The Dial and the Decade

What can be measured in a living person with Alzheimer's disease, and what cannot

Benjamin Aaron Gustafsson

AdultCognitiveDisease.com

November 2026

A standalone paper on method. The model describes a disease that runs for fifty years inside living neurons; almost none of that can be measured in a living person. Claims about levels are reached by the instruments. Claims about rates, about order, and about what could be undone are not — and the last of those can be settled only by treating somebody and watching.

Contents

How to read this paper
1. A Model About Things Nobody Can See
2. Four Kinds of Claim
3. What the Instruments Actually Reach
4. How the Field Solved This: It Redefined the Disease
5. The Order Problem
    5.1 The twenty-five-year curve
    5.2 The brains of the young
    5.3 Why post-mortem series cannot settle order
6. The Price of Our Own Redraw
7. The One Case Where a Read-Out Changed the Answer
8. The Best Decision the Corpus Made
9. Three Tiers: What Is Closable, What Is Reachable Only by Inference, What Is Closed
10. Why This Corpus Grades Everything
11. What Follows
12. Graded Ledger
13. What Would Change This Account
Companion papers in this corpus
References

Abstract

This paper asks a question about instruments rather than about biology, and it asks it of our own model as much as of anyone else's.

The account this site has built describes a disease that runs for fifty years inside living neurons. It says which cells fail first, in what order, how fast, and — most consequentially — that some of the damage is silence rather than death and might therefore be undone. Almost none of that can be measured in a living person.

The claims divide into four kinds, and the instruments reach only one. **Levels** — how much of something is present — are measured well, by scanner and by blood test. **Rates** — how fast a process runs — are almost never measured, because every human read-out returns a quantity rather than a speed. **Order** — what happened before what, inside one person — is not measured at all; the field's two most famous timing results are both assembled from people seen once each, and the authors of both say so in their own papers. **Reversibility** — whether a silenced cell could be woken — has no instrument even in principle, and can be demonstrated only by treating someone and watching.

The field's response to this has been to redefine the disease as the things that can be measured: amyloid, tau and neurodegeneration, the three axes of the research framework now in general use. That is a defensible move and it has made research communicable. It also means that none of the four processes our model is actually about has a standard read-out, and that our model is therefore, in the framework's terms, not about Alzheimer's disease at all.

One section is constructive. There is a case in this literature where a mechanism-level read-out existed, was used to choose patients, and changed what a trial found. It is worth studying, because it shows what the rest of the field is missing and roughly what it would cost to get it.

The paper ends where it has to: with the admission that the stacked redraw of our own model, which we think is more faithful to the biology, is *harder* to check than the relay it replaced — and with the reason we grade every claim rather than asserting it.

How to read this paper

This is a standalone paper about method. It is not part of the four-paper series on the stacked phases, though it is about them; it can be read before or after them, and on its own.

It contains no new biology. Everything it cites is in the corpus elsewhere, and the point of assembling it here is to ask a different question of the same material: not *what is true*, but *what could be found out, by whom, with what*.

Claims are graded with the four grades the corpus uses throughout. **Established** means replicated in human material by more than one group. **Probable** means good evidence with a gap. **Inference** means a conclusion drawn here from evidence gathered for another purpose. **Speculative** means coherent and untested. Readers will notice that the grades in this paper attach mostly to statements about *the state of the instruments*, which is an unusual thing to grade, and that several of the most important claims are graded by what is missing rather than by what is present.

Two companion papers go deeper on parts of this. *The Unmeasured Rate* works the rate problem through one case — autophagic flux — in far more detail than here. *The Biomarker Cascade* works out which existing biomarkers detect which phase. This paper is the floor underneath both.

1. A Model About Things Nobody Can See

Start with what the model claims, stated flatly.

It claims that the disease begins in the third decade of life in a few small nuclei at the base of the brain. It claims those neurons fail in *function* long before they die, so that the relevant quantity is connections lost rather than cells counted. It claims the brain's immune cells are held at rest by a set of restraints that are released one at a time across the sixth and seventh decades. It claims a protective matrix around inhibitory neurons is digested in the eighth. It claims that beneath all three, a sorting step inside the cell fails to

return what it should. And it claims — this is the part that matters to a patient — that a neuron which has lost its connections but is still alive is silenced rather than dead, and might in principle be recovered.

Now ask of each claim: what instrument, applied to a living person, would show it?

For most of them the answer is none. Not *none yet*; none. There is no scan that reports how long a protein sits in a compartment. There is no blood test for the rate at which a receptor is returned to a membrane. There is no way to ask of a particular silent neuron in a particular living person whether it could be woken.

This is not a complaint about underfunding. It is a structural feature of the subject, and it has consequences for what anyone — this corpus very much included — is entitled to say.

2. Four Kinds of Claim

The useful division is not by biology but by what sort of quantity the claim is about. There are four, and they are not equally accessible.

A level is a quantity present at a moment. How much amyloid is in the cortex. How much tau. How much cortical thickness remains. How much of a protein is in the spinal fluid or the blood. Levels are what instruments naturally return, because an instrument samples at a time.

A rate is a quantity per unit time. How fast material moves through a compartment. How fast a receptor is recycled. How fast waste is cleared. Rates are what most of the mechanisms in this disease actually are. A level and a rate can come apart completely: a compartment can be full because delivery into it rose, or because export out of it fell, and the two have opposite therapeutic implications. A single measurement cannot distinguish them.

An order claim is about sequence within one individual. Not "in this population, people with more X also have more Y," but "in this person, X happened before Y." Order is what a staging model is made of, and it is the hardest of the four, for a reason Section 5 sets out.

A reversibility claim is a counterfactual. Not what is, but what would happen if something were restored. These are the claims that carry all the hope in the field, and they are the only kind that no measurement can settle, because they are not about an existing state of affairs.

The model this site has built makes claims of all four kinds. The instruments in routine use address the first. That mismatch is the subject of this paper.

3. What the Instruments Actually Reach

It is worth being concrete and fair about this, because the field's instruments are genuinely impressive and the point is not that they are bad. It is that they measure a different category of thing from the one the model is about.

Table 1 — The principal human read-outs, and what each actually reports.

Instrument	What it returns	Kind	What it does not report
Amyloid PET	Fibrillar amyloid burden in cortex	Level	Production rate, clearance rate, intraneuronal pool
Tau PET	Aggregated tau in a regional distribution	Level	When it arrived; whether it is still accumulating in this person
Structural MRI	Tissue volume remaining	Level	Whether loss is cells or neuropil; whether it is recoverable
FDG-PET	Regional glucose uptake	Level (of a rate proxy)	Which cell type; whether the deficit is demand or supply

Instrument	What it returns	Kind	What it does not report
High-resolution fMRI	Regional metabolic defect, e.g. lateral entorhinal cortex in preclinical disease (Khan et al., 2014)	Level	The molecular lesion producing it
SV2A PET	Synapse density *in vivo* (Mecca et al., 2020)	Level — of the model's own unit	Which synapses; whether the remaining ones work
CSF and plasma assays	Concentrations: amyloid ratios, phosphorylated tau species, neurofilament	Level	Compartment of origin; the rate of the process that produced them
EEG / MEG	Network activity, epileptiform discharges (Vossel et al., 2016)	Function, in real time	Which cells; the lesion behind it
Neuromelanin-sensitive MRI	Locus coeruleus signal intensity	Level	Transmitter output; whether axons or cell bodies are affected

Three things stand out from that table.

Almost every entry is a **level**. The one genuine exception is electrophysiology, which returns network function in real time — and it is, as Section 7 shows, the one place in this disease where a mechanism-level read-out has actually changed what a trial found.

The entries that bear on the model's *own* units are thin. For Phase I there is neuromelanin MRI, which reports how much pigmented tissue is there, not how much noradrenaline is being delivered. For Phase II there is no standard read-out of microglial restraint at all. For Phase III there is EEG. For the endosomal axis there is nothing in routine use; the candidates are a spinal-fluid panel linked to retromer traffic (Simoes et al., 2020), soluble SORL1 as a marker of variant impact (de Waal et al., 2026), and a reporter assay that so far works in cells rather than people (Fazeli et al., 2026). All early. **Established by absence.**

And one entry is better than the corpus usually admits. SV2A PET measures synapse density in living people, and synapse density is the unit the model chose. Section 8 returns to that, because it is the strongest methodological position the corpus holds and it was arrived at almost by accident.

4. How the Field Solved This: It Redefined the Disease

Faced with a disease whose mechanisms are invisible, the field made a decision that deserves more attention than it gets.

The research framework now in general use defines Alzheimer's disease by its underlying pathologic processes **as documented by biomarkers in living people**, grouping them into three: β-amyloid deposition, pathologic tau, and neurodegeneration — the scheme known as AT(N) (Jack et al., 2018). The framework's authors are explicit that this shifts the definition of the disease in living people from a syndrome to a biological construct, and equally explicit that it is for research rather than for clinical care.

They are also, to their considerable credit, explicit about the risk. The framework states plainly that it is possible the plaques and tangles are not causal, and that they are used because they are what distinguishes this disease from other causes of dementia.

Read that twice. The field's operative definition of the disease rests on three measurable markers which the definition's own authors decline to claim are causal. This is not a scandal. It is what a serious field does when the causal processes are out of reach: it standardises on what can be measured, says so, and gets on with the work. Without it, no two studies could be compared.

But the cost is specific and it lands squarely on work like ours. **None of the four processes our model is about has a place in AT(N).** Brainstem transmitter supply is not A, T or N. Microglial restraint is not A, T or N. Perineuronal net integrity is not A, T or N. Endosomal return is not A, T or N. Synapse loss is arguably the N, but N is usually operationalised as atrophy or neurofilament, not as synapse density.

So in the framework's own terms, a model built on those four things is not a model of Alzheimer's disease. It is a model of something that correlates with it. We think the framework is drawing the boundary in the wrong place — around what is visible rather than around what is happening — but we should be honest that we are the ones outside the boundary, and that the burden of showing the boundary is wrong falls on us.

Grade: Established that AT(N) is the operative definition and that our units sit outside it. **Inference**, and ours, that the boundary is drawn around the measurable rather than around the causal.

5. The Order Problem

This is the section that matters most, and the cleanest way to make the point is to take the field's two most celebrated timing results and look at how they were built.

5.1 The twenty-five-year curve

Almost everyone in this field can recite the dominantly-inherited timeline: spinal-fluid amyloid begins to fall about twenty-five years before expected symptom onset; amyloid deposition is detectable at fifteen years; raised spinal-fluid tau and brain atrophy at fifteen; reduced glucose metabolism and impaired episodic memory at ten; global cognitive impairment at five; diagnosis about three years after expected onset (Bateman et al., 2012).

It is a magnificent study and the numbers above are its real findings. But look at how it was done. One hundred and twenty-eight people with a dominant mutation were assessed, and each person's distance from expected onset was estimated from their parent's age at onset. The analyses of order were then performed **across** those people — a cross-sectional analysis arranged along an estimated timeline. The authors state this in the paper, and they state the consequence: the results require confirmation with longitudinal data, and may not apply to the sporadic disease.

Nobody was followed for twenty-five years. The curve is a composite of people at different distances from an event estimated from a relative. It is the best available evidence about order in this disease, and it is still not a

measurement of order in anybody.

5.2 The brains of the young

The corpus's own Phase I start date rests on a second famous study. Forty-two brains, from people aged four to twenty-nine who died of other causes, were examined for the proteins of this disease. Thirty-eight of the forty-two had abnormal tau somewhere. In nineteen of the twenty-two cases that had none in the cortex, the abnormal material was confined to the noradrenergic locus coeruleus complex (Braak and Del Tredici, 2011). A separate stereological series found that even at the lowest Braak stage, around eight per cent of locus coeruleus neurons already carry tau inclusions (Ehrenberg et al., 2017).

That is the empirical floor under "the disease begins in the third decade." And the original authors name the assumption it needs, in their own abstract: *assuming the pretangle alterations are not transient and do not regress*.

That assumption cannot be checked. Checking it would require seeing the same brainstem twice in the same person, decades apart, which is exactly what cannot be done. It may well be right. It is an assumption, it is load-bearing, and our model inherits it.

5.3 Why post-mortem series cannot settle order

The general point behind both examples is worth stating on its own, because it applies to most of the human evidence in this disease.

You see each brain **once**, at the **end**. From a set of such observations you can construct a plausible sequence by assuming that the people with less pathology are earlier versions of the people with more. That assumption is usually reasonable and it is never verified. It fails whenever there is more than one route to the endpoint, which in a disease this heterogeneous is likely.

The sample is also not a sample of the living. Autopsy series are made of people who died, often of something else, and who agreed to donate. Whatever makes a person reach ninety with a brain available for study may be exactly what protected them, or what killed them early; either way they are not a random draw from the population the model is about.

And a single observation cannot distinguish a slow process caught early from a fast process caught late. Two brains with equal pathology may have taken five years or thirty to get there. Nothing in the tissue says which.

Grade: Established. None of this is controversial; it is simply not often said out loud next to a staging diagram.

6. The Price of Our Own Redraw

Honesty requires a section that cuts against us, and this is it.

The model was first drawn as a relay: Phase I hands over to Phase II, which hands over to Phase III. It is now drawn as stacked steps, because that is more faithful to the biology — the brainstem keeps failing underneath everything that follows, and a person of seventy-five is carrying all three phases at once.

We think the redraw is right. It is also **harder to check**, and that cost should be entered on the ledger rather than enjoyed quietly.

A relay model makes a hard prediction: there are hand-over moments, and at a hand-over something should change measurably. A missed hand-over would falsify it. A stacked model makes no such prediction. Every process is running in everybody over sixty, at different intensities, and a measurement taken at seventy-five sees the superposition. To decompose that into three contributions you would need either a read-out specific to each layer — which Section 3 says does not exist for three of the four — or the same person measured repeatedly across decades, which Section 5 says nobody has.

So the redraw trades falsifiability for fidelity. That is often the right trade in biology, and it is the trade we made. But a model that fits everything and forbids nothing is not a scientific advance, and the way to keep this one honest is to insist that each layer carry its own refutation conditions even though the stack as a whole is hard to test. That is what the per-paper ledgers are for.

Grade: Inference, and ours. We know of no formal treatment of the testability cost of stacked versus sequential staging models in this disease.

7. The One Case Where a Read-Out Changed the Answer

It would be easy to end in gloom. There is a case that shows what the other direction looks like, and it is worth more attention than it has had.

Network hyperexcitability is a Phase III mechanism. It has something almost nothing else in this disease has: a direct, cheap, real-time human read-out. Overnight video-EEG and magnetoencephalography detect subclinical epileptiform activity, and in a prospective series it was present in 42.4 per cent of Alzheimer patients against 10.5 per cent of controls. The patients with it were clinically indistinguishable at the time of monitoring, and then declined more than twice as fast — 3.9 points a year on a global cognitive measure against 1.6 (Vossel et al., 2016).

Then the read-out was used to choose. A randomised crossover trial gave levetiracetam, an antiseizure drug, to thirty-four people with Alzheimer's disease. **The primary outcome was negative**: across the whole group, the drug did nothing measurable to executive function. But in a prespecified analysis of the subgroup with epileptiform activity — thirteen of the thirty-four — performance improved on an interference task and on a spatial route-learning task (Vossel et al., 2021).

This is a small, exploratory result in a phase 2a trial and it should not be oversold; the authors call for further assessment rather than for practice change. What makes it important here is structural rather than clinical. A mechanism-level read-out existed. It identified the third of patients in whom the mechanism was actually operating. Those patients responded; the rest diluted the signal to nothing. **Without the read-out, the honest conclusion from that trial would have been that the mechanism does not matter.**

That is the shape of the problem across this whole disease. The field has run many trials on diluted populations and concluded that mechanisms do not matter. Some of those conclusions are probably wrong in exactly this way, and there is no way to know which, because the read-outs that would have stratified them did not exist.

Grade: Established for the measurements. **Probable** for the drug effect, which is a single small trial with a negative primary outcome and a positive prespecified subgroup. **Inference**, ours, that other negative trials in this disease suffer the same dilution.

8. The Best Decision the Corpus Made

One more constructive section, and it concerns a choice that was made for biological reasons and turned out to be the corpus's strongest methodological position.

The model measures function in **connections**, not in cells. The reason given is biological: synapse loss is the structural change that tracks cognitive decline most closely — better than plaque burden, better than tangle burden — established across three decades (DeKosky and Scheff, 1990; Terry et al., 1991).

The methodological consequence was not the reason for the choice, but it is substantial. Synapse density is measurable in living people. SV2A PET does it, and has demonstrated widespread synaptic loss *in vivo* (Mecca et al., 2020). Of the model's various units, the one it elected to treat as primary is the one with a working human instrument.

This is worth saying plainly because it is unusual. A great deal of mechanism in this disease is defined in terms that could only ever be checked in tissue. The corpus's central quantity is not. That means the model's headline claim — that what matters is connections lost, and that losing them is what produces the symptom — is in principle checkable in living people, and partly has been.

What SV2A PET does not do is say *which* synapses, or whether the ones remaining are working, or whether a lost one could come back. It is a level, like almost everything else. But it is a level of the right thing, which is more than most of the field's instruments can say.

Grade: Established.

9. Three Tiers: What Is Closable, What Is Reachable Only by Inference, What Is Closed

Not all of the gaps described here are the same kind of gap, and conflating them produces either false despair or false promises.

Tier one — closable with money and effort. These are missing assays, not impossibilities. A validated human read-out of endosomal recycling. A microglial state marker specific enough to say which restraints have been released. A measure of noradrenergic *output* rather than of pigment. Longer longitudinal cohorts with repeated sampling. None of these requires new physics; they require someone to decide they matter. The spinal-fluid and reporter work cited in Section 3 is the beginning of the first.

A worked example, and the cheapest one we know of. One entry stated concretely is worth more than a list, and the clearest comes from Estela Area-Gómez's programme. She proposes that the informative quantity is not the level of amyloid-β but a *ratio*: total amyloid-β to C99, the uncut fragment it is cut from, measured in white blood cells from an ordinary blood draw (Montesinos et al., 2020). A ratio of product to substrate reports whether the cutting enzyme is keeping up with its workload — a rate, read from a single sample, using two markers out of one tube. The reason it is not routine is a reagent problem the programme itself identified: many antibodies raised against amyloid-β cannot distinguish it from the larger fragment that contains it.

That example also shows the tiers are properties of steps rather than of projects. The *assay* is Tier one: a solvable antibody problem, no new physics. The *interpretation* — reading a rate off a static ratio — belongs to Tier two, and is only as good as the model licensing it.

We record it as an inventory item rather than as a recommendation. It has not been run at scale, the ratio's behaviour across the staging series is unknown, and a read-out that would be informative if the substrate reversal is right may be less so if it is not. **Inference**, ours, that this is the cheapest unbuilt read-out in the corpus; the proposal itself is the programme's, and was stated

in advance of the data.

Tier two — reachable only by inference from a chain. Some quantities will never be observed directly in a living brain but can be constrained by measuring their consequences, if the chain from mechanism to consequence is specified tightly enough. Rates are mostly here: you cannot watch a compartment empty in a living person, but a well-chosen pair of markers whose ratio depends on the rate can bound it. This is the territory *The Unmeasured Rate* works through in detail, and it is harder than tier one because the inference is only as good as the model licensing it — which means the model is being used to measure the thing that is supposed to test the model. That circularity is manageable but it has to be declared.

Tier three — closed to measurement in principle. Reversibility lives here, and it is the tier that matters most.

The claim that a silenced neuron might be recovered is not a claim about any present state of affairs. It is a claim about what would happen under an intervention that has not been performed. No scan can settle it, because there is nothing yet to scan. The only way to find out whether silence is recoverable is to restore the missing function in a person and see whether anything comes back.

Which means that for the most consequential claim in the model, **measurement and therapy are the same act**. You cannot measure first and treat second. The trial *is* the measurement.

That is an uncomfortable conclusion and it should not be softened. It means the model's hopeful claim cannot be de-risked in advance by better instruments, however many are built. It can only be tested by someone being brave enough, and well-enough funded, to try the intervention and watch.

Grade: Inference, and ours, for the three-tier division. **Established** that no present instrument addresses reversibility.

10. Why This Corpus Grades Everything

A reader who has got this far may notice that the four-grade ledger at the end of every paper in this corpus is not a stylistic tic. It is the direct consequence of everything above.

If the subject were observable, grading would be unnecessary; one would simply report measurements. Because most of it is not observable, nearly every claim in the model rests on something other than a direct human measurement — on an animal in which the lesion was induced, on a dish without age or circuitry, on a post-mortem series seen once at the end, or on an inference drawn across several of these. A framework built on such material can be respectable or it can be assertion, and the only thing that separates them is whether each claim carries, visibly, the kind of evidence it actually rests on and the result that would overturn it.

That is what the grades do. **Established** means human material, replicated. **Probable** means a real gap named. **Inference** means we joined things up and nobody has tested the join. **Speculative** means we find it coherent and that is all.

It is also why this corpus states its own errors in print rather than correcting them silently, and why claims that are ours rather than the literature's are marked as ours in every ledger. In a field where the primary objects cannot be seen, the discipline that substitutes for observation is declaring exactly what you are standing on.

11. What Follows

Four things follow, in descending order of how confident we are about them.

Instrumentation is the rate-limiting step, not biology. This disease has no shortage of plausible mechanisms; it has a shortage of ways to tell, in a person, which mechanism is running. On the evidence of Section 7, the cost of that shortage is not merely slow progress but wrong conclusions — trials that correctly measured a diluted population and were read as refuting a mechanism. **Inference**, ours, stated at that strength.

Read-outs should be funded as objects in their own right, not as appendages to drug programmes. A drug programme needs its read-out to

work by the time the trial starts; a read-out that takes a decade will therefore never be built inside one. The two cases in this paper where a read-out existed — EEG for hyperexcitability, SV2A PET for synapses — were both built by people who were not trying to qualify a particular compound.

Stratify before concluding. Where a read-out exists, a negative primary outcome in an unstratified population should be reported as uninformative about mechanism rather than as evidence against it. This is standard practice in oncology and remains unusual here.

And be exact about which tier a gap is in. Promising that a better scanner will answer the reversibility question is not true and will eventually be found out. The honest statement is that tier one is a funding problem, tier two is a modelling problem, and tier three can only be answered by treating someone.

12. Graded Ledger

#	Claim	Grade	Resting on	What would overturn it
1	Nearly all human read-outs in this disease return levels, not rates	Established	Table 1; the instruments' own specifications	A deployed human assay returning a rate
2	The field's operative definition (AT(N)) rests on markers its authors decline to call causal	Established	Jack 2018, stated in the framework itself	—
3	None of the model's four units has a standard read-out within AT(N)	Established (by absence)	Table 1; Section 4	A validated marker of brainstem output, microglial restraint, net integrity or endosomal return entering routine use

#	Claim	Grade	Resting on	What would overturn it
4	The dominantly-inherited timeline is a cross-sectional construction, not a longitudinal observation	Established	Bateman 2012, stated by its authors	Longitudinal confirmation across the same individuals
5	The Phase I start date rests on an explicit, uncheckable non-regression assumption	Established	Braak and Del Tredici 2011, stated in the paper	Evidence that early brainstem tau regresses or is transient
6	Post-mortem series cannot establish order within an individual	Established	General; Section 5.3	—
7	A mechanism-level read-out changed what one trial in this disease found	Probable (small, exploratory)	Vossel 2016; Vossel 2021, negative primary outcome with positive prespecified subgroup	Failure to replicate the stratified effect
8	Synapse density, the model's chosen unit, is measurable in living people	Established	Mecca 2020; DeKosky 1990; Terry 1991	—
9	The stacked redraw is less falsifiable than the relay it replaced	**Inference (ours, untested)**	Section 6	A layer-specific read-out that decomposes a single cross-sectional measurement
10	Other negative trials in this disease are diluted in the way the levetiracetam trial was	**Inference (ours, untested)**	Section 7 read across the trial record	Re-analysis of a negative trial with a mechanism read-out showing no responsive stratum

#	Claim	Grade	Resting on	What would overturn it
11	Reversibility is closed to measurement in principle and demonstrable only by treating	**Inference (ours)**	Section 9	An instrument that reports the recoverability of a silenced neuron

13. What Would Change This Account

A deployed human assay that returns a rate. This is the single development that would most change the paper, and it would move a large number of claims across the corpus from inference to testable. *The Unmeasured Rate* specifies what one would have to do for a single case.

A longitudinal cohort of sufficient duration. Not five years; thirty. Expensive, unglamorous, and the only thing that would convert the order claims in this corpus from constructions into observations. Nothing else does it.

A negative trial re-read with a read-out. If a mechanism trial that failed in an unstratified population were re-analysed with a marker of that mechanism and showed no responsive subgroup, claim 10 above would weaken considerably and the argument of Section 7 would lose its generality.

A layer-specific marker. Anything that reports on one stacked layer independently of the others would restore to the model the falsifiability Section 6 admits it gave up.

Companion papers in this corpus

The Unmeasured Rate — why interventions at the autophagic bottleneck have failed, and what a human assay of flux would have to measure. The rate problem, worked through one case. *The Biomarker Cascade* — which biomarkers detect which phase, and why current panels miss the first two. *What the Brainstem Stops Sending*, *The Hands That Let Go*, *The Cage and the Clock*, *The Long Way Back* — the four short papers whose claims this one asks after.

References

Bateman RJ, Xiong C, Benzinger TLS, Fagan AM, Goate A, Fox NC, Marcus DS, Cairns NJ, Xie X, Blazey TM, Holtzman DM, Santacruz A, Buckles V, Oliver A, Moulder K, Aisen PS, Ghetti B, Klunk WE, McDade E, Martins RN, Masters CL, Mayeux R, Ringman JM, Rossor MN, Schofield PR, Sperling RA, Salloway S, Morris JC. Clinical and biomarker changes in dominantly inherited Alzheimer's disease. *New England Journal of Medicine*. 2012;367(9):795-804. PMID 22784036.

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

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

de Waal MWJ, van der Lee SJ, Lunding M, et al., Andersen OM, Teunissen CE, Holstege H. Soluble SORL1 in cerebrospinal fluid as a marker for functional impact of rare SORL1 variants. *Alzheimer's and Dementia*. 2026;22(2):e71042. PMID 41685556.

Ehrenberg AJ, Nguy AK, Theofilas P, Dunlop S, Suemoto CK, Di Lorenzo Alho AT, Leite RP, Diehl Rodriguez R, Mejia MB, Rub U, Farfel JM, de Lucena Ferretti-Rebustini RE, Nascimento CF, Nitrini R, Pasquallucci CA, Jacob-Filho W, Miller B, Seeley WW, Heinsen H, Grinberg LT. Quantifying the accretion of hyperphosphorylated tau in the locus coeruleus and dorsal raphe nucleus: the pathological building blocks of early Alzheimer's disease. *Neuropathology and Applied Neurobiology*. 2017;43(5):393-408. PMID 28117917.

Fazeli E, Jan A, Huber AK, Jensen AMG, Fazeli E, Klein J, Merchant K, Andersen OM. A luminescent reporter assay to quantify SORL1 ectodomain shedding and retromer-dependent endosome recycling activity. *Journal of Biological Chemistry*. 2026;302(2):111136. PMID 41519200.

Jack CR Jr, Bennett DA, Blennow K, Carrillo MC, Dunn B, Haeberlein SB, Holtzman DM, Jagust W, Jessen F, Karlawish J, Liu E, Molinuevo JL, Montine T, Phelps C, Rankin KP, Rowe CC, Scheltens P, Siemers E, Snyder HM, Sperling R. NIA-AA Research Framework: toward a biological definition of Alzheimer's disease. *Alzheimer's and Dementia*. 2018;14(4):535-562. PMID 29653606.

Khan UA, Liu L, Provenzano FA, Berman DE, Profaci CP, Sloan R, Mayeux R, Duff KE, Small SA. Molecular drivers and cortical spread of lateral entorhinal cortex dysfunction in preclinical Alzheimer's disease. *Nature Neuroscience*. 2014;17(2):304-311. PMID 24362760.

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

Montesinos J, Pera M, Larrea D, Guardia-Laguarta C, Agrawal RR, Velasco KR, Yun TD, Stavrovskaya IG, Xu Y, Koo SY, Snead AM, Sproul AA, Area-Gómez E. The Alzheimer's disease-associated C99 fragment of APP regulates cellular cholesterol trafficking. *EMBO Journal*. 2020;39(20):e103791. PMID 32865299.

Simoes S, Neufeld JL, Triana-Baltzer G, Moughadam S, Chen EI, Kothiya M, Qureshi YH, Patel V, Honig LS, Kolb H, Small SA. Tau and other proteins found in Alzheimer's disease spinal fluid are linked to retromer-mediated endosomal traffic in mice and humans. *Science Translational Medicine*. 2020;12(571):eaba6334. PMID 33239387.

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

Vossel KA, Ranasinghe KG, Beagle AJ, Mizuiri D, Honma SM, Dowling AF, Darwish SM, Van Berlo V, Barnes DE, Mantle M, Karydas AM, Coppola G, Roberson ED, Miller BL, Garcia PA, Kirsch HE, Mucke L, Nagarajan SS. Incidence and impact of subclinical epileptiform activity in Alzheimer's disease. *Annals of Neurology*. 2016;80(6):858-870. PMID 27696483.

Vossel K, Ranasinghe KG, Beagle AJ, La A, Ah Pook K, Castro M, Mizuiri D, Honma SM, Venkateswaran N, Koestler M, Zhang W, Mucke L, Howell MJ, Possin KL, Kramer JH, Boxer AL, Miller BL, Nagarajan SS, Kirsch HE. Effect of levetiracetam on cognition in patients with Alzheimer disease with and without epileptiform activity: a randomized clinical trial. *JAMA Neurology*. 2021;78(11):1345-1354. PMID 34570177.