition, and the account's ordering survives it only by treating subcortical pretangle tau as a precondition rather than as the disease. That treatment must be argued, not assumed.

R2 — The genetic route. *If the deterministic APP and PSEN mutations are shown to act on the downstream cascade by a route independent of amyloid-β aggregation, the load-bearing wall falls.* Candidate routes exist in the literature the account itself cites — effects of these mutations on endosomal, autophagic and lysosomal trafficking that appear in some cases independent of amyloid processing.

R3 — Prevention at full engagement. *If an agent achieving near-complete and durable removal of amyloid in amyloid-positive but asymptomatic individuals produces no reduction in subsequent tau accumulation, neurodegeneration or dementia incidence over an adequate interval, the therapeutic core fails.* This is the account's own named test and it is running. A4 does not satisfy it, for the reason given in Chapter 8.

R4 — The scaffold. *If diffuse and neuritic plaques, matched for region, size and donor, are shown to carry the same accreted protein cargo, the scaffold model loses its explanation of deposit-type selectivity* — and with it its principal advantage over direct toxicity.

R5 — The interruption. *If the documented cases of extreme resistance are shown to have been protected by reduced amyloid exposure rather than by interruption between amyloid and tau, the inference of Chapter 21 fails* and the resilience cases revert to being evidence about the trigger rather than about the joint. The published characterisation — heavy amyloid, limited tau, three-decade deferral — currently says the opposite.

31. Eight Experiments, in Order

Ordered by what the answer to each unlocks for the next, not by feasibility.

1. Cargo differential between diffuse and neuritic plaques. Same brain, same region, matched deposits; quantitative proteomics or multiplexed imaging of the responsome module. This is the pivot experiment of Chapter 18 and everything downstream is easier if it resolves.

2. Cargo manipulation against tau seeding. In plaque-bearing, non-tau-overexpressing animals seeded with human-derived tau, ask whether modulating individual responsome proteins changes neuritic-plaque tau formation. This converts the correlation of experiment 1 into a mechanism, and the tools already exist for midkine and pleiotrophin.

3. The dystrophic neurite as a compartment, quantified in human tissue. How much of the axonal and dendritic arbor is lost to peri-plaque dystrophy at each disease stage, in surviving neurons? If the early lesion is compartmental rather than cellular, neuronal counting has been the wrong instrument.

4. The truncated-tau hypothesis, tested directly. The account's own speculation: does the cleavage that produces secreted amino-terminal tau leave an aggregation-prone intracellular remainder? Identify the protease, block it, measure both products.

5. Locus coeruleus tau conversion, prospectively. Does the transition from indolent subcortical pretangle tau to spreading cortical tau require amyloid, as the cross-sectional in vivo data suggest? A longitudinal design with locus coeruleus imaging, tau PET and amyloid status can answer this, and the answer decides whether Chapter 19's rescue of the account holds.

6. Prevention at full target engagement. Already running. Nothing in this evaluation should be read as a reason to wait for it before doing 1 through 5; the point of a well-specified vacancy is that it can be worked on independently of the trial programme.

7. Reserve, mechanised. Identify a measurable biological substrate for the tolerance layer of Chapter 3 in humans, prospectively, and test whether it predicts the pathology-to-symptom offset. Under a biological definition of the disease this is no longer an ancillary question.

8. The explanatory fraction, stated. In an autopsy series with longitudinal cognition, estimate what proportion of decline is attributable to the amyloid-initiated sequence specifically, as opposed to co-occurring pathologies. This is a study, not a rhetorical concession, and it would close row 23.

32. Limitations of This Evaluation

It reads a synthesis, not a laboratory. The primary object is a ten-page argument and its bibliography. The author's experimental programme is used to grade the synthesis and is not itself audited; a proper evaluation of the laboratory would be a different and larger undertaking.

It grades a 2020 document against a 2026 record. This is fair when a prediction is scored and unfair when an omission is charged. Chapter 19's omission is charged because the omitted work was nine years old at the time of writing. Chapters 6 through 9 score predictions, which is legitimate. Chapter 13 credits movement the author made after writing, which is generous but accurate.

The scaffold model is favoured, and the favouring is an argument rather than a result. Chapter 12's latency criterion is this evaluation's own reasoning applied to the account's own objections. It is not the field's consensus and it may be wrong. A reader who rejects the latency argument should discount Chapters 13, 18 and 31 accordingly.

Two of the four models are treated more briefly than they deserve. The glial and vascular routes each have large literatures that are not surveyed here. They are graded on what the account says about them and on the constraints of Chapter 14, which is a thinner basis than the scaffold model receives.

Citations were verified; interpretations were not verifiable. Every reference in this paper was checked against the indexed record for existence and content. That procedure catches misattribution. It does not catch a plausible reading of a real paper that its authors would reject.

33. Conclusion — The Honest Cascade

The amyloid cascade hypothesis has been declared dead so many times that the declarations have stopped carrying information. This account is a reason to be more careful, because it shows what the position looks like when it is held by someone unwilling to overstate it.

What it gets right is a matter of record. It said in 2020 that the homozygous ε4 genotype was approaching causality, and in 2024 a large series said so formally. It said that removing the trigger from a brain that had already degenerated would help a little and not much, and two trials that cleared essentially all of the plaque averted about a quarter of the decline. It said the disease should be defined biologically rather than clinically, and in 2024 it was. It named the diffuse-versus-neuritic problem, the latency problem, the rodent-correlation problem and the exogenous-application problem as the four things wrong with the direct-toxicity account of amyloid — an unusually complete indictment of the field's default position, written by one of its proponents.

What it gets wrong is narrower than the field's standard criticisms and harder to dismiss. It bounds everyone's data but its own. It leaves reserve — which decides when a person with the disease becomes a patient — explicitly unspecified, at a moment when a biological definition of the disease has made that the central clinical question. It infers from a variant that lowers amyloid for a lifetime to a therapy that removes it late, and does not mark the step. And in two hundred and forty-three references it does not look at the brainstem, where the first abnormal tau in the human nervous system appears decades before the first cortical plaque.

But the reason to read it is none of these. It is the paragraph in the middle where the argument stops and its author writes that he does not know how amyloid produces the disease, that this is the question that matters most, that there are four ways it might work, and that he is not going to choose between them.

That paragraph does something the rest of the literature mostly does not. It converts a contested claim into a specified problem. The trigger is established about as well as anything in this field is established, and the step from trigger to degeneration is vacant, and the vacancy has dimensions: it must produce a two-decade latency, distinguish diffuse deposits from neuritic ones, survive the genetics, explain why rodents do not degenerate, reach tau, and predict the trial results we now have. Any account that fills it is a theory of Alzheimer's disease. Any account that does not is a theory of something adjacent.

The most interesting development since 2020 is that the author's own laboratory went into the vacancy and came back with a candidate — a conserved module of proteins that accumulates on the deposit, modifies it, appears in the dystrophic processes around it, and does so with a regional

and deposit-type specificity that maps onto the very distinction he identified as unexplained. Whether the scaffold model is right is not yet decided. That it was proposed as one of four unranked options and then pursued by the person who listed it is the behaviour of a research programme rather than a position.

The cascade may or may not survive. What should survive is the way this document handles the difference between what it knows and what it does not — which is, in the end, the only thing that distinguishes a theory from a commitment.

References

Afshar S, Dammer EB, Bian S, Bennett DA, Mohs R, Beauregard D, Dwyer J, Hales CM, Goldstein FC, Parker MW, et al. Plasma proteomic associations with Alzheimer's disease endophenotypes. *Nature Aging*. 2025;5(10):2104-2124. PMID 40931114.

Arboleda-Velasquez JF, Lopera F, O'Hare M, Delgado-Tirado S, Marino C, Chmielewska N, Saez-Torres KL, Amarnani D, Schultz AP, Sperling RA, et al. Resistance to autosomal dominant Alzheimer's disease in an APOE3 Christchurch homozygote: a case report. *Nature Medicine*. 2019;25(11):1680-1683. PMID 31686034.

Bellenguez C, Küçükali F, Jansen IE, Kleineidam L, Moreno-Grau S, Amin N, Naj AC, Campos-Martin R, Grenier-Boley B, Andrade V, et al. New insights into the genetic etiology of Alzheimer's disease and related dementias. *Nature Genetics*. 2022;54(4):412-436. PMID 35379992.

Boyle PA, Yu L, Wilson RS, Leurgans SE, Schneider JA, Bennett DA. Person-specific contribution of neuropathologies to cognitive loss in old age. *Annals of Neurology*. 2018;83(1):74-83. PMID 29244218.

Braak H, Braak E. Neuropathological stageing of Alzheimer-related changes. *Acta Neuropathologica*. 1991;82(4):239-59. PMID 1759558.

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):960-9. PMID 22002422.

Fortea J, Pegueroles J, Alcolea D, Belbin O, Dols-Icardo O, Vaqué-Alcázar L, Videla L, Gispert JD, Suárez-Calvet M, Johnson SC, et al. APOE4 homozygozity represents a distinct genetic form of Alzheimer's disease. *Nature Medicine*. 2024;30(5):1284-1291. PMID 38710950.

Frisoni GB, Altomare D, Thal DR, Ribaldi F, van der Kant R, Ossenkoppele R, Blennow K, Cummings J, van Duijn C, Nilsson PM, et al. The probabilistic model of Alzheimer disease: the amyloid hypothesis revised. *Nature Reviews Neuroscience*. 2022;23(1):53-66. PMID 34815562.

Golde TE. Alzheimer's disease — the journey of a healthy brain into organ failure. *Molecular Neurodegeneration*. 2022;17(1):18. PMID 35248124.

Hardy JA, Higgins GA. Alzheimer's disease: the amyloid cascade hypothesis. *Science*. 1992;256(5054):184-5. PMID 1566067.

Hardy J, Selkoe DJ. The amyloid hypothesis of Alzheimer's disease: progress and problems on the road to therapeutics. *Science*. 2002;297(5580):353-6. PMID 12130773.

He Z, Guo JL, McBride JD, Narasimhan S, Kim H, Changolkar L, Zhang B, Gathagan RJ, Yue C, Dengler C, et al. Amyloid-β plaques enhance Alzheimer's brain tau-seeded pathologies by facilitating neuritic plaque tau aggregation. *Nature Medicine*. 2018;24(1):29-38. PMID 29200205.

Jack CR, Andrews JS, Beach TG, Buracchio T, Dunn B, Graf A, Hansson O, Ho C, Jagust W, McDade E, et al. Revised criteria for diagnosis and staging of Alzheimer's disease: Alzheimer's Association Workgroup. *Alzheimer's & Dementia*. 2024;20(8):5143-5169. PMID 38934362.

Jacobs HIL, Becker JA, Kwong K, Engels-Domínguez N, Prokopiou PC, Papp KV, Properzi M, Hampton OL, d'Oleire Uquillas F, Sanchez JS, et al. In vivo and neuropathology data support locus coeruleus integrity as indicator of Alzheimer's disease pathology and cognitive decline. *Science Translational Medicine*. 2021;13(612):eabj2511. PMID 34550726.

Jonsson T, Atwal JK, Steinberg S, Snaedal J, Jonsson PV, Bjornsson S, Stefansson H, Sulem P, Gudbjartsson D, Maloney J, et al. A mutation in APP protects against Alzheimer's disease and age-related cognitive decline. *Nature*. 2012;488(7409):96-9. PMID 22801501.

Karran E, De Strooper B. The amyloid hypothesis in Alzheimer disease: new insights from new therapeutics. *Nature Reviews Drug Discovery*. 2022;21(4):306-318. PMID 35177833.

Koller EJ, McFarland KN, Angelle C, Howard J, Ryu D, Dillon KD, Erquizi A, Beheray M, De La Cruz EG, Cruz PE, et al. Antagonizing Il10 and Il4 signaling via intracerebral decoy receptor expression attenuates Aβ accumulation. *Acta Neuropathologica Communications*. 2025;13(1):51. PMID 40055831.

Kunkle BW, Grenier-Boley B, Sims R, Bis JC, Damotte V, Naj AC, Boland A, Vronskaya M, van der Lee SJ, Amlie-Wolf A, et al. Genetic meta-analysis of diagnosed Alzheimer's disease identifies new risk loci and implicates Aβ, tau, immunity and lipid processing. *Nature Genetics*. 2019;51(3):414-430. PMID 30820047.

Lee S, Sato Y, Nixon RA. Lysosomal proteolysis inhibition selectively disrupts axonal transport of degradative organelles and causes an Alzheimer's-like axonal dystrophy. *Journal of Neuroscience*. 2011;31(21):7817-30. PMID 21613495.

Lessard CB, Rubio Rubio D, Tolton S, Criado-Marrero M, Ravi S, Garza TN, Koren J, Philips J, Bagchi P, McFarland K, et al. Progressive Supranuclear Palsy PERK Haplotype B Selectively Translates DLX1 Promoting Tau Toxicity. *Journal of Neuroscience*. 2026;46(13). PMID 41708330.

Levites Y, Dammer EB, Ran Y, Tsering W, Duong D, Abreha M, Gadhavi J, Lolo K, Trejo-Lopez J, Phillips J, et al. Integrative proteomics identifies a conserved Aβ amyloid responsome, novel plaque proteins, and pathology modifiers in Alzheimer's disease. *Cell Reports Medicine*. 2024;5(8):101669. PMID 39127040.

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. Resilience to autosomal dominant Alzheimer's disease in a Reelin-COLBOS heterozygous man. *Nature Medicine*.

2023;29(5):1243-1252. PMID 37188781.

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

Reiman EM, Arboleda-Velasquez JF, Quiroz YT, Huentelman MJ, Beach TG, Caselli RJ, Chen Y, Su Y, Myers AJ, Hardy J, et al. Exceptionally low likelihood of Alzheimer's dementia in APOE2 homozygotes from a 5,000-person neuropathological study. *Nature Communications*. 2020;11(1):667. PMID 32015339.

Salloway S, Farlow M, McDade E, Clifford DB, Wang G, Llibre-Guerra JJ, Hitchcock JM, Mills SL, Santacruz AM, Aschenbrenner AJ, et al. A trial of gantenerumab or solanezumab in dominantly inherited Alzheimer's disease. *Nature Medicine*. 2021;27(7):1187-1196. PMID 34155411.

Sharoar MG, Hu X, Ma XM, Zhu X, Yan R. Sequential formation of different layers of dystrophic neurites in Alzheimer's brains. *Molecular Psychiatry*. 2019;24(9):1369-1382. PMID 30899091.

Sims JR, Zimmer JA, Evans CD, Lu M, Ardayfio P, Sparks J, Wessels AM, Shcherbinin S, Wang H, Monkul Nery ES, et al. Donanemab in Early Symptomatic Alzheimer Disease: The TRAILBLAZER-ALZ 2 Randomized Clinical Trial. *JAMA*. 2023;330(6):512-527. PMID 37459141.

Sims R, van der Lee SJ, Naj AC, Bellenguez C, Badarinarayan N, Jakobsdottir J, Kunkle BW, Boland A, Raybould R, Bis JC, et al. Rare coding variants in PLCG2, ABI3, and TREM2 implicate microglial-mediated innate immunity in Alzheimer's disease. *Nature Genetics*. 2017;49(9):1373-1384. PMID 28714976.

Sperling RA, Donohue MC, Raman R, Rafii MS, Johnson K, Masters CL, van Dyck CH, Iwatsubo T, Marshall GA, Yaari R, et al. Trial of Solanezumab in Preclinical Alzheimer's Disease. *New England Journal of Medicine*. 2023;389(12):1096-1107. PMID 37458272.

Thal DR, Rüb U, Orantes M, Braak H. Phases of A beta-deposition in the human brain and its relevance for the development of AD. *Neurology*. 2002;58(12):1791-800. PMID 12084879.

Tsering W, Philips JL, Golde TE, Villareal JA, Prokop S. Brain Region-Specific Accumulation of Amyloidosis-Associated Proteins in Postmortem Brain Tissues of Alzheimer's Disease Patients. *Molecular Neurobiology*. 2025;62(12):15759-15778. PMID 40705253.

van Dyck CH, Swanson CJ, Aisen P, Bateman RJ, Chen C, Gee M, Kanekiyo M, Li D, Reyderman L, Cohen S, et al. Lecanemab in Early Alzheimer's Disease. *New England Journal of Medicine*. 2023;388(1):9-21. PMID 36449413.