
The Unmeasured Rate

*Why interventions at the autophagic bottleneck in Alzheimer's disease have failed, and
what a human assay of autophagic flux would have to measure*

Benjamin Aaron Gustafsson

AdultCognitiveDisease.com

September 2026

Eleven completed human trials have given drugs to patients with Alzheimer's disease on the grounds that those drugs correct the cell's disposal system. Not one of them measured the disposal system. The reason is not negligence: autophagic flux is a rate, every human readout is a level, and the assay that converts the second into the first requires a poison, a transgene, or the brain.

Contents

-
1. Introduction: A Mechanism Without a Dial
2. What Is Established About the Bottleneck
The scale of the asymmetry
3. The Direction of the Lesion: Induction Intact, Completion Failing
4. The Trial Record I — mTOR and the Rapalogues
5. The Trial Record II — AMPK and Metformin
6. The Trial Record III — TFEB, the Transcription Factor Without a Drug
7. The Trial Record IV — Trehalose, and the Largest Clean Test Anyone Has Run
8. The Trial Record V — Sigma-1, Phosphodiesterase-5, and the Autophagy Rationale at Its Most Elastic
9. The Trial Record VI — The Repurposing Wave
An aside on what a repurposing rationale is worth without a readout
10. What the Record Has in Common
11. Flux Is a Rate
A worked case
12. Every Human Readout Is a Level
The instrument the field did build
13. Three Ways a Level Misleads
14. Why Imaging Has Not Closed the Gap
15. The Preclinical Assays That Do Not Travel
16. A Target Where the Readout Exists: NLRP3
The two targets are not merely analogous; they are connected
17. What a Movable Readout Bought
18. The Limits of the Contrast
19. The Deeper Problem: Four Results from 2025 and 2026
20. Three Kinds of Biomarker, and Which One Autophagy Lacks
21. Specification: What a Deployable Human Flux Assay Must Measure
22. Four Candidate Architectures
22.1 Metabolic labelling of an autophagy-obligate cargo
22.2 A bounded pharmacological clamp
22.3 A peripheral cell challenge assay
22.4 An imaging readout of the degradative step
22.5 Ranking
23. How a Flux-Dosed Trial Would Differ
A concrete first study
24. Predictions
25. Strength of Evidence
26. Limitations
27. Conclusion
- References

Abstract

The endosomal–lysosomal–autophagic network is among the best-evidenced lesions in Alzheimer’s disease. Enlargement of the early endosome is the earliest disease-specific structural abnormality known in human tissue [1]; autophagic vacuoles are the dominant organelle inside the dystrophic neurites that define the neuritic plaque [2]; the terminal degradative step fails while autophagy induction is preserved or increased [3]; and common-variant genetics places endocytosis and lysosomal function among the pathways most strongly enriched across the disease’s risk loci [4]. In animals, correcting the lesion works. Restoring lysosomal proteolysis reverses amyloid pathology and memory deficits [5]; inducing autophagy pharmacologically abolishes cognitive deficits [6]; over-expressing the lysosomal master regulator TFEB clears aberrant tau and rescues neurotoxicity [7].

In human beings, the record is different, and the difference is not one of degree. This paper audits it. Across the compounds whose rationale in Alzheimer’s disease is autophagic or endolysosomal — rapamycin, metformin, trehalose, sigma-1 and phosphodiesterase-5 agents, nilotinib, lithium, neflamapimod — we find a uniform feature that has not, to our knowledge, been stated as a general property of the field: **not one completed trial measured autophagy in the people it treated.** Trials measured drug concentration in plasma or cerebrospinal fluid; they measured amyloid, phosphorylated tau, glial fibrillary acidic protein, hippocampal volume, cognition. None measured the process the drug was given to correct. In the one trial that set out to establish central nervous system penetrance of rapamycin in early Alzheimer’s disease, the drug was undetectable in cerebrospinal fluid before and after eight weeks of treatment, while phosphorylated tau-181 and glial fibrillary acidic protein rose [8].

We argue that this is not negligence but a consequence of a specific and under-acknowledged measurement problem. **Autophagic flux is a rate. Every available human readout is a level.** The steady-state abundance of an autophagosome marker is the quotient of two rates — the rate at which compartments are formed and the rate at which they are consumed — and a single level is therefore consistent with an unbounded family of rate pairs. Identifying the rate requires observing the system in two states, which in every validated laboratory assay means adding a lysosomal inhibitor. The measurement that defines flux requires administering a poison to the compartment being measured. That is why the assay has not crossed into humans, and why every human autophagy readout now in the literature — cerebrospinal fluid LC3A and LC3B [9], PINK1, BNIP3L and TFEB in serum and cerebrospinal fluid [10], ULK1 [11], lysosomal network proteins [12], p62 [13] — is a concentration, unperturbed, of unknown cellular provenance.

We contrast this with a target of comparable intracellular inaccessibility where the readout does exist. The NLRP3 inflammasome is, like autophagy, a cytoplasmic process in glia with no direct human assay; unlike autophagy, its output is a secreted cytokine. A brain-penetrant NLRP3

inhibitor, NT-0796 (ruvonoflast), achieves cerebrospinal fluid exposure and lowers cerebrospinal fluid interleukin- 1β alongside systemic C-reactive protein, fibrinogen, interleukin-6 and interleukin-18 in patients with Parkinson's disease [14], with the brain activation step subsequently traced to carboxylesterase-1 in microglia [15] and the anti-inflammatory effect confirmed against placebo in a randomised phase 1b [16]. The contrast is instructive and its limits must be stated: a movable readout has bought dose selection against effect rather than against exposure, and nothing yet beyond that.

That qualification opens the paper's second argument. Four results from 2025 and 2026 show that a movable biomarker is not sufficient. Valacyclovir not only failed to slow decline in herpes-seropositive early Alzheimer's disease but was associated with significantly *greater* cognitive worsening than placebo [17]; oral semaglutide failed on the Clinical Dementia Rating–Sum of Boxes across two phase 3 trials totalling 3,808 participants [18]; and, in results disclosed at conference and by company release rather than in the peer-reviewed literature, the anti-p217-tau antibody posdinemab abolished cerebrospinal fluid p-tau217 without altering decline, while the tau antisense oligonucleotide diranersen reduced tau pathology at every dose but missed its dose-response primary endpoint. We treat these as evidentially distinct from the autophagy failures, and the distinction is the paper's hinge. Posdinemab's failure is *informative*: it tells us that the quantity moved was not on the causal path, or that the drug reached the wrong pool. A rapamycin trial with no autophagy measurement is *uninformative*: it cannot distinguish a wrong hypothesis from an unengaged target. The field has been producing the second kind of failure and reading it as the first.

We close constructively. A deployable human flux assay must satisfy six requirements, of which the hardest is that it report a rate without requiring a toxic clamp. We set out four candidate architectures and argue that one of them is already assembled from published parts that have not been joined: neurofilament subunits are degraded by autophagy in vivo, with levels rising when autophagy is blocked and falling under rapamycin [19], and the turnover of neurofilament light in the human central nervous system has now been measured directly by stable-isotope labelling kinetics [20]. A labelling study of neurofilament turnover, performed on and off an autophagy-modulating drug, would yield a rate, in brain, for a protein whose degradation is autophagy-dependent. We specify how a trial dosed against such a readout would differ from those that have failed, give nine predictions that would falsify the account, and grade every claim.

1. Introduction: A Mechanism Without a Dial

There is an asymmetry in Alzheimer's therapeutics that is easy to state and hard to look at directly. The cell-biological lesion with the strongest claim to being early, human, and disease-specific is a failure of the disposal system — and it is the lesion against which the fewest interpretable human experiments have been performed.

The asymmetry is not for want of mechanism. The endosomal–lysosomal–autophagic network in Alzheimer's disease has been characterised for three and a half decades in human post-mortem and biopsy tissue, in isogenic human neurons, in a dozen animal models and at the resolution of the individual organelle. It has produced findings of a kind the rest of the field would envy: a structural abnormality visible decades before the defining lesion of the disease [1]; an organelle census inside the neuritic plaque's dystrophic neurites showing that they consist almost entirely of stalled autophagic compartments [2]; a direction of failure established in laser-captured human CA1 neurons, in which autophagy induction is *up* while clearance is *down* [3]; a molecular agent, the β -cleaved carboxy-terminal fragment of the amyloid precursor protein, identified independently by a laboratory with no stake in the account [21]; and a mechanism by which that fragment, once phosphorylated, binds and inhibits the vacuolar H⁺-ATPase [22]. Genome-wide association places endocytosis and lysosomal function among the most strongly enriched pathways across the disease's risk architecture [4]. Reviews of autophagy in human disease treat neurodegeneration as one of its clearest applications [23], and the case for it as a therapeutic target has been made at length [24].

Nor is it for want of preclinical therapeutic success. The list of manipulations that correct the lesion and improve the animal is long and the effect sizes are large. Genetic deletion of cystatin B, an endogenous inhibitor of lysosomal cysteine proteases, restores autophagic–lysosomal function, reduces amyloid pathology and improves memory in a transgenic model [5], with the same intervention subsequently shown to correct defective macroautophagic turnover of brain lipids [25]. Rapamycin abolishes cognitive deficits and reduces amyloid- β [6], an effect since reproduced widely enough to support a formal meta-analysis of preclinical studies [26]. Neuronal over-expression of TFEB, the transcription factor that coordinates lysosomal biogenesis [27] and links it to autophagy [28] through a shared promoter motif network [29], selectively clears aberrant tau and rescues neurotoxicity [7], accelerates lysosomal degradation of the amyloid precursor protein [30], reduces the intraneuronal C99 fragment [31], and — expressed in astrocytes instead — facilitates amyloid- β clearance and attenuates plaque pathogenesis [32]. Restoring mitophagy inhibits both amyloid and tau pathology and reverses cognitive deficits [33], an effect reproduced with urolithin A [34]. Acidic

nanoparticles restore lysosomal acidification defects [35]. A pharmacological chaperone stabilising the retromer limits amyloidogenic processing [36]. Raising ULK1, the initiating kinase of macroautophagy, stimulates autophagic flux, reduces pathology and delays cognitive decline [11]. Basal autophagy is not optional in neurons: its suppression alone causes neurodegeneration in mice [37].

What is missing is the middle term. Between a mechanism this well-specified and a clinical programme this thin lies a gap that is usually attributed to the ordinary difficulty of central nervous system drug development — blood–brain barrier penetration, chronic dosing, the length and expense of dementia trials. Those difficulties are real. They are not, we will argue, the binding constraint here.

The binding constraint is that nobody can measure the thing.

This paper makes three claims. The first is empirical and, we believe, new as a general statement: across the clinical record of autophagy-rationale compounds in Alzheimer’s disease, *no completed trial has measured autophagy in its participants*. We establish this drug by drug in Sections 4 to 9 and tabulate it in Section 10. The second claim is conceptual and explains the first: autophagic flux is a rate, every available human readout is a level, and the transformation from level to rate requires a perturbation that cannot ethically be applied to a human brain. We develop this in Sections 11 to 15. The third claim is constructive: the requirement is not in principle unsatisfiable, and one satisfying architecture can be assembled from published human methods that have not been combined. We specify it in Sections 21 to 23.

Section 16 to 18 supply a control. If the argument is that a missing readout explains a missing translation, then a comparably inaccessible target *with* a readout should have a different trajectory. The NLRP3 inflammasome is that control, and the comparison is informative in both directions — for what a readout buys, and for what it does not.

Sections 19 and 20 address the strongest objection to the whole enterprise. Alzheimer’s disease has lately produced several trials in which a biomarker moved decisively and the patients did not. If biomarkers do not predict benefit, why should the absence of one matter? The answer is that these are different epistemic situations, and conflating them is the error the field has been making. We argue that an uninformative failure is worse than an informative one, and that the autophagy programme has been generating the worse kind.

2. What Is Established About the Bottleneck

Before auditing the therapeutic record it is necessary to fix what that record was attempting to treat, and at what level of confidence. This section states the propositions we take as established, because they are the ones the trials assumed.

The compartment is abnormal, and abnormal early. Enlargement of the neuronal early endosome, with increased delivery of hydrolases to those compartments, is present in sporadic Alzheimer's disease before amyloid deposition; it is present in Down syndrome decades before the associated dementia; it is accentuated by the *APOE* $\epsilon 4$ allele; and it was not found in the comparison neurodegenerative diseases examined alongside it [1]. A pathological change visible in a brain long before the disease it is associated with is not plausibly a consequence of that disease. Among the structural abnormalities in Alzheimer's disease this has the best claim to primacy in time.

The agent is a fragment of the amyloid precursor protein, not the peptide. Working in a large isogenic panel of human induced-pluripotent-stem-cell neurons carrying familial *APP* and *PSEN1* mutations against a common background, one group established that the shared endosomal abnormalities are mediated by β -carboxy-terminal fragments rather than by amyloid- β [21]. (A published erratum was subsequently issued for this paper [38]; we note it and have relied on the version of record.) The fragment mechanism was extended downstream when phosphorylation of the fragment at tyrosine 682 was shown to produce a species that binds subunits of the vacuolar H^+ -ATPase and inhibits the pump, so that one molecule produces lesions at two compartments [22]. The sufficiency of the endosomal node has been tested directly: neuronal over-expression of the Rab5 effector adaptor APPL1 alone, with no manipulation of the amyloid precursor protein, replicates Alzheimer-related endosomal and synaptic dysfunction with cholinergic neurodegeneration [39].

The terminal degradative step is a contested node with a long history. Presenilin 1 was reported in 2010 to be required, independently of γ -secretase activity, for maturation and lysosomal targeting of the V0a1 subunit of the pump, so that loss of presenilin function alkalinises the lysosome and blocks autophagy [40]. The claim was contradicted within two years by a group reporting frankly normal lysosomal acidification in cells lacking both presenilins [41], and it has never been independently replicated in the form originally stated. We record this as unsettled, and note that the metabolic route through the phosphorylated fragment [22] supplies an alternative path to the same terminal lesion that does not require presenilin at all.

Autophagic vacuoles are the content of the neuritic plaque’s neurites. Immunogold labelling with compartmental markers and electron microscopy, performed on antemortem human neocortical biopsy tissue, established that autophagosomes, multivesicular bodies, multilamellar bodies and cathepsin-containing autolysosomes are the predominant organelles inside the dystrophic neurites that give the neuritic plaque its name [2]. This is a statement about the lesion by which the disease is neuropathologically staged.

The direction of failure is known, and it is the end of the pathway. This is the proposition that governs everything in the rest of this paper, and Section 3 is devoted to it.

The genetics point at the organelle. Pathway analysis of the largest common-variant association study of the disease places endocytosis and lysosomal function among the pathways most strongly enriched, alongside immunity and lipid metabolism [4]. More recently, genome-wide association analysis of cerebrospinal fluid biomarkers has identified loci regulating autophagy directly [42]. We attach the usual caution to pathway enrichment — it is a statement about annotation databases as much as about brains — but the signal is not marginal and it is not new.

In model systems, correcting the lesion works. Section 1 listed the evidence. We restate the form of the claim precisely, because the trials were built on it: in rodents, interventions at four distinct nodes of this network — protease inhibition relieved [5], mTOR inhibited [6], TFEB raised [7,30,32], ULK1 raised [11] — each improve both pathology and behaviour. This is an unusually consistent preclinical portfolio. It is also, in every case, a portfolio built in animals in which the independent variable was set by construction.

The scale of the asymmetry

It is worth putting a number, however rough, on the imbalance this paper describes, because the qualitative statement “preclinically saturated, clinically thin” is made about many targets and is usually an exaggeration.

The preclinical side does not need quantifying by us; the rodent literature on rapamycin alone in Alzheimer models is large enough to have supported a systematic review and meta-analysis of cognitive outcomes across studies [26], which is a threshold most therapeutic hypotheses in this disease never reach.

The clinical side is the striking number. The annual census of the Alzheimer’s disease drug development pipeline catalogues agents in active trials by mechanism and phase [43]. Against that inventory, the count of *completed* trials in Alzheimer’s disease with an explicit autophagic or endolysosomal rationale is the eleven studies of Table 1 — and the largest of them, excluding the sigma-1 programme whose autophagy claim is made only in its mechanism sentence, enrolled 168 participants [44], a trial completed in 2001 of a drug that *raises* lysosomal pH. Set beside the 3,808 participants randomised to a single GLP-1 hypothesis [18], or the 523 randomised to a single tau

antibody, the disparity is not a matter of the target being early in development. It is a target that has been mechanistically mature for two decades and has never had a properly powered, properly instrumented human test.

The obvious explanation — that the compounds are unavailable — does not survive inspection. Rapamycin, metformin, trehalose, lithium, hydroxychloroquine and nilotinib are all approved drugs with decades of human safety data. The compounds have been available the whole time. What has not been available is any way to tell whether giving them did the thing they were given to do.

3. The Direction of the Lesion: Induction Intact, Completion Failing

The single most important fact for therapeutic design in this area is not that autophagy is impaired in Alzheimer's disease. It is *where* it is impaired.

The definitive human measurement was made in laser-captured CA1 pyramidal neurons from early- and late-stage Alzheimer subjects and non-demented controls, analysed on a custom microarray of 578 neuropathology- and neuroscience-associated genes and complemented by protein and morphological readouts [3]. The result has three parts. First, autophagy-related genes are strikingly upregulated — more so than any other gene ontology group — beginning at early disease stages, with elevated expression of autophagosome components and increased LC3-positive puncta. Second, lysosomal biogenesis is likewise increased, with activation of the MiTF/TFE family of transcriptional regulators, and — a detail we will return to — with TFEB activation associated more strongly with glia than with neurons. Third, and despite this early compensatory response, substrate clearance progressively fails, reflected in autolysosomal accumulation of LC3-II and p62/SQSTM1 and in expansion of autolysosomal size and total area. Autophagosome-lysosome fusion was not evidently altered. The authors' conclusion is that sustained induction in the face of declining lysosomal clearance is what produces the disease's characteristically florid autophagic pathology.

This has a therapeutic corollary that is straightforward to state and has been almost universally ignored in practice. **If induction is already elevated and completion is what fails, then a drug that induces autophagy is acting on the step that is working.** In the best case it does nothing. In a plausible case it fills a neuron faster with compartments it cannot empty.

The direction has independent support in mouse work using a tandem-fluorophore LC3 reporter, in which one fluorophore is quenched at acidic pH and the other is not, so that the acidification state of every autophagic compartment can be read by colour in situ. In two amyloid precursor protein models, autolysosomal acidification declined in neurons well before extracellular amyloid deposition, in step with a measured fall in vacuolar-ATPase activity, and amyloid species then accumulated selectively within the enlarged, de-acidified compartments [45,46]. It has support from the opposite direction as well: blocking autophagosome formation by deleting *Atg7* in the neurons of an amyloid model markedly reduced extracellular amyloid deposition while raising intracellular amyloid and accelerating neurodegeneration [47]. Removing the route by which the deposit forms made the neuron worse.

One qualification belongs here and is developed in Section 21, because it complicates the naive prescription that follows from the above. “Restore acidification” is not the same instruction as “lower lysosomal pH.” Lysosomal pH is a regulated set-point held between opposing conductances, and recent reviews of neuronal lysosomal pH argue that dynamic instability in either direction is pathogenic rather than that low pH is protective [48]; the microglial case has been reviewed separately [49]. A therapy directed at the terminal step is a therapy aimed at a set-point, and the assay that would guide it must therefore report deviation in either direction. This is an additional reason that a readout is indispensable, not an argument against acting at the terminal step.

With the direction fixed, we turn to what was actually done in people.

4. The Trial Record I — mTOR and the Rapalogues

Rapamycin is the compound with the strongest preclinical case and the weakest clinical instrumentation, and the gap between the two is a microcosm of the paper's argument.

The preclinical case is old and robust. Inhibition of mTOR by rapamycin abolishes cognitive deficits and reduces amyloid- β levels in a mouse model [6]; the molecular interplay between mTOR, amyloid- β and tau, and its effect on cognitive impairment, was worked out in parallel [50]; and the animal literature is now large enough to have been meta-analysed, with a pooled effect on cognitive deficits across models [26]. The mechanistic rationale is that mTOR inhibition de-represses autophagy initiation and, through TFEB nuclear translocation, raises lysosomal biogenesis [28].

The clinical record consists, as of this writing, of two small open-label studies, and neither measured autophagy.

The first was a single-site open-label phase 1 trial in ten participants aged 55–85 with mild cognitive impairment or early-stage dementia (Global Clinical Dementia Rating 0.5–1.0), given rapamycin 1 mg/day for eight weeks [8]. Its primary aim was to evaluate central nervous system penetrance by assaying drug levels in cerebrospinal fluid before and after treatment. The finding was that **rapamycin was not detectable in cerebrospinal fluid at either time point**. Twenty mostly mild adverse events occurred; systolic blood pressure and haemoglobin A1c rose; multiple erythrocyte parameters fell; there were no significant cognitive changes; and cerebrospinal fluid phosphorylated tau-181 and glial fibrillary acidic protein *increased*.

It is worth being precise about what this trial establishes and what it does not. It establishes that at 1 mg/day for eight weeks, rapamycin does not reach measurable concentrations in the cerebrospinal fluid of these patients. It does not establish that mTOR in the brain was unengaged — cerebrospinal fluid concentration is an imperfect proxy for parenchymal exposure, and the subsequent demonstration that a different compound is metabolised to its active form *inside* brain parenchyma rather than in cerebrospinal fluid [15] shows how badly that proxy can mislead. Nor does it establish that autophagy was unaffected, because autophagy was not measured. The trial's primary endpoint was a pharmacokinetic quantity standing in for a pharmacodynamic one, and when the surrogate came back null there was no second instrument to consult.

The second study addressed dose and schedule rather than effect: an open-label pilot of weekly oral rapamycin in thirteen patients with early-stage Alzheimer’s disease, eleven at 7 mg weekly and two at reduced doses, over 26 weeks, with blood concentrations sampled at four timepoints in week 13 [51]. It reported moderate interindividual pharmacokinetic variability, a terminal half-life consistent with prior estimates, trough concentrations below immunosuppressive levels in all participants, and a recommendation that the 48-hour post-dose sample is optimal for monitoring. This is careful and useful work. It is also, entirely, about how much drug is in the blood.

The pattern is the one this paper is about. Both trials optimised *exposure*. Neither could ask, and neither asked, whether the exposure achieved did anything to autophagy. A dose-ranging exercise conducted without a pharmacodynamic endpoint is a search over a parameter whose optimum is unknown and unknowable by the method employed.

5. The Trial Record II — AMPK and Metformin

Metformin's rationale in Alzheimer's disease is doubly indirect. It arrives from epidemiology — the association between type 2 diabetes and dementia — and its autophagy connection runs through AMPK activation and consequent mTORC1 inhibition. Both steps are plausible. Neither has been measured in a patient's brain.

The pilot human study randomised twenty non-diabetic subjects with mild cognitive impairment or mild dementia due to Alzheimer's disease to metformin then placebo for eight weeks each, or the reverse [52]. Its most quoted finding is a pharmacological one: metformin was measurable in cerebrospinal fluid at an average steady-state concentration of 95.6 ng/mL. This is a genuine and useful result — it establishes central access, which the rapamycin study could not. Executive functioning improved, with trends in learning/memory and attention; cerebral blood flow by arterial spin labelling did not significantly change. The biomarkers collected were cerebrospinal fluid Alzheimer markers, neuroimaging and cognition. There was no autophagy measurement, and the rationale linking the drug to the disease was therefore untested in the trial that tested the drug.

The larger programme now under way inherits the same instrumentation. Metformin in Alzheimer Dementia Prevention (MAP) is a randomised phase II/III double-blind placebo-controlled trial of extended-release metformin in amnesic mild cognitive impairment [53]; MET-FINGER combines metformin repurposing with a multidomain lifestyle intervention [54]. Neither protocol, as published, specifies an autophagic or lysosomal pharmacodynamic endpoint.

The broader evidence on glucose-lowering drugs and cognition is meanwhile not encouraging. The GRADE randomised clinical trial compared glucose-lowering medications and glycaemia against cognitive outcomes and found no differential benefit of the kind the repurposing hypothesis anticipates [55]. This is not decisive against metformin — GRADE was not designed as a dementia-prevention trial and its participants were not selected for Alzheimer pathology — but it is the largest randomised evidence available and it points the wrong way for the hypothesis.

The instructive feature of the metformin programme for our argument is that it shows the problem is not simply one of drug access to brain. Metformin gets into cerebrospinal fluid at a quantified concentration. The programme still cannot say whether that concentration does what the hypothesis requires, because the hypothesis is about a rate and the trial measured a concentration.

6. The Trial Record III — TFEB, the Transcription Factor Without a Drug

TFEB is the most attractive node in the network and the one with no clinical-stage agent aimed at it. The attraction is well earned. TFEB is the master transcriptional regulator of a coordinated lysosomal gene network [27], acting through a shared promoter motif that integrates the cell's clearance pathways [29] and linking autophagy to lysosomal biogenesis as a single programme [28]. (We note that a published erratum to the last of these was issued in 2025 [56]; see Section 26.) In Alzheimer-relevant models, raising TFEB selectively clears aberrant tau species and rescues neurotoxicity [7]; neuronally targeted TFEB accelerates lysosomal degradation of the amyloid precursor protein and reduces plaque pathogenesis [30]; TFEB reduces the intraneuronal C99 fragment, which is the species the endosomal account identifies as the agent [31]; and astrocyte-targeted TFEB facilitates amyloid- β clearance [32].

That last result is more consequential for trial design than it is usually given credit for, and it pairs with a detail from the human tissue data. In laser-captured human CA1 neurons, MiTF/TFE activation in Alzheimer's disease was associated more strongly with glia than with neurons [3]. If the compensatory transcriptional response is largely glial while the degradative failure that matters is largely neuronal, then a bulk fluid measurement of TFEB reports predominantly the wrong cell. Serum TFEB has in fact been measured across the Alzheimer continuum and found *lower* in patients than in controls [10] — a level, in a mixed compartment, of a transcription factor whose activity is defined by its subcellular localisation rather than by its abundance.

There is, at present, no selective small-molecule TFEB agonist in clinical development for Alzheimer's disease. Every human intervention with a TFEB rationale reaches it indirectly: through mTOR (Section 4), through AMPK (Section 5), through cGMP/PKG (Section 8), or as a nutraceutical. This matters for our thesis in a particular way. The absence of a direct agent is often read as a medicinal-chemistry problem — transcription factors are hard targets. It is also a measurement problem. A programme cannot rationally invest in a hard target when it has no way to confirm that a candidate has hit it in the organ of interest. The two failures are not independent: the missing assay raises the cost of the missing drug.

7. The Trial Record IV — Trehalose, and the Largest Clean Test Anyone Has Run

Trehalose is the compound that comes closest to a fair test of the autophagy-induction hypothesis in a human neurodegenerative disease, and it did not happen in Alzheimer's disease.

Oral trehalose has a well-known problem: it is a disaccharide cleaved by intestinal trehalase, which limits systemic exposure and has confounded interpretation of oral studies. The HEALEY ALS Platform Trial regimen avoided this by giving trehalose intravenously — 0.75 g/kg weekly over 24 weeks, randomised 3:1 against matching placebo, stratified by concomitant edaravone and riluzole, at 60 sites, within an adaptive phase 2/3 master protocol with prespecified futility rules and interim analyses every twelve weeks [57]. The primary outcome was a Bayesian composite of the relative rate of functional decline on the revised ALS Functional Rating Scale and survival. The stated rationale was explicit: trehalose activates autophagy pathways in animal models of neurodegenerative disease, with the potential to catalyse clearance of misfolded proteins in motor neurons.

This is a well-constructed trial. It is adequately powered, properly randomised, placebo-controlled, delivered by a route that solves the bioavailability objection, and run inside a platform designed to make negative results interpretable. It is also, as far as the autophagy hypothesis is concerned, uninterpretable — because the trial had no measurement of autophagy. If the result is negative, there are at least four readings: trehalose does not induce autophagy in human motor neurons at this exposure; it does, but induction is not therapeutic in this disease; it does, and induction is harmful because completion is the failing step; or the hypothesis is right and the trial was too short. Nothing in the data distinguishes them.

The same structure recurs in the smaller trehalose studies. A randomised double-blind placebo-controlled trial of oral trehalose in spinocerebellar ataxia type 3 reported an interim analysis on clinical scales [58]. A phase II randomised study of intravenous trehalose in Alzheimer's disease analysed circulating microRNAs by direct hybridisation and bioinformatics [59] — a pharmacodynamic readout, and an imaginative one, but a readout of transcriptional consequence rather than of degradative rate.

Trehalose therefore occupies a specific place in the argument. It is the disproof of the hypothesis that the autophagy programme has failed merely for want of money, scale or rigour. Here was scale and rigour, in the right route, under a platform protocol. What was absent was the dial.

8. The Trial Record V — Sigma-1, Phosphodiesterase-5, and the Autophagy Rationale at Its Most Elastic

Two recent Alzheimer’s programmes carry an autophagy rationale in their published descriptions while measuring nothing autophagic. They are worth separating from the repurposing wave because in both cases the autophagy claim appears in the primary trial report itself.

Blarcamesine (ANAVEX2-73) is an orally available small-molecule activator of the sigma-1 receptor. The phase IIb/III trial, ANAVEX2-73-AD-004, randomised 508 participants with early Alzheimer’s disease to 30 mg, 50 mg or placebo once daily for 48 weeks across 52 centres in five countries [60]. Its own abstract states the mechanism as acting “through restoration of cellular homeostasis including autophagy enhancement.” The endpoints were the co-primary ADAS-Cog13 and ADCS-ADL, the secondary CDR-SB, plasma A β 42/40 ratio, and global brain volume by MRI. Autophagy appears in the mechanism sentence and in no endpoint.

The trial’s interpretation has been publicly contested, in a letter raising concerns about the conduct and analysis of the study published in the same journal [61], and the programme has since been reviewed systematically [62]. We take no position here on the clinical claim. Our point is narrower and holds whichever way that dispute resolves: a trial whose stated mechanism is autophagy enhancement, and which measures amyloid ratio, brain volume and cognition, has left its own mechanism untested. If the clinical result is positive, the field cannot attribute it to autophagy. If negative, the field cannot exonerate autophagy.

AR1001 (mirodenafil) is a phosphodiesterase-5 inhibitor whose Alzheimer rationale runs through cGMP/PKG signalling to autophagic and lysosomal function. Its phase 2 trial randomised participants aged 55–80 with mild-to-moderate dementia (MMSE 16–26) at 21 US sites to placebo, 10 mg or 30 mg once daily for 26 weeks with an optional 26-week extension, against co-primary endpoints of ADAS-Cog13 and ADCS-ADL change at week 26 [63]. Preclinical support includes amelioration of Alzheimer-like pathology by multimodal actions [64]. Again: no autophagic or lysosomal pharmacodynamic endpoint.

The sigma-1 and PDE5 programmes illustrate a specific hazard created by the missing assay, which we call **rationale elasticity**. When a mechanism cannot be measured, it cannot constrain. A compound may be described as an autophagy enhancer in its introduction and developed entirely against clinical scales, and no data at any point in its life will contradict the description. Mechanisms

that cannot be measured become rhetorical rather than experimental resources, and the field loses the ability to tell an autophagy drug from a drug with an autophagy paragraph.

9. The Trial Record VI — The Repurposing Wave

Four further compounds have been tested in Alzheimer's disease with rationales that are autophagic or endolysosomal in whole or in part. Each is instructive, and one is instructive in a way that ought to be uncomfortable.

Nilotinib is a tyrosine kinase inhibitor whose Alzheimer rationale is autophagic clearance of misfolded protein. The phase 2 randomised, double-blind, placebo-controlled study enrolled 37 participants with mild-to-moderate dementia due to Alzheimer's disease, biomarker-confirmed, taking 150 mg daily for 26 weeks then 300 mg for a further 26 weeks [65]. The trial measured a great deal: cerebrospinal fluid drug concentration, amyloid burden by positron emission tomography, cerebrospinal fluid A β 40 and A β 42, hippocampal volume, phosphorylated tau-181. It reported reduced frontal-lobe amyloid burden, reduced cerebrospinal fluid A β 40 at six months and A β 42 at twelve, attenuated hippocampal volume loss and reduced phospho-tau-181. A companion analysis examined cerebrospinal fluid microRNAs and inflammatory and vascular fibrosis markers through discoidin domain receptor-1 [66]. This is, by the standards of the field, an unusually well-instrumented small trial. It measured cargo, structure, inflammation and drug level. It did not measure clearance.

Lithium has an autophagy rationale through inositol monophosphatase and mTOR-independent autophagy induction, alongside GSK-3 β inhibition. A randomised trial of long-term lithium in older adults with amnesic mild cognitive impairment reported clinical and biological effects [67], with long-term follow-up thirteen years after the original trial subsequently published [68]; the literature has been meta-analysed across Alzheimer's and Parkinson's disease [69]. The biological measures were cerebrospinal fluid Alzheimer biomarkers. No autophagy endpoint appears.

Neflamapimod is the one compound in this list derived directly from the endosomal account: a p38 α kinase inhibitor that reverses Rab5 over-activation. The phase 2 REVERSE-SD study randomised 161 biomarker-confirmed participants with mild Alzheimer's disease to 40 mg twice daily or placebo for 24 weeks against a primary endpoint of episodic memory improvement [70], and did not meet it. The preclinical rationale and a randomised clinical evaluation in basal forebrain cholinergic degeneration were reported together [71]; the programme then moved to dementia with Lewy bodies, where a phase 2a informed the design of the phase 2b RewinD-LB trial [72], and where response was found to be associated with plasma phosphorylated tau status — that

is, the programme progressed by identifying and excluding participants with Alzheimer co-pathology [73]. The endosomal target was never measured in any of these patients; the stratification that made the programme work was performed on a tau marker.

An aside on what a repurposing rationale is worth without a readout

Four of the compounds in this section — nilotinib, lithium, metformin, hydroxychloroquine — reached Alzheimer’s disease by repurposing. This is now a large enterprise: formal pipelines exist for prioritising repurposing candidates for multi-arm platform trials [74], and real-world evidence for repurposable medicines in the Alzheimer pipeline has been critically evaluated using target trial emulation [75].

The methodological attention devoted to *selecting* repurposing candidates is, on the evidence of Table 1, not matched by attention to *instrumenting* them. This is a specific and correctable asymmetry. A prioritisation pipeline that ranked candidates partly by whether their proposed mechanism is measurable in humans would rank this set very differently, and would have ranked hydroxychloroquine — whose mechanism is measurable in no human, and whose direction is arguably the wrong one — below almost everything else.

Hydroxychloroquine is the uncomfortable one. It is a lysosomotropic weak base. Its pharmacology is to accumulate in acidic compartments and raise their pH — that is, to do to the lysosome what the disease is proposed to do. It was given to patients with early Alzheimer’s disease for eighteen months in a randomised, double-blind, placebo-controlled trial of 168 participants, on an anti-inflammatory rationale, and had no effect on progression [44]. Nobody measured lysosomal pH, in that trial or since, because there is no human method for doing so. A drug whose mechanism is the direct pharmacological reproduction of the proposed lesion was administered for a year and a half to the relevant population, and the experiment yielded no information about the lesion at all. If any single trial demonstrates the cost of the missing instrument, it is this one.

10. What the Record Has in Common

Table 1 puts the audit in one place. The column that carries the argument is the last one.

Table 1 — Human trials in Alzheimer’s disease and related neurodegeneration with an autophagic or endolysosomal rationale, and the pharmacodynamic measurements they made.

Compound	Node claimed	Design, N	Drug level measured	Other effect measures made	Flux or lysosomal readout
Rapamycin [8]	mTORC1 → induction	Open-label phase 1, n=10, 8 wk	CSF drug level (undetectable)	CSF p-tau181 (rose), GFAP (rose), cognition	None
Rapamycin [51]	mTORC1 → induction	Open-label PK pilot, n=13, 26 wk	Blood levels, 4 timepoints	None	None
Metformin [52]	AMPK → mTORC1	Randomised crossover, n=20, 8+8 wk	CSF drug level (95.6 ng/mL)	CSF AD markers, ASL-MRI, cognition	None
Trehalose [57]	Induction (TFEB-linked)	Randomised phase 2/3, IV, 24 wk	Not reported in abstract	ALSFRS-R, survival	None
Trehalose [59]	Induction	Randomised phase II, IV, in AD	—	Circulating microRNAs	None
Blarcamesine [60]	Sigma-1 → “autophagy enhancement”	Randomised phase IIb/III, n=508, 48 wk	—	ADAS-Cog13, ADCS-ADL, CDR-SB, plasma Aβ42/40, MRI volume	None
AR1001 [63]	PDE5 → cGMP/PKG	Randomised phase 2, 26 wk	—	ADAS-Cog13, ADCS-ADL	None
Nilotinib [65]	Autophagic clearance	Randomised phase 2, n=37, 52 wk	CSF drug level	Amyloid PET, CSF Aβ40/42, p-tau181, hippocampal volume	None

Compound	Node claimed	Design, N	Drug level measured	Other effect measures made	Flux or lysosomal readout
Lithium [67]	mTOR-free induction	Randomised, aMCI, long-term	Serum lithium	CSF AD biomarkers, cognition	None
Neflamapimod [70]	p38 α $\rightarrow$ Rab5	Randomised phase 2, n=161, 24 wk	—	Episodic memory, CSF AD biomarkers	None
Hydroxychloroquine [44]	(De-acidifies lysosomes)	Randomised, n=168, 18 mo	—	Clinical progression	None

Three features of this table deserve comment.

First, the last column is uniform. We have not found an exception. Across eleven completed human studies spanning twenty-five years, four mechanistic nodes and two diseases, autophagic or lysosomal function was measured in none of the participants. This is not a claim that the trialists were careless; it is a claim about what was available to them.

Second, the pharmacokinetic column is doing the pharmacodynamic column's work. Where any mechanistic measurement was made at all, it was almost always drug concentration. The rapamycin phase 1 states this explicitly — its *primary aim* was to evaluate central nervous system penetrance [8]. When a field routinely elevates exposure to a primary endpoint, it is telling you that it has no effect endpoint.

Third, the measurements that were made are measurements of cargo. Amyloid PET, cerebrospinal fluid A β 40/42, phosphorylated tau, hippocampal volume: every one of these reports the accumulated product of a process, not the rate of the process. This is a coherent choice in a production theory, where the product is the object of interest. It is close to uninterpretable in a clearance theory, where the same product level can arise from a high production rate with normal clearance or a normal production rate with impaired clearance. The clearance hypothesis has been tested, throughout, with instruments built for the production hypothesis.

Why this happened is the subject of the next section.

II. Flux Is a Rate

The core of the problem is arithmetic, and stating it in arithmetic makes clear why no amount of better assay chemistry will solve it.

Let A be the abundance of autophagosomes in a cell — operationally, the quantity of lipidated LC3 (LC3-II) on autophagosomal membranes, or the count of LC3-positive puncta. Let $\mathcal{J}_{\text{in}}$ be the rate at which autophagosomes are formed and k_{out} the first-order rate constant for their consumption by fusion with lysosomes and degradation of their contents. Then

$$dA/dt = \mathcal{J}_{\text{in}} - k_{\text{out}} \cdot A$$

and at steady state,

$$A = \mathcal{J}_{\text{in}} / k_{\text{out}}$$

Autophagic flux is $\mathcal{J}_{\text{in}}$ at steady state — the quantity of material passing through the pathway per unit time. The measurable quantity A is a ratio. An elevated A is equally consistent with a doubled formation rate and with a halved consumption rate. A normal A is consistent with both rates doubled, both halved, or both normal. There is no function from a single observation of A to either rate, and no improvement in the precision of the A measurement creates one. The system is unidentified in the formal sense: the parameter of interest is not determined by the data.

This is not an obscure point; it is the founding methodological principle of the autophagy field, stated at length in successive editions of the community guidelines, which devote substantial space to the distinction between “autophagy markers” and “autophagic flux” and to the errors that follow from conflating them [76,77]. The guidelines’ operational prescription is unambiguous: to measure flux, compare LC3-II (or an equivalent marker) in the presence and absence of an agent that blocks lysosomal degradation — bafilomycin A1, chloroquine, leupeptin, or an equivalent. The difference between the clamped and unclamped states is the flux.

The reason this works is worth making explicit, because it determines what any human assay must reproduce. Adding a lysosomal inhibitor sets $k_{\text{out}} \approx 0$. The system then departs steady state and A accumulates at rate $\mathcal{J}_{\text{in}}$. Measuring the slope of that accumulation — or, in the common two-timepoint version, the increment over a fixed interval — estimates $\mathcal{J}_{\text{in}}$ directly. The clamp converts an unidentified ratio into an identified rate by removing one of the two unknowns.

Every alternative laboratory approach does the same thing by other means. The tandem-fluorophore reporters (mRFP-GFP-LC3 and its relatives) exploit differential pH-quenching so that the acidification state of each compartment reports its position in the pathway, converting a static image into a two-state observation within a single cell — this is the reporter with which autolysosomal acidification failure was demonstrated to precede amyloid deposition in mouse models [45]. The GFP-LC3-RFP-LC3ΔG probe is cleaved by endogenous ATG4 into a degradable and a stable moiety, so that the ratio of the two reports cumulative flux against an internal control generated in the same cell [78,79]. Metabolic labelling approaches measure the disappearance of labelled substrate. In each case the assay obtains a second state: a blocked condition, a pH-distinguished compartment, a co-generated control, or a time-resolved label.

The general principle: a rate cannot be recovered from any number of level measurements taken in a single state. Recovering it requires either a perturbation of known sign, an internal reference generated by the same process, or time resolution of a label. Any candidate human assay must supply one of these three, and Sections 21 and 22 are organised around that requirement.

A worked case

The abstraction is easier to hold with numbers, and the numbers are not academic: they correspond to the two competing readings of the human Alzheimer data.

Take a control neuron with $\mathcal{J}_{\text{in}} = 100$ compartments per hour and $k_{\text{out}} = 1.0$ per hour, giving $A = 100$. Now consider two diseased neurons.

Neuron X has doubled its formation rate as a compensatory response and retained normal degradation: $\mathcal{J}_{\text{in}} = 200$, $k_{\text{out}} = 1.0$, so $A = 200$.

Neuron Y has a normal formation rate and has lost half its degradative capacity: $\mathcal{J}_{\text{in}} = 100$, $k_{\text{out}} = 0.5$, so $A = 200$.

The two neurons return an identical marker level, and they are in opposite therapeutic situations. Neuron X is already doing the thing an inducer would make it do more of, and its problem, if it has one, lies elsewhere. Neuron Y needs its degradative step repaired, and an inducer given to it raises A to 400 without moving a single additional molecule through to degradation — the cell accumulates faster and clears no more.

Now make it realistic. The human CA1 data indicate that both changes are present simultaneously [3]: induction is up *and* clearance is down. A neuron with $\mathcal{J}_{\text{in}} = 200$ and $k_{\text{out}} = 0.5$ returns $A = 400$, which correctly signals severe abnormality but says nothing about which term to treat. And a neuron with $\mathcal{J}_{\text{in}} = 50$ and $k_{\text{out}} = 0.5$ — halved induction, halved clearance — returns $A = 100$, indistinguishable from the healthy control while passing material at half the rate. This last

case is the one that should worry anyone proposing to use a level as a screening or stratification variable: **the marker is normal and the flux is halved.**

The therapeutically relevant quantity in every one of these cells is $\tilde{J}_{\text{in}}$, the throughput. It is 100, 200, 100, 200 and 50 respectively, against marker levels of 100, 200, 200, 400 and 100. There is no monotone relationship between the two columns, and no calibration curve can create one.

12. Every Human Readout Is a Level

With the requirement stated, the audit of what exists in humans becomes short.

Cerebrospinal fluid LC3A and LC3B. The largest such study measured LC3A and LC3B concentrations in cerebrospinal fluid in 568 participants from the Alzheimer’s Disease Neuroimaging Initiative and related them to cognition, dementia risk, Alzheimer biomarkers, microglial activity, brain volume and metabolism, cross-sectionally and longitudinally, with subgroup analyses by ATN classification and receiver-operating-characteristic analyses of discriminative performance [9]. This is a large, careful, well-powered study, and it reports a *concentration*, in a single state, in a compartment whose relationship to intracellular LC3 is unestablished.

Serum and cerebrospinal fluid mitophagy proteins. PINK1, BNIP3L and TFEB were quantified in cerebrospinal fluid and serum from 246 individuals across the continuum from cognitive unimpairment through mild cognitive impairment due to Alzheimer’s disease to Alzheimer dementia [10]. Serum and cerebrospinal fluid PINK1 and serum BNIP3L were higher in patients; serum TFEB was lower. Again: concentrations, single state.

Cerebrospinal fluid and serum ULK1. ULK1 was found reduced in serum and cerebrospinal fluid with age in cognitively unimpaired participants and in patients with Alzheimer’s disease across two cohorts totalling 391 people, in a study whose mouse arm demonstrated that raising ULK1 stimulates autophagic flux and reduces pathology [11]. The contrast within that single paper is the paper’s own demonstration of our point: *flux* was measured in the mouse; a *level* was measured in the human.

Cerebrospinal fluid lysosomal network proteins. An early and under-cited study measured a panel of lysosomal network proteins in cerebrospinal fluid as candidate Alzheimer biomarkers [12]. Concentrations.

Cerebrospinal fluid p62/SQSTM1. Measured in dementia with the explicit question of whether it could serve as an *in vivo* autophagy biomarker [13]. A concentration — and one with a particularly acute interpretive problem, since p62 is both an autophagy receptor consumed by the pathway and a transcriptional target induced by its stress.

That is the inventory. Five classes of measurement, all concentrations, none perturbed, none compartment-resolved, and — a point taken up in the next section — none with an established cell of origin.

The strongest form of the objection to this section runs: *these are early days; longitudinal within-person change in a level is a rate*. It is not, and the distinction matters. The longitudinal change in a cerebrospinal fluid concentration is the rate of change of a ratio, which depends on the trajectories of both $\mathcal{J}_{\text{in}}$ and k_{out} and on the trajectory of whatever governs release of the marker into cerebrospinal fluid. Longitudinal sampling improves precision on $A(t)$. It does not identify $\mathcal{J}_{\text{in}}$.

A second objection: *a drug's effect on the level is informative even if the level is ambiguous, because the direction of the drug's action is known*. This is the best argument available for the current instruments, and it is partially right — if one knew with certainty that rapamycin acts only on $\mathcal{J}_{\text{in}}$, then a change in A under rapamycin would be informative about k_{out} . But the premise fails: mTOR inhibition raises TFEB activity and therefore acts on lysosomal biogenesis as well as on initiation [28], so rapamycin moves both parameters in the same direction, and the sign of its effect on A is genuinely indeterminate a priori. The drugs the field most wants to dose are precisely the drugs for which the single-state inference fails.

The instrument the field did build

It would be wrong to conclude from any of this that fluid biomarker development in Alzheimer's disease is immature or that the field is incapable of building instruments. The opposite is true, and the contrast tells us what the conditions for success are.

Over roughly fifteen years the field built, validated and deployed a blood test for Alzheimer pathology. Phosphorylated tau at residue 217 discriminates Alzheimer's disease from other conditions with diagnostic accuracy sufficient for clinical use, a conclusion now supported by systematic review and meta-analysis across assay platforms and cohorts [80] (a correction to which has since been published [81]); head-to-head comparisons of commercial platforms have been performed [82]; and the analyte is used for enrolment, for stratification and, in at least one programme, for exclusion — the neflamapimod development path in dementia with Lewy bodies proceeded by using plasma phosphorylated tau to identify participants whose response differed, effectively screening out Alzheimer co-pathology [73].

Two features made that achievement possible and neither is available for autophagy. The analyte is a **soluble protein exported from neurons into the extracellular space and thence into cerebrospinal fluid and blood**, so that a peripheral measurement has a defined relationship to a central process. And the quantity of interest is genuinely a **concentration** — the question “how much phosphorylated tau is there” is the question the pathology poses, and a level answers it.

The same two properties account for the only unambiguously instrumented therapeutic programme this disease has produced. The anti-amyloid antibodies were developed against amyloid positron emission tomography as a target-engagement and dose-selection readout, and both pivotal

trials report plaque removal as a quantity with a dose-response and a time course alongside their clinical endpoints [83,84]. Whatever view one takes of the size of the clinical effect, the programme could say what its drug did to its target, in the organ, in the patient. That is the capacity at issue in this paper, and it is the capacity the autophagy programme has never had.

Autophagy has neither property. Its defining event is intracellular and its product is the absence of material. The field's demonstrated competence at building fluid biomarkers is therefore not transferable, and the fifteen years that produced p-tau217 would not, applied to autophagy, produce anything, because the effort would be directed at finding a better level.

13. Three Ways a Level Misleads

The formal objection of Section 11 is that a level underdetermines a rate. Three further problems are specific to the human Alzheimer setting, and each has an empirical demonstration in the literature already cited.

The sign ambiguity is not hypothetical in this disease — both terms move, in opposite directions, at once. The human CA1 data show autophagy-related gene expression, autophagosome-component protein and LC3-positive puncta all *increased*, and substrate clearance simultaneously *decreased*, with autolysosomal accumulation of LC3-II and p62 [3]. Any single-state marker in such a brain is the resultant of two opposed changes of unknown relative magnitude. It is the one circumstance in which a level is maximally uninformative, and it is the circumstance that obtains.

The compartment problem: the readout is bulk and the lesion is local. The account of the disease under test is specific about *where* the failure is — the terminal, acidification-dependent step, inside the autolysosome, in neurons [45,22]. A cerebrospinal fluid concentration has no compartment resolution whatever. Worse, the compensatory response that would dominate such a measurement appears to be in the wrong cell: MiTF/TFE activation in Alzheimer tissue was associated more strongly with glia than with neurons [3], and astrocytic and neuronal TFEB manipulations have distinct consequences in models [32,30]. A fluid measurement of a lysosomal transcriptional programme in Alzheimer’s disease may be predominantly a glial signal reporting compensation, in a disease whose relevant lesion is neuronal and uncompensated.

The provenance problem: how the marker got into the fluid determines what it means. A protein appears in cerebrospinal fluid by one of three routes — regulated secretion, release in extracellular vesicles, or escape from dying or damaged cells — and the three have opposite relationships to intracellular flux. If LC3 reaches cerebrospinal fluid chiefly from dying neurons, its concentration is a neurodegeneration marker wearing an autophagy label, and would be expected to correlate with neurofilament light rather than with any degradative rate. This is not a speculative concern: the success of neurofilament light as a biomarker rests precisely on the third route [85], and the recent demonstration that even neurofilament light shows a delayed, apparently controlled release into the extracellular milieu rather than simple passive leakage [20] shows that provenance is not settled even for the field’s best-characterised fluid marker. For LC3, PINK1, BNIP3L, TFEB and p62 in cerebrospinal fluid, provenance is not merely unsettled; it is, as far as we can determine, unexamined.

14. Why Imaging Has Not Closed the Gap

If fluid cannot supply the measurement, imaging is the obvious alternative, and it is worth saying plainly what is and is not available.

There is no positron emission tomography tracer in human use that reports autophagic flux, and none that reports lysosomal pH. Searching the indexed literature for a human brain autophagy tracer returns no primary reports. What exists is a set of adjacent measurements, each reporting something other than the rate.

Fluorodeoxyglucose reports regional glucose metabolism — an energetic readout that would be expected to fall with, among many other things, a failing ATP-consuming proton pump, but which is not remotely specific to it.

Translocator protein tracers report microglial activation. They have been used as pharmacodynamic markers in neuroinflammatory programmes, and preclinical work with a brain-penetrant NLRP3 inhibitor used 18F-DPA-714 alongside dopaminergic tracers and gadolinium-contrast MRI to demonstrate multi-modal target engagement in Parkinson models [86]. This is an instructive precedent for what an imaging pharmacodynamic package looks like when one exists.

Amyloid and tau tracers report cargo. They are the most successful molecular imaging agents in neurology and they are, for our purposes, measurements of the accumulated product — the numerator problem of Section 10 in imaging form.

The reason no flux tracer exists is the reason of Section 11 in a different medium. A tracer reports the concentration of a bound or trapped species, which is a level. To make it report a rate one must either use the tracer's own kinetics (requiring a compartmental model whose parameters are themselves unidentified without a perturbation) or perturb the system. The pharmacological clamp reappears, and with it the ethical objection.

15. The Preclinical Assays That Do Not Travel

It is worth being precise about *why* the standard assays stop at the species boundary, because the reason is not technical difficulty and cannot be engineered away by better instruments.

The clamp assay requires a lysosomal poison. Bafilomycin A1 is a vacuolar H⁺-ATPase inhibitor; chloroquine and its relatives are lysosomotropic bases; leupeptin is a protease inhibitor. Each works by disabling the degradative step in order to measure the rate at which material arrives at it. Administering such an agent to a human being in order to measure autophagy would mean inducing, deliberately and systemically, the exact lesion the disease is proposed to consist of. There is a hard version of this objection (one should not poison a patient's lysosomes for a measurement) and a soft version that is scientifically more interesting: even if it were acceptable, a clamp applied to a brain already at the edge of degradative failure would be expected to have non-linear effects, so the perturbation would not be small and the estimate would not be of the unperturbed rate.

We note in passing that a partial natural experiment exists. Chronic lysosomotropic agents are prescribed in large populations — hydroxychloroquine most obviously, and the eighteen-month randomised trial of it in early Alzheimer's disease [44] is an eighteen-month human lysosomal clamp with a clinical readout attached. The trial was neutral. Because no lysosomal measurement accompanied it, the result constrains nothing about the lysosomal hypothesis.

The reporter assays require a transgene. The tandem-fluorophore LC3 reporters and the GFP-LC3-RFP-LC3ΔG internal-control probe [78,79] are genetically encoded. Their elegance consists precisely in generating the second state inside the cell being measured. That property is what makes them untransferable: one cannot express a reporter in the neurons of a living patient, and the readout is microscopic rather than remote.

The morphological assays require tissue. Immunogold electron microscopy on biopsy material [2] and laser-capture microdissection of identified neurons [3] are the two techniques that have produced the best human data in this area, and both require the brain. Biopsy is not a clinical-trial instrument, and post-mortem tissue supplies a single frame from the end of a process running for decades, at an unknown and variable post-mortem interval to which autophagic compartments are notoriously sensitive.

This is the crux of the paper. The three assay families that identify flux each require something a living patient cannot provide: a poison, a transgene, or their brain. That is why the measurement has not crossed over, and it is why closing the gap requires a different assay architecture rather than a better version of an existing one.

16. A Target Where the Readout Exists: NLRP3

If the argument is that a missing readout explains a missing translation, the argument owes a control. The NLRP3 inflammasome supplies one, and the comparison is close enough to be informative.

The similarities are substantial. NLRP3 is an intracellular multiprotein complex, assembled in the cytoplasm, predominantly in myeloid cells, whose activity is a rate rather than a level and whose assembly is not directly observable in a living human. Its evidential position in Alzheimer's disease is comparable in strength to autophagy's: NALP3 was shown to mediate the innate immune response to amyloid- β through lysosomal damage [87]; NLRP3 was shown to be activated in human Alzheimer brain and to contribute to pathology in APP/PS1 mice [88]; and inflammasome activation was subsequently shown to drive tau pathology [89]. A well-characterised tool inhibitor, MCC950, was developed and its binding site mapped to the NLRP3 ATP-hydrolysis motif [90]. Like the autophagy programme, this one is preclinically saturated.

The critical difference is a single structural fact about the target. **NLRP3's output is a secreted cytokine.** Activation cleaves pro-interleukin-1 β and pro-interleukin-18 into mature, exported forms. The consequence is that the activity of an invisible intracellular complex produces a measurable extracellular species, in the compartment a lumbar puncture samples. Autophagy has no equivalent: its output is the *disappearance* of intracellular material, and disappearance leaves nothing to assay.

That difference has permitted a development path the autophagy programme has not had. NT-0796 was designed as a brain-penetrant NLRP3 inhibitor and reported as a clinical candidate for neuroinflammatory disorders [91]. Its unusual pharmacology was then characterised in detail: it is an ester prodrug whose potency in human monocytes derives from cell-penetrance followed by intracellular conversion to the active acid NDT-19795 by carboxylesterase-1 [92,93]. In a high-fat-diet model, it and a related compound reversed obesity, systemic inflammation and astrogliosis [94].

The clinical step is the one that matters here. A phase 1b open-label study administered 150 mg twice daily to four elderly healthy volunteers for seven days and to ten subjects with Parkinson's disease for 28 days, with serial blood sampling and lumbar punctures [14]. It reported significant cerebrospinal fluid exposure in both groups; reductions in systemic high-sensitivity C-reactive protein, fibrinogen, interleukin-6 and interleukin-18; and reductions in neuroinflammatory markers

including interleukin-1 β in the cerebrospinal fluid of the Parkinson's participants, which the authors interpret as evidence of central target engagement. The pharmacological basis of the central effect was subsequently established: human and non-human primate cerebrospinal fluid cannot metabolise the prodrug, but brain parenchymal homogenates can, and carboxylesterase-1 in the brain is expressed by a subpopulation of cells largely positive for the microglial marker IBA1 [15].

The compound has since been given an international non-proprietary name, ruvonoflast, and tested in a properly randomised double-blind phase 1b in 63 participants with elevated inflammatory and cardiovascular risk [16]. It met its primary endpoint with a posterior probability exceeding 99%: geometric least-squares mean reduction in high-sensitivity C-reactive protein of 82.2% (95% CI 75.9–86.8) versus 37.2% (7.0–57.6) on placebo at day 28, with separation from day 3, return to baseline within seven days of discontinuation, and significant reductions in interleukin-6 and fibrinogen. Ten per cent of treated participants discontinued for transient reversible adverse events, against none on placebo.

The class has moved elsewhere too: a selective NLRP3 inhibitor, selnoflast, completed a randomised phase 1b in ulcerative colitis with a dose chosen to achieve 90% plasma interleukin-1 β inhibition [95] — note the dose-selection logic, which is exactly what the autophagy programme cannot perform — and a second brain-penetrant inhibitor, MCC7840 (inzomelid/emlenoflast), has been characterised preclinically with a multi-tracer PET-MRI pharmacodynamic package [86].

The two targets are not merely analogous; they are connected

There is a detail in the founding inflammasome paper that makes the comparison sharper than a simple analogy, and it deserves to be stated because it cuts against any suggestion that we have chosen a conveniently unrelated control.

The mechanism by which amyloid- β activates the NLRP3 inflammasome is **lysosomal**. Phagocytosed fibrillar amyloid causes lysosomal swelling and rupture, releasing cathepsin B into the cytosol, and it is that release which licenses inflammasome assembly [87]. The inflammasome signal in Alzheimer's disease is therefore, in substantial part, a *readout of a lysosomal event* — the same organelle, the same failure of containment, the same disease.

This has two consequences for the argument.

The first is that the control is a strong one. If two targets sit on the same organelle in the same disease, and one has produced a development programme with central target engagement demonstrated in fourteen subjects over 28 days while the other has produced eleven uninstrumented trials over twenty-five years, the difference is unlikely to be about the organelle's tractability, the disease's difficulty, or the availability of compounds. It is about what the target emits.

The second is more useful. If inflammasome activation is downstream of lysosomal membrane permeabilisation, then cerebrospinal fluid interleukin-1 β and interleukin-18 are, among other things, *indirect reporters of lysosomal failure* — not of flux, but of one specific catastrophic mode of it. This does not satisfy any of the six requirements of Section 21: it reports a level, of a downstream consequence, in a myeloid cell, with no rate and no compartment resolution, and it will move for many reasons unrelated to autophagy. But it is the closest thing to a lysosomal readout that human medicine currently possesses, and it arrived by accident, as a by-product of an inflammation programme. That a field hunting for a lysosomal measurement should find its nearest approximation in a cytokine assay built for a different purpose is an accurate picture of where the instrumentation stands.

A caution attaches, and it is the one that prevents this observation from being turned into a proposal. Lysosomal rupture is the *end state* of the failure mode the disease account describes, not its process. A marker of rupture will be at its most informative in the tissue and at the stage where cells are already dying, which is where a preventive therapeutic most needs not to be aimed.

17. What a Movable Readout Bought

It is worth itemising what the NLRP3 programme obtained that the autophagy programme did not, because the list is shorter and more specific than enthusiasm for biomarkers usually implies.

It bought dose selection against effect rather than exposure. The selnoflast study chose 450 mg once daily to achieve 90% interleukin-1 β inhibition [95]. That is a dose chosen against a pharmacodynamic target. Compare the rapamycin phase 1, whose primary aim was to determine whether any drug reached the cerebrospinal fluid at all [8], and the rapamycin pharmacokinetic pilot, whose recommendation concerned the optimal timing of a blood sample [51].

It bought a go/no-go decision at phase 1. The NT-0796 phase 1b could conclude that central target engagement had occurred, on evidence internal to the study, in fourteen subjects over 28 days [14]. No autophagy trial in Alzheimer's disease has ever been able to reach that conclusion, at any sample size or duration.

It bought a route by which failure would be informative. If a properly powered NLRP3 trial in neurodegeneration now fails with cerebrospinal fluid interleukin-1 β suppressed, the field learns something definite: that central inflammasome activity is not on the causal path to decline in the population and stage tested. This is precisely the inference the autophagy programme cannot draw from any of its eleven negative or neutral studies.

It bought a mechanism-specific safety signal. A pharmacodynamic marker that returns to baseline within seven days of discontinuation [16] tells a developer about reversibility, which is information about risk as well as about efficacy.

What it did not buy is anything clinical. That is the subject of the next section, and the honesty of the comparison depends on it.

18. The Limits of the Contrast

Four qualifications must be attached to Section 16, and we state them at full strength because the argument does not need the contrast to be flattering.

The Parkinson’s study was small, open-label and uncontrolled. Fourteen subjects, no placebo arm, open-label administration [14]. Biomarker changes in such a design are vulnerable to regression to the mean, to the effects of lumbar puncture itself on inflammatory markers, and to expectancy. The randomised evidence for the compound’s anti-inflammatory effect is in a cardiovascular population with a systemic endpoint [16], not in a neurodegenerative population with a central one.

The most robust readout is peripheral. High-sensitivity C-reactive protein is a hepatic acute-phase protein. An 82% reduction in it is a striking demonstration that the drug does something to systemic inflammation; it is not a measurement of the brain. The central readouts — cerebrospinal fluid interleukin-1 β above all — are the ones that matter for neurodegeneration, and interleukin-1 β sits near or below the lower limit of quantification of many immunoassays in normal cerebrospinal fluid, which makes a reported reduction technically demanding and dependent on assay particulars.

No NLRP3 inhibitor has shown clinical benefit in any neurodegenerative disease. The compounds are unapproved for these indications, the human data are preliminary, and a 2026 review of the neuroimmune network in Alzheimer’s and Parkinson’s diseases states as much explicitly while characterising the NT-0796 phase 1b/2a as showing preliminary reductions in biomarkers of axonal damage and T-cell activation in Parkinson’s disease [96]. We record that characterisation as a review’s summary rather than as a primary result: the abstract of the primary study names high-sensitivity C-reactive protein, fibrinogen, interleukin-6, interleukin-18 and cerebrospinal fluid interleukin-1 β , and does not name neurofilament light [14].

Soluble TREM2 has not, to our knowledge, been reported as a readout for this compound. It is worth naming because it is the obvious candidate. Cerebrospinal fluid soluble TREM2 is one of the best-validated microglial-state markers available: it rises early in Alzheimer’s disease in association with tau-related neurodegeneration rather than with amyloid pathology [97], and higher levels are associated with *reduced* subsequent cognitive and clinical decline [98]. Paired with neurofilament light, whose serum dynamics predict neurodegeneration and clinical progression years before symptom onset [85], it would give a neuroinflammation programme a two-sided readout — a state marker and a damage marker. We have not found a published report of either in an NT-0796 trial. The point of naming them is therefore not to claim that they were moved but to

observe that in the inflammasome programme such readouts *exist and could be applied*, whereas in the autophagy programme no analogous pair exists to apply.

The corrected form of the contrast is this. NLRP3 is not a success story; it is an *instrumented* story. What the readout bought is the ability to be wrong informatively. That is a smaller claim than the usual one and it is the only one the evidence supports.

19. The Deeper Problem: Four Results from 2025 and 2026

The obvious objection to everything above is that biomarkers have lately been a disappointment in this disease, and that a field which has watched biomarkers move without patients improving might reasonably decline to build another one. Four results from the last eighteen months make the objection concrete. Two are peer-reviewed; two are not, and we handle them differently.

Valacyclovir (peer-reviewed). The VALAD randomised clinical trial tested valacyclovir 4 g/day against placebo in 120 participants with early symptomatic Alzheimer’s disease and herpes simplex virus seropositivity, across three US memory-disorder clinics, with recruitment from January 2018 to May 2022 and final follow-up in September 2024 [17]. The primary outcome was change at 78 weeks on the 11-item ADAS-Cognitive subscale, with secondary outcomes including ADCS-ADL, 18F-florbetapir amyloid PET standardised uptake value ratios across six brain regions and 18F-MK-6240 tau PET across four medial temporal regions.

The result was not neutral. Least-squares mean change on the ADAS-Cognitive subscale at 78 weeks was 10.86 (95% CI 8.80 to 12.91) on valacyclovir against 6.92 (4.88 to 8.97) on placebo — a between-group difference of 3.93 (1.03 to 6.83; $P = .01$) **in the direction of greater cognitive worsening on active treatment.** Neither amyloid nor tau PET differed between groups. The authors conclude that valacyclovir was not efficacious, that it was accompanied by cognitive worsening, and that it is not recommended in this population.

Its rationale — an infectious contribution to Alzheimer pathogenesis — was strong enough to justify the study, and the study was designed to test it in the enriched population where it should be detectable. That it returned a harm signal rather than a null is the kind of result only an adequately instrumented, adequately controlled trial can produce, and it is worth more to the field than a neutral result would have been.

Oral semaglutide (peer-reviewed). The evoke and evoke+ trials randomised 3,808 participants with amyloid-confirmed mild cognitive impairment or mild dementia due to Alzheimer’s disease, aged 55–85, across 566 sites in 40 countries, to oral semaglutide up to 14 mg once daily or placebo for up to 156 weeks, with a primary endpoint of change in CDR-SB at week 104 [18]. Mean CDR-SB change from baseline to week 104 was 2.3 with semaglutide against 2.3 with placebo in evoke (estimated difference -0.08 , 95% CI -0.35 to 0.20 ; $P = 0.57$) and 2.2 against 2.1 in evoke+ (0.10 , -0.17 to 0.38 ; $P = 0.46$). Both trials were discontinued for negative clinical outcome. The rationale — reduced dementia risk after GLP-1 receptor agonist exposure in animal,

clinical and real-world data — was among the strongest epidemiological cases in the repurposing literature, and the trials were among the largest ever conducted in the disease.

Two features of these trials are worth drawing out before turning to the unpublished results, because they show that the field’s instrumentation problem is selective rather than general.

Both trials were adequately instrumented for their own hypotheses. VALAD enrolled only herpes-seropositive participants — that is, it stratified on the presence of the factor its hypothesis implicated — and measured amyloid PET across six regions as a secondary [17]. evoke and evoke+ required amyloid confirmation for entry [18]. These are well-designed studies whose negative results are therefore interpretable: the field now knows something it did not know before.

Neither could have been designed that way for autophagy. The equivalent of “herpes-seropositive” for a clearance hypothesis is “has demonstrable degradative impairment,” and there is no test for it. This is the enrolment problem of Section 23 seen from the other end: the reason no autophagy trial has enrolled on its own lesion is not that the idea has not occurred to anyone.

Posdinemab (conference disclosure; not peer-reviewed). Posdinemab (JNJ-63733657) is a humanised IgG1/ κ monoclonal antibody binding phosphorylated tau at amino acid 217. Its peer-reviewed record comprises phase 1 studies in healthy participants and participants with Alzheimer’s disease, which established linear pharmacokinetics and dose-dependent reductions in free and total p217+tau in cerebrospinal fluid [99]; a bridging phase 1 study in healthy Chinese adults [100]; and a mechanistic population pharmacokinetic–pharmacodynamic model built on those data to guide phase 2 dose selection, which predicted that 1,000 mg and 3,000 mg every four weeks would achieve greater than 90% reduction in interstitial-fluid tau seeds by day 391 and day 154 respectively [101].

The phase 2b result is **not available as a peer-reviewed publication and we do not cite one**. According to reporting of the presentation at AD/PD 2026 (Copenhagen, 17–21 March 2026), the Autonomy trial (NCT04619420) randomised 523 participants to 1,000 mg or 3,000 mg posdinemab every four weeks or placebo, of whom 412 completed two years; the primary endpoint, change in the integrated Alzheimer’s Disease Rating Scale for mild cognitive impairment, was not met at either dose; cerebrospinal fluid p-tau217 was abolished at both doses by the one-year timepoint and remained so; tau PET showed a slight reduction of accumulation in the meta-temporal region but no difference from placebo in tau-naïve regions; and the programme was discontinued. **This account rests on conference reporting and company communication, not on a peer-reviewed paper, and should be treated accordingly.** What can be verified from the indexed literature is the antibody’s pharmacology and its capacity to suppress cerebrospinal fluid p217+tau [99,101] — that is, the target-engagement half of the story, and not the clinical half.

Diranersen (company disclosure; not peer-reviewed). Diranersen (BIIB080) is an antisense oligonucleotide lowering tau expression. Its phase 2 CELIA study (NCT05399888) reported topline results by company press release on 14 May 2026, with detailed results presented at the Alzheimer’s Association International Conference in July 2026. According to those disclosures, the trial **did not meet its primary endpoint**, which was dose response for change from baseline in CDR-SB at week 76; it reported reductions in cerebrospinal fluid tau and in tau pathology by PET across all doses studied, without a clear dose–response relationship; and pre-specified analyses reported slowing of clinical decline across dose groups, most notably at the lowest dose tested (60 mg every 24 weeks). The sponsor announced continuation into registrational development. **We have found no peer-reviewed publication of these results and cite none.** The evidentiary status is company disclosure plus conference presentation; readers should weight it accordingly, and the discordance between a missed primary endpoint and a positive interpretation is itself a reason for caution.

20. Three Kinds of Biomarker, and Which One Autophagy Lacks

The four results of Section 19 do not undermine the argument of this paper. They sharpen it, provided one distinction is held.

Biomarkers in therapeutic development do three distinct jobs, and conflating them is the source of most of the confusion in this area.

A **target-engagement marker** answers: did the drug reach and act on its molecular target in the organ of interest? Cerebrospinal fluid p217+tau under posdinemab is a near-perfect example [99,101]. Cerebrospinal fluid interleukin-1 β under NT-0796 is a partial one [14].

A **pathway or pharmacodynamic marker** answers: did the biological process the target controls change? These are harder, because they require a measurement downstream of the target and upstream of the clinical outcome. Tau PET under diranersen is one. Autophagic flux would be one.

A **surrogate endpoint** answers: does a change in this quantity reliably predict a change in clinical outcome? This is the hardest of the three, requires validation across multiple interventions, and is the standard almost nothing in this field has met.

The 2025–26 results are, read correctly, failures of the *third* kind while the first two succeeded. Posdinemab abolished its target-engagement marker and failed clinically; diranersen moved its pathway marker at every dose and failed its dose-response primary. The correct conclusion is not “biomarkers are useless.” It is that **cerebrospinal fluid p-tau217 is not a validated surrogate for clinical decline**, which is a real and important piece of knowledge, obtained at great expense, that the field did not have before and now does. The same holds for semaglutide and valacyclovir: those trials tell us that GLP-1 agonism does not slow decline in early Alzheimer’s disease, and that antiviral suppression in herpes-seropositive patients is not merely useless but accompanied by faster cognitive worsening [18,17]. Both results constrain live hypotheses, and the second constrains one rather sharply.

Now apply the distinction to autophagy. **The autophagy programme lacks all three.** It has no target-engagement marker: no trial can say whether the drug acted on mTOR, AMPK, TFEB or the lysosome in the patient’s brain. It has no pathway marker: that is this paper’s subject. It has no surrogate, which is unsurprising since one cannot validate a surrogate one cannot measure.

This yields the paper’s central methodological claim. **An informative failure is expensive; an uninformative failure is worse than expensive, because it is repeatable.** When posdinemab failed, the field could remove a hypothesis from the board. When a rapamycin trial in Alzheimer’s disease returns null, nothing can be removed, because the trial cannot distinguish between “autophagy induction does not help,” “the drug did not reach the brain,” “it reached the brain and did not induce autophagy,” “it induced autophagy and the induction was harmful because completion is what fails,” and “the dose or duration was wrong.” Five hypotheses enter and five leave. The programme can, and does, run this experiment indefinitely without converging.

This is why the missing readout is not one deficiency among several. It is the deficiency that prevents the others from being diagnosed.

21. Specification: What a Deployable Human Flux Assay Must Measure

We now state the requirement formally. An assay adequate to dose an autophagy-directed drug in Alzheimer's disease must satisfy six conditions. They are not equally hard, and we rank them by difficulty.

R1 — It must report a rate. From Section 11: it must supply a second state, by perturbation of known sign, by an internally generated reference, or by time resolution of a label. An assay that returns a single concentration fails at the first requirement regardless of its precision, reproducibility or biological plausibility.

R2 — Its second state must be tolerable. This is the binding constraint and the reason the laboratory assays have not crossed over. The perturbation must be either non-pharmacological (a label), or a drug whose action is bounded, reversible and independently justified. A systemic lysosomal poison is excluded.

It is worth being exact about what “tolerable” permits, because the requirement is less restrictive than it first appears and the distinction determines which architectures survive. Three classes of second state are acceptable. A **metabolic label** perturbs nothing: a stable isotope is chemically identical to its abundant counterpart and its administration is an observation, not an intervention; this is why the stable-isotope labelling approach in this disease has been used in cognitively unimpaired volunteers [102,20]. A **therapeutic drug the participant is receiving anyway** supplies a second state for free — if a participant is being given an autophagy modulator for therapeutic reasons, measuring before and during treatment is a within-person contrast that costs nothing in additional risk. And an **ex vivo perturbation** of cells removed from the participant is unrestricted, since the poison never enters the person; this is the entire basis of architecture 22.3.

What is excluded is only the *in vivo systemic block*, and it is excluded twice over: once on safety, and once — the more interesting objection — because in a brain near degradative failure a block is not a small perturbation and the resulting estimate would not be of the unperturbed rate.

R3 — It must resolve the degradative step. The account under test locates the lesion at the terminal, acidification-dependent step [45,22], and the direction of the human data is that induction is preserved or increased while completion fails [3]. An assay that reported total pathway throughput without distinguishing the two would miss the lesion's defining feature, and would, in a

brain with both raised induction and impaired clearance, return a misleadingly normal number.

R4 — It must be attributable to a cell type, or at minimum to neurons versus glia.

The compensatory transcriptional response in Alzheimer tissue appears to be predominantly glial [3] while the account's lesion is neuronal, and the relevant manipulations have cell-type-specific consequences [32,30]. A bulk measurement that averages a compensating glial population with a failing neuronal one may show no net abnormality in a brain in which both are severely abnormal.

R5 — It must report brain. A peripheral measurement of autophagy in leucocytes is a measurement of autophagy in leucocytes. It may be a useful covariate; it is not the organ.

R6 — It must have within-person test-retest reproducibility sufficient for dose titration. This is the requirement most often forgotten and it is quantitative. If the assay is to be used as the rapamycin pharmacokinetic study used blood levels — to select a dose — then its within-person coefficient of variation must be small relative to the effect size the drug is expected to produce. The rapamycin pilot reported interindividual pharmacokinetic coefficients of variation of 0.28–0.40 and identified a sampling timepoint that minimised them [51]; an analogous characterisation would be required of any flux assay before it could carry the same weight.

A seventh consideration is not a requirement of the assay but a constraint on its interpretation, and it follows from Section 3. Because lysosomal pH is a regulated set-point held between opposing conductances rather than a quantity to be minimised [48,49], an assay guiding a terminal-step intervention must be able to report deviation in *either* direction. An assay calibrated on the assumption that more degradation is always better would mis-dose the very intervention it was built to guide.

22. Four Candidate Architectures

We consider four routes to R1–R6, rank them, and specify what each would require. One of them, we argue, is already assembled from published human methods that have not been joined.

22.1 Metabolic labelling of an autophagy-obligate cargo

This is the architecture we rank first, and the reason is that both of its halves already exist in humans.

The principle. Administer a stable isotope label — the established approach in this disease uses $^{13}\text{C}_6$ -leucine — and follow the labelled fraction of a specific protein in cerebrospinal fluid over time. The rise and fall of labelled protein gives the synthesis and clearance rates of that protein in the central nervous system directly. This is stable-isotope labelling kinetics, and it is not a proposal: it is an established human method in this exact disease, applied to tau, where it showed that tau is a long-lived protein whose kinetics differ between neurons and the human central nervous system and which is altered in Alzheimer’s disease [102] (with a published erratum [103]), and to the soluble ectodomains of the amyloid precursor protein, whose isoform-specific metabolism was shown to be quantifiable by the same approach [104].

The missing step, and why it is now available. To convert a protein turnover measurement into an autophagy measurement, the protein must be one whose degradation is autophagy-dependent. That link has been established for neurofilament subunits. Blocking autophagy — pharmacologically with 3-methyladenine, or genetically by depleting ATG5 — raises NEFL, NEFM and NEFH levels substantially, while activating autophagy with rapamycin lowers them; intracerebroventricular infusion of 3-methyladenine raises neurofilament subunit levels in mouse brain *in vivo*; and neurofilament proteins have been localised inside autophagosomes, autolysosomes and lysosomes by tomographic confocal microscopy, immunoelectron microscopy and biochemical fractionation [19].

The convergence. Neurofilament light kinetics in the human central nervous system have now been measured directly by stable-isotope labelling in 35 participants across three cohorts — primary tauopathies, a cognitively unimpaired and mild-cognitive-impairment cohort, and suspected normal-pressure hydrocephalus — establishing that NfL turnover is slow, comparable to tau, and showing an unexpected 3–6 day delay between synthesis and release into the extracellular medium in human induced-pluripotent-stem-cell-derived neurons, which the authors interpret as evidence for controlled rather than purely passive release [20].

The proposal. Measure neurofilament turnover by stable-isotope labelling kinetics in the same participant before and during administration of an autophagy-modulating drug. The change in the clearance rate constant is a rate, in brain, of a protein whose degradation is autophagy-dependent. The drug is the second state (satisfying R1 and R2), the measurement is central (R5), and the protein is neuronal and axonal rather than glial (R4).

Against it, honestly. Four objections are serious. *First*, neurofilament is not an exclusive autophagy substrate: calpain cleavage during Wallerian degeneration is well established and is named in the same work [19], and proteasomal contributions are likely, so the measured clearance rate is autophagy-dependent rather than autophagy-specific. The assay would report a composite whose autophagic fraction requires separate estimation. *Second*, R3 is not satisfied: a turnover rate reports total throughput, not the terminal step, and so would not by itself distinguish reduced induction from failed completion — the very distinction Section 3 says matters most. A two-drug design, using an agent known to act at initiation and one known to act at the lysosome, would be required to decompose it. *Third*, the release delay [20] means the cerebrospinal fluid appearance rate conflates degradation with regulated release, which is the provenance problem of Section 13 arriving in the assay itself. *Fourth*, in a disease whose defining feature is neuronal death, the neurofilament pool is being fed by a third route that has nothing to do with turnover.

None of these is fatal, and each is addressable by design. They should be stated in any protocol, and we state them here.

22.2 A bounded pharmacological clamp

The principle. Restore the laboratory assay’s logic with an acceptable agent: give a drug that partially and reversibly inhibits lysosomal degradation, and measure the accumulation of an autophagosomal marker during the window.

In its favour. It is the only architecture that directly satisfies R1 and R3 together, because a lysosomal block is by construction a perturbation of the terminal step.

Against it. R2 is violated in spirit even if a tolerable agent could be found. Worse, the pharmacology is non-linear in exactly the population of interest: a partial lysosomal block applied to a neuron already near degradative failure is a large perturbation, not a small one, and the estimate returned would not be of the unperturbed rate. And the natural experiment that most resembles it has already been run without instrumentation: eighteen months of hydroxychloroquine in early Alzheimer’s disease, with no lysosomal measurement [44].

We rank this last and mention it chiefly to be explicit that the obvious solution has been considered and rejected.

22.3 A peripheral cell challenge assay

The principle. Isolate monocytes or peripheral blood mononuclear cells from a participant, apply a bafilomycin or chloroquine clamp *ex vivo*, and measure LC3-II accumulation. This is the standard laboratory assay, performed outside the body, and it returns a true flux number.

In its favour. It satisfies R1, R2 (the poison is applied to cells in a dish, not a person), R3 and R6. It is technically mature, cheap relative to imaging or labelling, and repeatable at every visit.

Against it. It fails R5 outright and R4 partially. It measures autophagy in a circulating myeloid cell, not in a cortical neuron, and there is no established quantitative relationship between the two. Its defensible use is therefore as a *covariate and a systemic pharmacodynamic marker* — evidence that a given oral dose produces a measurable autophagic effect in some human cell — rather than as a brain readout.

Two considerations raise its value above the dismissal it usually receives. The first is that myeloid cells may not be the wrong compartment: carboxylesterase-1, the enzyme that activates the NLRP3 prodrug inside brain parenchyma, is expressed by a subpopulation of cells largely positive for the microglial marker IBA1 [15], which is a demonstration that a myeloid brain compartment can be the pharmacologically relevant one. The second is that a peripheral flux assay would at least prevent the specific failure mode of the rapamycin phase 1 [8]: a trial that could show no drug in cerebrospinal fluid *and* no autophagic effect in blood is in a different epistemic position from one that could show only the former.

22.4 An imaging readout of the degradative step

The principle. A positron emission tomography tracer that is trapped in a pH-dependent manner would report the acidification state of the degradative compartment, and its washout kinetics under a paired-scan design would approach a rate.

Requirements. Such a tracer would need a pKa positioned to discriminate across the pathological range (roughly pH 4.5 to 5.5), sufficient brain penetration, rapid clearance of the untrapped fraction, and a specific-to-non-specific ratio adequate for parametric imaging. No such agent is in human use. The precedent for what a pharmacodynamic imaging package looks like exists in the preclinical NLRP3 work, where four tracers plus contrast MRI were used in the same animal to demonstrate multi-node engagement [86], and recent work on lysosomal surface chemistry using DNA nanodevices indicates the measurement is being approached from the reagent side [105].

Assessment. This architecture would satisfy R3, R4 (regionally, if not by cell type), R5 and potentially R1, and is the only one that would give a spatial map. It is also the furthest from existing, and its development timeline is measured in years rather than trials. We rank it second on merit and

third on availability.

22.5 Ranking

Architecture	R1 rate	R2 safe	R3 end step	R4 cell	R5 brain	R6 repeat	Availability
Labelling of an autophagy cargo	Yes	Yes	Partial	Partial	Yes	Unknown	Published in two halves; never joined
Imaging of the degradative step	Potential	Yes	Yes	Regional	Yes	Unknown	No tracer exists
Peripheral cell challenge	Yes	Yes	Yes	No	No	Likely	Available now
Bounded pharmacological clamp	Yes	No	Yes	No	Yes	Unknown	Excluded on safety

The practical recommendation that follows is a combination rather than a winner. A programme could, today, run a peripheral challenge assay as a systemic pharmacodynamic marker and a labelling study as a central one, in the same participants, at a cost that is small relative to a phase 2 trial and trivial relative to a phase 3. Nothing in the current state of the art prevents it. What has prevented it is that the requirement has not been stated.

23. How a Flux-Dosed Trial Would Differ

Suppose the assay existed. Four features of trial design would change, and it is worth specifying them, because the changes are not cosmetic.

Dose would be selected against effect, not exposure. The current procedure is visible in the record of Section 4: a dose is chosen from tolerability and pharmacokinetics, administered, and its adequacy inferred from drug concentration [8,51]. The flux-dosed procedure is the one the inflammasome field already uses: choose the dose that produces a specified fractional change in the pharmacodynamic marker [95]. The difference is not marginal. Rapamycin’s dose–response for autophagy induction is non-monotonic in several systems and its immunosuppressive and metabolic effects rise with exposure, so the therapeutic window — if there is one — is defined by a pharmacodynamic optimum that current methods cannot locate.

Participants would be enrolled by lesion rather than by syndrome. Every trial in Table 1 enrolled on clinical stage plus, in the better ones, an amyloid or tau biomarker. None enrolled on the abnormality the drug corrects. If autophagic flux were measurable, the natural design would enrol participants with demonstrated impairment and stratify on its severity. This matters more here than in most fields because of the direction of the lesion: a patient whose induction is intact and whose completion is failing is a different pharmacological subject from one in whom both are intact, and current trials contain unknown mixtures of both.

A direction check would become mandatory. This is the design change with the largest expected effect and it follows directly from Section 3. If induction is already elevated and completion is failing [3], then an inducer may accumulate undegraded compartments rather than clear them, and the animal evidence is consistent with this being harmful rather than merely futile [47]. A flux-instrumented phase 1 could test the sign of the drug’s effect before a phase 2 commits. The current procedure cannot, and a trial that cannot detect that its drug is making the lesion worse will read the resulting clinical outcome as evidence against the mechanism.

The phase 2 primary endpoint would be the rate. This is the change that would most alter the economics. A dementia trial powered on a clinical scale requires hundreds of participants and 18 to 24 months; the two largest trials in Section 19 required 3,808 participants and 156 weeks [18] and 523 participants and two years respectively. A trial powered on a within-person pharmacodynamic change requires tens of participants and weeks, which is what the inflammasome programme achieved with fourteen subjects over 28 days [14]. The field could then afford to be wrong quickly, which is the condition under which programmes converge.

A fifth difference is a matter of interpretation rather than design, and it is the one that motivates the whole exercise. A flux-instrumented negative trial would *mean something*. The five-hypothesis problem of Section 20 would collapse to one or two, and the mechanism could be removed from the board or promoted on the evidence. That is the object.

A concrete first study

Abstract requirements are easy to assent to and hard to act on, so we specify the smallest study that would test the paper’s constructive proposal. It is not a therapeutic trial and should not be run as one.

Design. Open-label, within-person, two-period. Participants serve as their own controls.

Participants. Twelve to sixteen people with biomarker-confirmed mild cognitive impairment or mild dementia due to Alzheimer’s disease, and six to eight cognitively unimpaired age-matched controls. This is the sample size at which the inflammasome programme reached a target-engagement conclusion [14] and it is adequate for a within-person pharmacodynamic question.

Period 1 — baseline kinetics. Stable-isotope labelling with $^{13}\text{C}_6$ -leucine, with serial cerebrospinal fluid sampling by indwelling catheter over 24–36 hours and continued intermittent sampling thereafter, following the design established for tau kinetics in this population [102] and applied to neurofilament light [20]. Primary analyte: neurofilament light. Secondary analytes, measured in the same samples at no additional labelling cost: tau, to provide a within-study comparator with published kinetics, and the soluble amyloid precursor protein ectodomains, whose isoform-specific kinetics are quantifiable by the same method [104].

Interval. Four to eight weeks, during which the participant receives an autophagy-modulating agent at a dose with established central exposure. Metformin is the pragmatic choice for a first study because its cerebrospinal fluid concentration in this population has been measured [52]; rapamycin is the mechanistically preferred choice but its central exposure at tolerated doses is in doubt [8].

Period 2 — on-drug kinetics. Repeat the labelling protocol.

Primary endpoint. Within-person change in the fractional clearance rate of neurofilament light.

Pre-specified secondary endpoints. Change in tau clearance rate, as a specificity control: if the drug moves neurofilament and tau clearance identically, the result is more consistent with a general change in protein turnover or in cerebrospinal fluid dynamics than with an autophagy-specific effect. Peripheral monocyte autophagic flux by ex vivo clamp (architecture 22.3) at both timepoints, as a systemic pharmacodynamic covariate.

What each outcome would mean. A drug-associated acceleration of neurofilament clearance, not mirrored in tau, would be the first measurement of a human autophagic rate response to a drug, and would license the stratified trial of Section 24, prediction 6. No change, with peripheral flux demonstrably moved, would indicate that the drug does not reach or does not act on the central compartment — a definite negative that the rapamycin phase 1 could not obtain. No change anywhere would indicate inadequate exposure or an inactive dose, which is also definite. A *deceleration* of neurofilament clearance would be the most interesting result of all, since it would be the predicted signature of an inducer worsening a completion-limited system, and would justify halting the therapeutic programme it was meant to enable.

What it would cost. Two labelling visits with serial cerebrospinal fluid sampling in approximately twenty participants is a study of the scale routinely run in academic centres for tau and amyloid kinetics. It is a small fraction of the cost of any of the trials in Table 1 and an extremely small fraction of the trials in Section 19. The obstacle has never been the money.

24. Predictions

The account above is falsifiable in specific ways. We list nine, with what each would show.

1. **No autophagy or lysosomal pharmacodynamic endpoint will be found in any completed Alzheimer’s disease trial of an autophagy-rationale compound.** We have audited eleven (Table 1) and found none. A single counterexample of a completed trial with a pre-specified autophagic-flux endpoint would substantially weaken the paper’s empirical claim.
2. **A measured level of an autophagosome marker in human cerebrospinal fluid will not correlate with a directly measured degradative rate in the same individuals.** This becomes testable as soon as any rate measurement exists, and is the cleanest test of Section 11’s arithmetic in human material.
3. **Cerebrospinal fluid LC3A and LC3B will correlate more strongly with markers of neuronal injury than with any lysosomal-functional measure.** If provenance is dominated by cell damage rather than by regulated release, the correlation structure of the existing ADNI measurements [9] should already show this against neurofilament light. This can be tested today, in existing data, at no cost.
4. **Cerebrospinal fluid or serum TFEB will track glial rather than neuronal markers.** Predicted by the human tissue finding that MiTF/TFE activation is glia-weighted [3], and testable in the existing serum and cerebrospinal fluid TFEB dataset [10].
5. **In a flux-instrumented study, mTOR inhibition at clinically tolerated exposures will raise autophagosome formation rate without raising degradative rate.** This is the paper’s substantive mechanistic prediction. If instead mTOR inhibition raises both — through TFEB-mediated lysosomal biogenesis [28,27] — then the induction-versus-completion framing of Section 3 is too simple for pharmacological purposes and the therapeutic pessimism about inducers should be revised.
6. **Stratifying an autophagy-drug trial by a valid flux measure will produce a treatment-by-stratum interaction, with benefit confined to participants whose degradative capacity is preserved.** This is the operational form of the direction argument and the single most consequential prediction here.
7. **Neurofilament turnover measured by stable-isotope labelling will slow under pharmacological autophagy inhibition and accelerate under induction, in the**

direction established in cells and mouse brain [19]. If it does not, architecture 22.1 fails and the paper's constructive proposal fails.

- 8. Trials of NLRP3 inhibitors in neurodegeneration will reach phase 2 go/no-go decisions faster, and at smaller sample size, than trials of autophagy modulators of equivalent preclinical support.** This is a prediction about the sociology of development, and it is measurable against the record over the next five years.
- 9. A brain-penetrant NLRP3 inhibitor may fail clinically in neurodegeneration while suppressing cerebrospinal fluid interleukin-1 β .** We predict this is more likely than not on current evidence, and we state it because the paper's argument does not require the inflammasome programme to succeed. If it fails with its marker suppressed, the failure will be informative, which is the entire point of Sections 16 to 20.

25. Strength of Evidence

Every substantive claim in this paper is graded below. **Established** means multiple independent laboratories, or a single result with internal replication and an intervention. **Probable** means one good primary result, or a consistent body of work, not yet independently replicated in the form stated. **Inferred** means assembled here from published parts, with no experiment addressing the composite. **Speculative** means consistent with the evidence and not tested by it.

Table 2 — The four architectures against the six requirements.

#	Claim	Grade	Load-bearing evidence	What would move it
1	Early-endosome enlargement is the earliest disease-specific structural abnormality in human Alzheimer's disease	Established	Human sporadic and Down syndrome tissue, <i>APOE</i> -modified, absent in comparison diseases [1]; gene-expression confirmation in captured neurons [3]	—
2	Autophagic vacuoles are the predominant organelle in the dystrophic neurites of the neuritic plaque	Established	Immunogold electron microscopy on human antemortem biopsy [2]	—
3	In human Alzheimer neurons, autophagy induction is increased while substrate clearance fails	Probable	Laser-captured CA1 neurons, transcript, protein and morphological readouts in one study [3]; consistent mouse reporter data [45]	Independent replication in human tissue with a flux-capable method
4	The pathogenic species acting on the endosome is the β -carboxy-terminal fragment, not amyloid- β	Established	Isogenic human iPSC panel from an independent laboratory [21]; downstream v-ATPase inhibition by the phosphorylated fragment [22]	—
5	Presenilin 1 is required for lysosomal acidification	Contested	Original report [40] contradicted [41] and not independently replicated as stated	Independent replication, or a decisive negative

#	Claim	Grade	Load-bearing evidence	What would move it
6	Correcting the endolysosomal lesion improves pathology and behaviour in rodents	Established	Four independent nodes: protease de-inhibition [5,25], mTOR [6], TFEB [7,30,32], ULK1 [11]	—
7	No completed human trial of an autophagy-rationale compound in Alzheimer's disease measured autophagy	Probable	Audit of eleven studies, Table 1	A single documented counterexample
8	Autophagic flux is formally unidentified from single-state level measurements	Established	Elementary; and the operational basis of the field's own guidelines [76,77]	—
9	Every existing human autophagy readout is a single-state level	Probable	Inventory of five classes [9,10,11,12,13]	A published human flux measurement
10	No human PET tracer reports autophagic flux or lysosomal pH	Probable	Absence in the indexed literature; adjacent tracers report metabolism, microglia or cargo	A first-in-human tracer report
11	The flux assay has not crossed to humans because its second state requires a poison, a transgene, or the brain	Inferred	Assembled from the three assay families [77,78,2]	A fourth assay family not requiring any of the three
12	NLRP3's development path was enabled by its output being a secreted cytokine	Inferred	Structural fact about the target plus the dose-selection logic in the trial record [95,14]	—
13	NT-0796 achieves CSF exposure and lowers CSF interleukin-1 β in Parkinson's disease	Probable	Open-label phase 1b, n=14, no placebo [14]; mechanism corroborated [15]	A randomised placebo-controlled central replication
14	NT-0796 (ruvonoflast) lowers systemic inflammatory markers against placebo	Established	Randomised double-blind phase 1b, n=63, primary endpoint met [16]	—
15	No NLRP3 inhibitor has demonstrated clinical benefit in any neurodegenerative disease	Established	Absence in the record; stated in a 2026 review [96]	A positive randomised clinical outcome

#	Claim	Grade	Load-bearing evidence	What would move it
16	CSF p-tau217 is not a validated surrogate for clinical decline	Probable	Posdinemab phase 2b (conference disclosure, not peer-reviewed); target engagement itself is established [99,101]	Peer-reviewed publication of the phase 2b data
17	Valacyclovir and oral semaglutide do not slow decline in early symptomatic Alzheimer's disease, and valacyclovir is associated with greater cognitive worsening	Established	Two adequately powered randomised trials; valacyclovir ADAS-Cog between-group difference 3.93 (1.03–6.83), $P = .01$ [17,18]	—
18	An uninformative failure is epistemically worse than an informative one because it is repeatable	Inferred	The five-hypothesis argument of §20	—
19	Neurofilament subunits are degraded by autophagy in vivo	Probable	Pharmacological, genetic and morphological convergence in cells and mouse brain from one laboratory [19]	Independent replication; quantification of the autophagic fraction of total turnover
20	Neurofilament turnover can be measured in the human CNS by stable-isotope labelling	Probable	First human quantification, $n=35$ across three cohorts [20]	Independent replication; test–retest characterisation
21	A labelling study of neurofilament turnover on and off an autophagy drug would yield a human autophagic rate	Inferred	Rows 19 and 20 joined; never performed	Performing it
22	Autophagy inducers may worsen a brain in which completion has failed	Speculative	Direction of the human lesion [3]; <i>Atg7</i> deletion raising intracellular amyloid and accelerating degeneration [47]	A flux-stratified trial showing a treatment-by-stratum interaction
23	Lysosomal pH is a set-point whose deviation in either direction is pathogenic, not a quantity to minimise	Probable	Recent reviews of neuronal and microglial lysosomal pH regulation [48,49]	A therapeutic demonstration that monotonic acidification is beneficial

26. Limitations

The audit is of the indexed literature, and absence of evidence is doing work. Several of the paper’s claims — that no completed trial measured autophagy, that no human flux measurement exists, that no PET tracer reports lysosomal pH — are negative claims established by not finding something. We have searched the indexed literature systematically and found nothing; we cannot exclude an unpublished measurement, a measurement buried in a supplementary table we did not retrieve, or a trial whose registry entry names an endpoint its publication does not report. The claim should be read as: *if such measurements exist, they have not entered the literature in a form that has influenced the field’s interpretation of its own trials.*

Two of the four 2025–26 results are not peer-reviewed. The posdinemab and diranersen accounts in Section 19 rest on conference presentation and company disclosure. We have said so at each point of use, cited no publication for either, and graded the conclusion drawn from posdinemab as probable rather than established (Table 3, row 16). If either account is materially revised on publication, Section 19’s illustrative force changes; Section 20’s argument does not, since it rests on the VALAD and evoke results, which are peer-reviewed.

One foundational citation carries a recent erratum whose content we have not examined. The 2011 report linking TFEB to lysosomal biogenesis [28] had an erratum published in 2025 [56]. We have not been able to retrieve the erratum’s content. We have therefore not relied on any specific quantitative claim from that paper, and the propositions for which we cite it are independently supported [27,29]. Readers building on the primary paper should consult the erratum. A further seven cited works carry errata, listed in the accompanying verification record.

The neurofilament proposal rests on a single laboratory for one of its two halves. The demonstration that autophagy degrades neurofilament proteins in vivo comes from one group [19], and to our knowledge has not been independently replicated. We have graded it probable rather than established (row 19). If it does not replicate, architecture 22.1 loses its substrate and the remaining candidates are weaker.

The NLRP3 contrast may be flattered by selection. We chose the inflammasome because it is the closest structural analogue with a readout. A sceptic could reply that the relevant difference is not the readout but the target’s tractability, the commercial pull of a class with cardiovascular and inflammatory indications, or simple chronology. We think the secreted-output argument is the dominant one and have graded it as inferred (row 12) rather than established, and Section 18 states the contrast’s limits at full strength.

The paper does not establish that autophagy is a valid therapeutic target. It establishes that the question has not been asked properly. Those are different claims, and the second does not imply the first. It remains entirely possible that a valid flux assay would be built, applied, and show that autophagy modulation does nothing useful in Alzheimer’s disease. That outcome would be a success of the programme proposed here, not a failure of it.

The flux framing may be too coarse for the biology. Autophagy is not one process. Macroautophagy, chaperone-mediated autophagy and the selective forms — mitophagy, aggrephagy, lipophagy — have different machinery, different cargo and different relationships to the disease, and the mitophagy literature in particular has developed its own human markers [10,11,33]. A single scalar “flux” may be the wrong abstraction. We use it because it is the abstraction the therapeutic programme has used, and because the arithmetic of Section 11 applies to each selective pathway separately.

27. Conclusion

The proposition that Alzheimer's disease involves a failure of the cell's disposal system is among the better-supported mechanistic claims in the field. The proposition that this failure is therapeutically addressable has been tested, in animals, dozens of times, and the animals have almost always improved. The proposition has never been tested in a human being.

That is a strong statement and we mean it literally. Eleven completed human studies have administered compounds whose rationale is autophagic or endolysosomal to patients with Alzheimer's disease or a related neurodegeneration. They have measured drug concentrations, cognition, amyloid, phosphorylated tau, glial fibrillary acidic protein, hippocampal volume, cerebral blood flow, microRNAs, and survival. Not one of them measured whether the disposal system worked any better after treatment than before. The trials were tests of clinical benefit; they were never tests of the mechanism, and they cannot be read as such.

The reason is specific and it is not remediable by effort. Autophagic flux is a rate. Every human readout is a level. A level is the quotient of two rates and identifies neither. The laboratory solves this by adding a lysosomal poison, by expressing a genetically encoded internal control, or by looking at the tissue under an electron microscope, and each of those solutions requires something a living patient cannot supply. So the field has proceeded with the instruments it had, which are instruments for measuring the accumulated products of a process — amyloid, tau, atrophy — in a hypothesis whose entire content concerns the rate at which products are removed.

The consequence is a programme that cannot learn from its own results. When a tau antibody abolishes its target in cerebrospinal fluid and the patients decline anyway, the field learns that the quantity moved was not on the causal path; that is expensive knowledge and it is knowledge. When a rapamycin trial in Alzheimer's disease returns null, five hypotheses remain standing, and the next trial will return the same five. A field that produces uninformative failures can run forever without converging, and this one has been running for fifteen years.

The contrast with the inflammasome is not that the inflammasome has succeeded. It has not; no NLRP3 inhibitor has shown clinical benefit in any neurodegenerative disease, and it may well be that none will. The contrast is that when the inflammasome programme fails, it will fail with its central marker suppressed and its dose chosen against effect, and the field will be able to say what was refuted. That capacity was bought by an accident of the target's structure: NLRP3's output is secreted, and autophagy's output is an absence.

What follows is not a call for more mechanism. The mechanism is saturated. It is a call to build one instrument. We have specified what it must do — report a rate, without a poison, resolving the degradative step, attributable to a cell type, in brain, reproducibly enough to titrate a dose — and we have argued that one architecture satisfying most of these requirements is already lying in the literature in two pieces. Neurofilament subunits are degraded by autophagy *in vivo*. Neurofilament turnover in the human central nervous system can now be measured. Nobody has put a drug between the two measurements.

Until someone does, the therapeutics of the autophagic bottleneck will continue to consist of well-designed trials of interesting compounds, conducted in the right patients, at unknown doses, against an unmeasured target, and read as though their results meant something.

References

1. Cataldo AM, Peterhoff CM, Troncoso JC, et al. Endocytic pathway abnormalities precede amyloid beta deposition in sporadic Alzheimer's disease and Down syndrome: differential effects of APOE genotype and presenilin mutations. *Am J Pathol* 2000;157(1):277-86. PMID 10880397.
2. Nixon RA, Wegiel J, Kumar A, et al. Extensive involvement of autophagy in Alzheimer disease: an immuno-electron microscopy study. *J Neuropathol Exp Neurol* 2005;64(2):113-22. PMID 15751225.
3. Bordi M, Berg MJ, Mohan PS, et al. Autophagy flux in CA1 neurons of Alzheimer hippocampus: Increased induction overburdens failing lysosomes to propel neuritic dystrophy. *Autophagy* 2016;12(12):2467-2483. PMID 27813694.
4. Bellenguez C, Küçükali F, Jansen IE, et al. New insights into the genetic etiology of Alzheimer's disease and related dementias. *Nat Genet* 2022;54(4):412-436. PMID 35379992.
5. Yang DS, Stavrides P, Mohan PS, et al. Reversal of autophagy dysfunction in the TgCRND8 mouse model of Alzheimer's disease ameliorates amyloid pathologies and memory deficits. *Brain* 2011;134(Pt 1):258-77. PMID 21186265.
6. Spilman P, Podlitskaya N, Hart MJ, et al. Inhibition of mTOR by rapamycin abolishes cognitive deficits and reduces amyloid-beta levels in a mouse model of Alzheimer's disease. *PLoS One* 2010;5(4):e9979. PMID 20376313.
7. Polito VA, Li H, Martini-Stoica H, et al. Selective clearance of aberrant tau proteins and rescue of neurotoxicity by transcription factor EB. *EMBO Mol Med* 2014;6(9):1142-60. PMID 25069841.
8. Gonzales MM, Garbarino VR, Kautz TF, et al. Rapamycin treatment for Alzheimer's disease and related dementias: a pilot phase 1 clinical trial. *Commun Med (Lond)* 2025;5(1):189. PMID 40394335.
9. Guo HH, Su D, Song B, et al. CSF microtubule-associated protein 1 light chain 3A and 3B levels are associated with tau pathology and Alzheimer's disease risk through amyloid deposition and microglial signaling. *J Neuroinflammation* 2026;23(1):94. PMID 41795086.
10. Veverová K, Laczó J, Katonová A, et al. Alterations of human CSF and serum-based mitophagy biomarkers in the continuum of Alzheimer disease. *Autophagy*

2024;20(8):1868-1878. PMID 38695174.

- 11.** Pan JP, Wang PJ, Zhang J, et al. Reduced ULK1 links impaired autophagy and mitophagy to Alzheimer's disease pathology. *Nat Aging* 2026;6(5):1079-1102. PMID 42135576.
- 12.** Armstrong A, Mattsson N, Appelqvist H, et al. Lysosomal network proteins as potential novel CSF biomarkers for Alzheimer's disease. *Neuromolecular Med* 2014;16(1):150-60. PMID 24101586.
- 13.** Rubino E, Boschi S, Roveta F, et al. Investigating p62 Concentrations in Cerebrospinal Fluid of Patients with Dementia: A Potential Autophagy Biomarker In Vivo?. *Brain Sci* 2022;12(10):1414. PMID 36291347.
- 14.** Clarke N, Thornton P, Reader V, et al. Anti-Neuroinflammatory and Anti-Inflammatory Effects of the NLRP3 Inhibitor NT-0796 in Subjects with Parkinson's Disease. *Mov Disord* 2025;40(10):2199-2208. PMID 40792655.
- 15.** Digby Z, Diamond C, Wescott H, et al. Ester Prodrug NLRP3 Inflammasome Inhibitor NT-0796 is Brain Active due to Activation by Local Expression of Carboxylesterase-1. *ACS Chem Neurosci* 2026;17(3):599-609. PMID 41586719.
- 16.** Ray KK, Clarke N, Thornton P, et al. Anti-Inflammatory Effects of Oral NLRP3 Inhibition With Ruvonoflast Among Individuals at Elevated Cardiovascular Risk. *J Am Coll Cardiol* 2026 (online ahead of print; doi:10.1016/j.jacc.2026.05.014). PMID 42187339.
- 17.** Devanand DP, Wisniewski T, Razlighi Q, et al. Valacyclovir Treatment of Early Symptomatic Alzheimer Disease: The VALAD Randomized Clinical Trial. *JAMA* 2026;335(6):511-522. PMID 41405855.
- 18.** Cummings JL, Atri A, Sano M, et al. Efficacy and safety of oral semaglutide 14 mg (flexible dose) in early-stage symptomatic Alzheimer's disease (evoke and evoke+): two phase 3, randomised, placebo-controlled trials. *Lancet* 2026;407(10544):2167-2179. PMID 41865758.
- 19.** Rao MV, Darji S, Stavrides PH, et al. Autophagy is a novel pathway for neurofilament protein degradation in vivo. *Autophagy* 2023;19(4):1277-1292. PMID 36131358.
- 20.** Leckey CA, Giovannucci TA, Coulton JB, et al. Stable isotope labelling kinetics of neurofilament light. *Brain Commun* 2025;7(6):fcaf468. PMID 41416248.
- 21.** Kwart D, Gregg A, Scheckel C, et al. A Large Panel of Isogenic APP and PSEN1 Mutant Human iPSC Neurons Reveals Shared Endosomal Abnormalities Mediated by APP β -CTFs, Not A β . *Neuron* 2019;104(2):256-270.e5. PMID 31416668.

22. Im E, Jiang Y, Stavrides PH, et al. Lysosomal dysfunction in Down syndrome and Alzheimer mouse models is caused by v-ATPase inhibition by Tyr682-phosphorylated APP β CTF. *Sci Adv* 2023;9(30):eadg1925. PMID 37494443.
23. Mizushima N, Levine B. Autophagy in Human Diseases. *N Engl J Med* 2020;383(16):1564-1576. PMID 33053285.
24. Menzies FM, Fleming A, Caricasole A, et al. Autophagy and Neurodegeneration: Pathogenic Mechanisms and Therapeutic Opportunities. *Neuron* 2017;93(5):1015-1034. PMID 28279350.
25. Yang DS, Stavrides P, Saito M, et al. Defective macroautophagic turnover of brain lipids in the TgCRND8 Alzheimer mouse model: prevention by correcting lysosomal proteolytic deficits. *Brain* 2014;137(Pt 12):3300-18. PMID 25270989.
26. Cai J, Xie D, Kong F, et al. Effect and Mechanism of Rapamycin on Cognitive Deficits in Animal Models of Alzheimer's Disease: A Systematic Review and Meta-analysis of Preclinical Studies. *J Alzheimers Dis* 2024;99(1):53-84. PMID 38640155.
27. Sardiello M, Palmieri M, di Ronza A, et al. A gene network regulating lysosomal biogenesis and function. *Science* 2009;325(5939):473-7. PMID 19556463.
28. Settembre C, Di Malta C, Polito VA, et al. TFEB links autophagy to lysosomal biogenesis. *Science* 2011;332(6036):1429-33. PMID 21617040.
29. Palmieri M, Impey S, Kang H, et al. Characterization of the CLEAR network reveals an integrated control of cellular clearance pathways. *Hum Mol Genet* 2011;20(19):3852-66. PMID 21752829.
30. Xiao Q, Yan P, Ma X, et al. Neuronal-Targeted TFEB Accelerates Lysosomal Degradation of APP, Reducing A β Generation and Amyloid Plaque Pathogenesis. *J Neurosci* 2015;35(35):12137-51. PMID 26338325.
31. Bécot A, Pardossi-Piquard R, Bourgeois A, et al. The Transcription Factor EB Reduces the Intraneuronal Accumulation of the Beta-Secretase-Derived APP Fragment C99 in Cellular and Mouse Alzheimer's Disease Models. *Cells* 2020;9(5):1204. PMID 32408680.
32. Xiao Q, Yan P, Ma X, et al. Enhancing astrocytic lysosome biogenesis facilitates A β clearance and attenuates amyloid plaque pathogenesis. *J Neurosci* 2014;34(29):9607-20. PMID 25031402.
33. Fang EF, Hou Y, Palikaras K, et al. Mitophagy inhibits amyloid- β and tau pathology and reverses cognitive deficits in models of Alzheimer's disease. *Nat Neurosci* 2019;22(3):401-412. PMID 30742114.

34. Hou Y, Chu X, Park JH, et al. Urolithin A improves Alzheimer's disease cognition and restores mitophagy and lysosomal functions. *Alzheimers Dement* 2024;20(6):4212-4233. PMID 38753870.
35. Bourdenx M, Daniel J, Genin E, et al. Nanoparticles restore lysosomal acidification defects: Implications for Parkinson and other lysosomal-related diseases. *Autophagy* 2016;12(3):472-83. PMID 26761717.
36. Mecozzi VJ, Berman DE, Simoes S, et al. Pharmacological chaperones stabilize retromer to limit APP processing. *Nat Chem Biol* 2014;10(6):443-9. PMID 24747528.
37. Hara T, Nakamura K, Matsui M, et al. Suppression of basal autophagy in neural cells causes neurodegenerative disease in mice. *Nature* 2006;441(7095):885-9. PMID 16625204.
38. Kwart D, Gregg A, Scheckel C, et al. A Large Panel of Isogenic APP and PSEN1 Mutant Human iPSC Neurons Reveals Shared Endosomal Abnormalities Mediated by APP β -CTFs, Not A β . *Neuron* 2019;104(5):1022. PMID 31805257.
39. Jiang Y, Sachdeva K, Goulbourne CN, et al. Increased Neuronal Expression of the Early Endosomal Adaptor APPL1 Replicates Alzheimer's Disease-Related Endosomal and Synaptic Dysfunction with Cholinergic Neurodegeneration. *J Neurosci* 2025;45(29):e2331242025. PMID 40514243.
40. Lee JH, Yu WH, Kumar A, et al. Lysosomal proteolysis and autophagy require presenilin 1 and are disrupted by Alzheimer-related PS1 mutations. *Cell* 2010;141(7):1146-58. PMID 20541250.
41. Zhang X, Garbett K, Veeraraghavalu K, et al. A role for presenilins in autophagy revisited: normal acidification of lysosomes in cells lacking PSEN1 and PSEN2. *J Neurosci* 2012;32(25):8633-48. PMID 22723704.
42. Timsina J, Jiang C, McCartney DL, et al. GWAS meta-analysis of cerebrospinal fluid Alzheimer's biomarkers reveals loci regulating lipids, brain volume and autophagy. *Nat Commun* 2026;17(1):7385. PMID 42014397.
43. Cummings JL, Zhou Y, Yang Y, et al. Alzheimer's disease drug development pipeline: 2026. *Alzheimers Dement (NY)* 2026;12(2):e70251. PMID 42095064.
44. Van Gool WA, Weinstein HC, Scheltens P, et al. Effect of hydroxychloroquine on progression of dementia in early Alzheimer's disease: an 18-month randomised, double-blind, placebo-controlled study. *Lancet* 2001;358(9280):455-60. PMID 11513909.
45. Lee JH, Yang DS, Goulbourne CN, et al. Faulty autolysosome acidification in Alzheimer's disease mouse models induces autophagic build-up of A β in neurons, yielding senile plaques.

- Nat Neurosci* 2022;25(6):688-701. PMID 35654956.
- 46.** Lee JH, Nixon RA. Autolysosomal acidification failure as a primary driver of Alzheimer disease pathogenesis. *Autophagy* 2022;18(11):2763-2764. PMID 35947489.
- 47.** Nilsson P, Loganathan K, Sekiguchi M, et al. A β secretion and plaque formation depend on autophagy. *Cell Rep* 2013;5(1):61-9. PMID 24095740.
- 48.** Tang Y, Guo T, He H, et al. Beyond acid-base dyshomeostasis: Dynamic instability of neuronal lysosomal pH as a pathogenic mechanism and therapeutic target in neurological diseases. *Biochem Pharmacol* 2026;250(Pt 1):117990. PMID 42002234.
- 49.** Jaganathan R, Vijayakumar S, Chen Y, et al. New insights on microglial lysosomal acidification: A therapeutic target of neurodegenerative diseases. *Neurobiol Dis* 2026;228:107555. PMID 42546981.
- 50.** Caccamo A, Majumder S, Richardson A, et al. Molecular interplay between mammalian target of rapamycin (mTOR), amyloid-beta, and Tau: effects on cognitive impairments. *J Biol Chem* 2010;285(17):13107-20. PMID 20178983.
- 51.** Wallgren HA, Kivipelto M, Plavén-Sigray P, et al. Pharmacokinetic analysis of intermittent rapamycin administration in early-stage Alzheimer's Disease. *Geroscience* 2026;48(3):3295-3304. PMID 41046300.
- 52.** Koenig AM, Mechanic-Hamilton D, Xie SX, et al. Effects of the Insulin Sensitizer Metformin in Alzheimer Disease: Pilot Data From a Randomized Placebo-controlled Crossover Study. *Alzheimer Dis Assoc Disord* 2017;31(2):107-113. PMID 28538088.
- 53.** Luchsinger JA, Devanand D, Goldberg TE, et al. Protocol for a Randomized Phase II/III Double-Blind Placebo-Controlled Trial to Evaluate the Safety and Efficacy of Extended-Release Metformin in Amnesic Mild Cognitive Impairment: Metformin in Alzheimer Dementia Prevention (MAP). *Alzheimer Dis Assoc Disord* 2025;39(2):123-133. PMID 40434891.
- 54.** Barbera M, Lehtisalo J, Perera D, et al. A multimodal precision-prevention approach combining lifestyle intervention with metformin repurposing to prevent cognitive impairment and disability: the MET-FINGER randomised controlled trial protocol. *Alzheimers Res Ther* 2024;16(1):23. PMID 38297399.
- 55.** Luchsinger JA, Rosin SP, Kazemi EJ, et al. Glucose-Lowering Medications, Glycemia, and Cognitive Outcomes: The GRADE Randomized Clinical Trial. *JAMA Intern Med* 2025;185(7):778-787. PMID 40388190.

56. [No authors listed]. Erratum for the Report “TFEB links autophagy to lysosomal biogenesis” by C. Settembre et al. *Science* 2025;390(6779):eae5883. PMID 41411452.
57. [No authors listed]. Safety and efficacy of trehalose in amyotrophic lateral sclerosis (HEALEY ALS Platform Trial): an adaptive, phase 2/3, double-blind, randomised, placebo-controlled trial. *Lancet Neurol* 2025;24(6):500-511. PMID 40409314.
58. Yap KH, Azmin S, Abdul Manan H, et al. Randomized double-blind placebo-controlled trial of the effects of oral trehalose in spinocerebellar ataxia type 3: An interim analysis. *Parkinsonism Relat Disord* 2024;124:107013. PMID 38843619.
59. Radbakhsh S, Abrego-Guandique DM, Bacchetti T, et al. Direct hybridization and bioinformatics analysis of circulating microRNAs in patients with Alzheimer’s disease under intravenous trehalose treatment. *Brain Res* 2025;1857:149607. PMID 40187517.
60. Macfarlane S, Grimmer T, Teo K, et al. Blarcamesine for the treatment of Early Alzheimer’s Disease: Results from the ANAVEX2-73-AD-004 Phase IIB/III trial. *J Prev Alzheimers Dis* 2025;12(1):100016. PMID 39800452.
61. Brodtkin J. Concerns about Anavex’s clinical trial of Blarcamesine. *J Prev Alzheimers Dis* 2025;12(5):100137. PMID 40190003.
62. Brimson S, Thitilertdecha P, Kumaree KK, et al. Sigma-1 Receptor Ligand Blarcamesine (ANAVEX 2-73) for Alzheimer’s Disease: A Systematic Review. *CNS Drugs* 2026;40(9):1225-1241. PMID 42560467.
63. Greeley D, Nash M, Herskowitz B, et al. A phase 2 randomized, placebo-controlled study on the efficacy and safety of AR1001, a phosphodiesterase-5 inhibitor, in patients with mild-to-moderate Alzheimer’s disease. *J Prev Alzheimers Dis* 2025;12(9):100337. PMID 40912996.
64. Kang BW, Kim F, Cho JY, et al. Phosphodiesterase 5 inhibitor mirodenafil ameliorates Alzheimer-like pathology and symptoms by multimodal actions. *Alzheimers Res Ther* 2022;14(1):92. PMID 35804462.
65. Turner RS, Hebron ML, Lawler A, et al. Nilotinib Effects on Safety, Tolerability, and Biomarkers in Alzheimer’s Disease. *Ann Neurol* 2020;88(1):183-194. PMID 32468646.
66. Stevenson M, Varghese R, Hebron ML, et al. Inhibition of discoidin domain receptor (DDR)-1 with nilotinib alters CSF miRNAs and is associated with reduced inflammation and vascular fibrosis in Alzheimer’s disease. *J Neuroinflammation* 2023;20(1):116. PMID 37194065.
67. Forlenza OV, Radanovic M, Talib LL, et al. Clinical and biological effects of long-term lithium treatment in older adults with amnesic mild cognitive impairment: randomised

- clinical trial. *Br J Psychiatry* 2019;215(5):668-674. PMID 30947755.
68. Damiano RF, Loureiro JC, Pais MV, et al. Revisiting global cognitive and functional state 13 years after a clinical trial of lithium for mild cognitive impairment. *Braz J Psychiatry* 2023;45(1):46-49. PMID 36049127.
 69. Singulani MP, Ferreira AFF, Figueroa PS, et al. Lithium and disease modification: A systematic review and meta-analysis in Alzheimer's and Parkinson's disease. *Ageing Res Rev* 2024;95:102231. PMID 38364914.
 70. Prins ND, Harrison JE, Chu HM, et al. A phase 2 double-blind placebo-controlled 24-week treatment clinical study of the p38 alpha kinase inhibitor neflamapimod in mild Alzheimer's disease. *Alzheimers Res Ther* 2021;13(1):106. PMID 34044875.
 71. Jiang Y, Alam JJ, Gomperts SN, et al. Preclinical and randomized clinical evaluation of the p38 α kinase inhibitor neflamapimod for basal forebrain cholinergic degeneration. *Nat Commun* 2022;13(1):5308. PMID 36130946.
 72. Prins ND, de Haan W, Gardner A, et al. Phase 2A Learnings Incorporated into RewinD-LB, a Phase 2B Clinical Trial of Neflamapimod in Dementia with Lewy Bodies. *J Prev Alzheimers Dis* 2024;11(3):549-557. PMID 38706271.
 73. Alam JJ, Maruff P, Doctrow SR, et al. Association of Plasma Phosphorylated Tau With the Response to Neflamapimod Treatment in Patients With Dementia With Lewy Bodies. *Neurology* 2023;101(17):e1708-e1717. PMID 37657939.
 74. Iqbal S, Olivares CB, Rizzo L, et al. Repurposed drug prioritization pipeline for a multi-arm platform trial in clinical Alzheimer's disease. *Alzheimers Dement* 2026;22(5):e71495. PMID 42186115.
 75. Tonegawa-Kuji R, Karavani E, Danziger M, et al. Critical evaluation of real-world evidence of repurposable medicines in the Alzheimer's disease drug development pipeline using a target trial emulation. *Alzheimers Dement (NY)* 2026;12(1):e70193. PMID 41473419.
 76. Klionsky DJ, Abdelmohsen K, Abe A, et al. Guidelines for the use and interpretation of assays for monitoring autophagy (3rd edition). *Autophagy* 2016;12(1):1-222. PMID 26799652.
 77. Klionsky DJ, Abdel-Aziz AK, Abdelfatah S, et al. Guidelines for the use and interpretation of assays for monitoring autophagy (4th edition)1. *Autophagy* 2021;17(1):1-382. PMID 33634751.
 78. Kaizuka T, Morishita H, Hama Y, et al. An Autophagic Flux Probe that Releases an Internal Control. *Mol Cell* 2016;64(4):835-849. PMID 27818143.

79. Morishita H, Kaizuka T, Hama Y, et al. A new probe to measure autophagic flux in vitro and in vivo. *Autophagy* 2017;13(4):757-758. PMID 28121224.
80. Therriault J, Brum WS, Trudel L, et al. Blood phosphorylated tau for the diagnosis of Alzheimer's disease: a systematic review and meta-analysis. *Lancet Neurol* 2025;24(9):740-752. PMID 40818474.
81. [No authors listed]. Correction to Lancet Neurol 2025; 24: 740-52. *Lancet Neurol* 2026;25(7):e12. PMID 42309092.
82. Pilotto A, Quaresima V, Trasciatti C, et al. Plasma p-tau217 in Alzheimer's disease: Lumipulse and ALZpath SIMOA head-to-head comparison. *Brain* 2025;148(2):408-415. PMID 39679606.
83. van Dyck CH, Swanson CJ, Aisen P, et al. Lecanemab in Early Alzheimer's Disease. *N Engl J Med* 2023;388(1):9-21. PMID 36449413.
84. Sims JR, Zimmer JA, Evans CD, et al. Donanemab in Early Symptomatic Alzheimer Disease: The TRAILBLAZER-ALZ 2 Randomized Clinical Trial. *JAMA* 2023;330(6):512-527. PMID 37459141.
85. Preische O, Schultz SA, Apel A, et al. Serum neurofilament dynamics predicts neurodegeneration and clinical progression in presymptomatic Alzheimer's disease. *Nat Med* 2019;25(2):277-283. PMID 30664784.
86. Albornoz EA, Mardon K, Bhalla R, et al. PET-MRI biomarkers reveal efficacy of a novel NLRP3 inhibitor in Parkinson's disease models. *Brain* 2026;149(4):1289-1301. PMID 41099134.
87. Halle A, Hornung V, Petzold GC, et al. The NALP3 inflammasome is involved in the innate immune response to amyloid-beta. *Nat Immunol* 2008;9(8):857-65. PMID 18604209.
88. Heneka MT, Kummer MP, Stutz A, et al. NLRP3 is activated in Alzheimer's disease and contributes to pathology in APP/PS1 mice. *Nature* 2013;493(7434):674-8. PMID 23254930.
89. Ising C, Venegas C, Zhang S, et al. NLRP3 inflammasome activation drives tau pathology. *Nature* 2019;575(7784):669-673. PMID 31748742.
90. Coll RC, Hill JR, Day CJ, et al. MCC950 directly targets the NLRP3 ATP-hydrolysis motif for inflammasome inhibition. *Nat Chem Biol* 2019;15(6):556-559. PMID 31086327.
91. Harrison D, Billinton A, Bock MG, et al. Discovery of Clinical Candidate NT-0796, a Brain-Penetrant and Highly Potent NLRP3 Inflammasome Inhibitor for Neuroinflammatory Disorders. *J Med Chem* 2023;66(21):14897-14911. PMID 37874905.

92. Smolak P, Nguyen M, Diamond C, et al. Target Cell Activation of a Structurally Novel NOD-Like Receptor Pyrin Domain-Containing Protein 3 Inhibitor NT-0796 Enhances Potency. *J Pharmacol Exp Ther* 2024;388(3):798-812. PMID 38253384.
93. Doedens JR, Diamond C, Harrison D, et al. The ester-containing prodrug NT-0796 enhances delivery of the NLRP3 inflammasome inhibitor NDT-19795 to monocytic cells expressing carboxylesterase-1. *Biochem Pharmacol* 2024;227:116455. PMID 39069136.
94. Thornton P, Reader V, Digby Z, et al. Reversal of High Fat Diet-Induced Obesity, Systemic Inflammation, and Astrogliosis by the NLRP3 Inflammasome Inhibitors NT-0249 and NT-0796. *J Pharmacol Exp Ther* 2024;388(3):813-826. PMID 38336379.
95. Klughammer B, Piali L, Nica A, et al. A randomized, double-blind phase 1b study evaluating the safety, tolerability, pharmacokinetics and pharmacodynamics of the NLRP3 inhibitor selnoflast in patients with moderate to severe active ulcerative colitis. *Clin Transl Med* 2023;13(11):e1471. PMID 37962000.
96. Fawzy MN. The neuroimmune network in Alzheimer's and Parkinson's diseases: from mechanistic insights to biomarker-guided immunotherapies and clinical translation. *Inflammopharmacology* 2026;34(9):5975-6001. PMID 42627617.
97. Suárez-Calvet M, Morenas-Rodríguez E, Kleinberger G, et al. Early increase of CSF sTREM2 in Alzheimer's disease is associated with tau related-neurodegeneration but not with amyloid- β pathology. *Mol Neurodegener* 2019;14(1):1. PMID 30630532.
98. Ewers M, Franzmeier N, Suárez-Calvet M, et al. Increased soluble TREM2 in cerebrospinal fluid is associated with reduced cognitive and clinical decline in Alzheimer's disease. *Sci Transl Med* 2019;11(507):eaaw6221. PMID 31462511.
99. Galpern WR, Triana-Baltzer G, Li L, et al. Phase 1 Studies of the Anti-Tau Monoclonal Antibody JNJ-63733657 in Healthy Participants and Participants with Alzheimer's Disease. *J Prev Alzheimers Dis* 2024;11(6):1592-1603. PMID 39559872.
100. Wang F, Liu J, Bai D, et al. A Phase 1 Study of the Pharmacokinetics, Safety, and Tolerability of Posdinemab (JNJ-63733657) in Healthy Chinese Adults. *J Clin Pharmacol* 2026;66(2):e70116. PMID 41078145.
101. Pérez-Ruixo C, Li L, Galpern WR, et al. Mechanistic Population Pharmacokinetic-Pharmacodynamic Model of the Tau-Targeted Antibody Posdinemab in Healthy Participants and Participants with Alzheimer's Disease. *Clin Pharmacol Ther* 2026;119(4):979-989. PMID 41431125.
102. Sato C, Barthélemy NR, Mawuenyega KG, et al. Tau Kinetics in Neurons and the Human Central Nervous System. *Neuron* 2018;97(6):1284-1298.e7. PMID 29566794.

- 103.** Sato C, Barthélemy NR, Mawuenyega KG, et al. Tau Kinetics in Neurons and the Human Central Nervous System. *Neuron* 2018;98(4):861-864. PMID 29772204.
- 104.** Dobrowolska Zakaria JA, Bateman RJ, Lysakowska M, et al. The metabolism of human soluble amyloid precursor protein isoforms is quantifiable by a stable isotope labeling-tandem mass spectrometry method. *Sci Rep* 2022;12(1):14985. PMID 36056033.
- 105.** Zhang Y, Hu M, Meng Y, et al. DNA nanodevices detect an acidic nanolayer on the lysosomal surface. *Nat Cell Biol* 2026;28(2):285-295. PMID 41565777.