LOW SIGNAL HIGH NOISE

The Excitatory–Inhibitory Balance in Alzheimer's Disease — An Adjudication of Chronic Excitatory Insufficiency Against Terminal Disinhibition, Two Opposite-Signed Mechanisms That Claim the Same Epileptiform Datum

The Contested Datum · Two Models at Their Strongest · The Precedent That Is Real
The Sign of the Drug · Two Axes, Not One · Resolution by Epoch

A Doctoral Synthesis Prepared in Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy in Neuroscience

Prepared under the Organic Network Synthesis methodology

AdultCognitiveDisease.com

August 2026

Two models of the Alzheimer network claim the same epileptiform datum through opposite mechanisms, and this corpus held both without noticing. This volume adjudicates them, records the fidelity error that concealed the contradiction, and argues that the disagreement was an artefact of collapsing two variables — the signal a circuit can carry, and the noise it makes when its restraints fail — into one.

"Excitatory insufficiency (meaning too much inhibition)."

— Moosmann & Sohre, defining their own terms

"…increased intrinsic excitability, reduced GABAergic inhibition…"

— this corpus, purporting to describe them

The two sentences are about the same phenomenon and they carry opposite signs. The second was printed as an exposition of the first, and stood for a year. This volume is what was owed once the discrepancy was noticed — not a patch to the sentence, but the argument the corpus had been assuming it already had.

THE ADJUDICATION VOLUMES

One Sensor, Two Collapses (the two NLRP3 arms, graded on one scale)

The Three Rivals (the three-phase architecture, stress-tested adversarially)

The Denominator (what fraction of dementia the architecture actually describes)

The Case for the Cascade (amyloid primacy, steelmanned without rebuttal)

Low Signal, High Noise ← this volume

The adjudication volumes exist because a synthesis that only accumulates is a synthesis that cannot be wrong. Each takes a place where the corpus holds two positions at once, or holds one it has not earned, and forces the question. This one was prompted by a correction: the discovery that the corpus had rendered a rival hypothesis with its sign reversed, and that the reversal had concealed a contradiction inside the corpus's own account of excitation and inhibition. It supersedes the treatment of chronic excitatory insufficiency in *Convergent Synaptic Collapse* §4 as the corpus's account of that question.

Contents

Abstract

Alzheimer's disease carries a large excess of epileptiform activity. Subclinical discharges are recorded in a substantial fraction of patients, clinical seizures in perhaps one in ten, and the excess rises steeply in early-onset autosomal-dominant and Down-syndrome forms. That datum is not in dispute. What is in dispute is its sign, and the corpus has until now held both answers at once without noticing.

Two models claim the same observation through opposite mechanisms. The first is **chronic excitatory insufficiency** (CEI), argued by Moosmann and Sohre in their 2020 Oskar Fischer Prize submission, which holds that Alzheimer's disease is at bottom a failure of glutamatergic, NMDA-type excitatory neurotransmission, and that amyloid-β and phospho-tau are not toxins but compensatory sensitisers — the brain's long, initially successful, finally lost attempt to restore excitatory tone. On this reading the network is tipped *toward inhibition*, and the seizures are what an over-inhibited network does when its sparse surviving excitatory signals are pathologically over-amplified — the mechanism established for absence and for autosomal dominant nocturnal frontal lobe epilepsy. The second is **terminal disinhibition**, which this corpus has argued across at least four volumes: perineuronal-net digestion and the progressive failure of parvalbumin-positive fast-spiking interneurons remove the network's brake, pyramidal cells fire without restraint, and the seizures are the audible edge of an excitotoxic terminus. One datum, two mechanisms, opposite signs, and no document in the corpus has previously set them against each other.

This volume is that adjudication, and it begins by recording the corpus's own error. Correction C-016 found that the *Convergent Synaptic Collapse* thesis rendered CEI with its sign inverted: it explained the AD seizure excess by "reduced GABAergic inhibition" and the "loss of inhibitory GABAergic interneurons," and attributed that explanation to Moosmann — who derives the excess the opposite way, and derives it that way precisely to forestall this misreading. The error is worth stating carefully, because it is not the error it appears to be. The corpus did not adopt a wrong mechanism; it adopted the *better-supported* mechanism and misattributed it to a hypothesis arguing against it. This is a **fidelity** failure, not a mechanistic one, and it is the more insidious of the two: a corpus that quietly rewrites a rival into agreement with itself has stopped being able to learn from it.

The steelman is taken seriously, and one leg of it survives. Moosmann's precedent is real. Seizures genuinely can arise from too much inhibition: tonic GABA-A inhibition is *enhanced* in thalamocortical neurons across genetic and pharmacological models of typical absence epilepsy, through compromised GABA uptake by GAT-1 (Cope and colleagues, 2009), and two mouse models of human autosomal dominant nocturnal frontal lobe epilepsy show frequent spontaneous seizures alongside *enhanced* cortical GABAergic inhibition (Klaassen and colleagues, 2006). The physiological core of CEI's amyloid claim is real too: picomolar Aβ42 markedly

increases hippocampal long-term potentiation and enhances reference and contextual fear memory (Puzzo and colleagues, 2008). Neither of these is a fringe result, and the corpus has been too quick to treat CEI as merely wrong.

The adjudication turns on a test the sign war supplies for free: the direction of the effective drug tells you the direction of the lesion. In absence epilepsy, where inhibition is genuinely excessive, the rational and demonstrated therapeutic direction is to *reduce* tonic inhibition. In Alzheimer's disease, every intervention that works runs the other way. Restoring interneuron function rescues the phenotype: Nav1.1 restoration in parvalbumin cells corrects network dysfunction and cognition in an amyloid model (Verret and colleagues, 2012). Lowering excitability rescues it: levetiracetam alone among a panel of antiepileptics suppressed spike activity and reversed hippocampal remodelling, synaptic dysfunction and learning deficits in hAPP mice (Sanchez and colleagues, 2012), and a low dose reduced dentate/CA3 hyperactivation and improved memory in amnestic mild cognitive impairment (Bakker and colleagues, 2012). Hyperactivity is observed directly, and its cause is measured as a relative *decrease* in synaptic inhibition (Busche and colleagues, 2008). If Alzheimer's seizures were absence-like, these results should have the opposite sign. They do not.

But the ladder of levetiracetam trials also refuses the simple answer, and this volume follows it to the end. The LEV-AD crossover trial found improvement in the subgroup with epileptiform activity and not in the group as a whole (Vossel and colleagues, 2021); the ILiAD pilot missed its primary endpoint entirely, with a small executive-domain signal confined to the nine participants with baseline epileptiform EEG (Sen and colleagues, 2024). Anti-excitability treatment works where hyperexcitability is documented and not otherwise. That is not a failure of the disinhibition model — it is its correct dosing, and it is the first clinically actionable statement this adjudication produces.

The resolution proposed is that the sign war is an artefact of collapsing two variables into one. "Insufficient meaningful excitation" — throughput, the signal a circuit can actually carry — is not the same quantity as "hyperexcitability" — the noise a circuit generates when its restraints fail. A brain can be low on the first and high on the second simultaneously, and the Alzheimer brain demonstrably is: globally hypometabolic and locally seizing, in the same tissue, in the same year. **Low signal, high noise.** Read on two axes, the models stop contradicting each other on the electrophysiology and remain genuine rivals only where they always were — on the *valence of amyloid*, which is a question about causation and not about excitability at all. The volume closes there, grading CEI's inversion as a defensible provocation with an indefensible mechanism, and stating exactly which of its claims the corpus should adopt.

I. The Contested Datum

Begin with the observation both models are trying to explain, stated without either's vocabulary.

The excess is large and it is real. People with Alzheimer's disease have seizures far more often than their age-matched peers. Clinical seizures occur in something like one patient in ten, against a population baseline well under one per cent, and subclinical epileptiform discharges — the kind that never produce a visible event and are found only because somebody recorded — occur in a substantially larger fraction still. The excess is steeper in early-onset autosomal-dominant disease and steeper again in Down syndrome, which is to say it tracks amyloid dose. It is not an artefact of end-stage brains: it appears early, and in some series it precedes diagnosis.

The most important version of the datum is the one you cannot see from the scalp. Lam and colleagues (2017) placed foramen ovale electrodes in two patients with Alzheimer's disease and no clinical seizure history and recorded clinically silent hippocampal seizures and epileptiform spikes arising from mesial temporal structures — activity that was largely invisible to simultaneous scalp electroencephalography, and that occurred preferentially during sleep, the period in which it is most likely to interfere with memory consolidation. This matters for the adjudication in a way that is easy to miss. It means the true rate of the phenomenon is unknown and certainly higher than the recorded rate; it means the phenomenon is anatomically concentrated exactly where the disease begins; and it means the seizures are not a late complication of a ruined cortex but an early event in a structure that still works.

Why this datum is the right place to arbitrate. Most disagreements in this field are about interpretation of a picture — which cell died first, what a signal on a scan means, whether a mouse resembles a person. The seizure datum is different. It is a direct electrophysiological measurement of the variable both models are about. If two models of excitatory–inhibitory balance cannot be distinguished by an electrophysiological recording, they are not really different models. And if they can, the recording settles it.

What each model must therefore explain. Any account of Alzheimer's disease that speaks about excitation and inhibition must produce: the excess itself; its early appearance; its anatomical concentration in the mesial temporal lobe; its scaling with amyloid dose; its nocturnal preference; and — the constraint that turns out to be decisive — the *direction* of the interventions that modify it. The remainder of this volume works through those six requirements twice, once for each model, and then asks what is left standing.

A note on what is not in dispute. Neither model denies the datum, and neither model denies that the late Alzheimer brain is globally hypometabolic — that fluorodeoxyglucose uptake falls in temporoparietal cortex years before diagnosis is one of the most reproducible findings in the imaging literature. Both models accept both facts. They differ on which fact is the disease and which is the consequence.

II. The Corpus's Own Error, and Why Its Type Matters

Before adjudicating between two positions it is necessary to state that the corpus previously misreported one of them, and to be precise about the kind of misreport it was.

What was written. The *Convergent Synaptic Collapse* thesis, §4, explained the Alzheimer seizure excess by writing that neurons "undergo compensatory upregulation of excitability — increased intrinsic excitability, reduced GABAergic inhibition," and that "the loss of inhibitory GABAergic interneurons in AD … removes the normal inhibitory constraint, exacerbating seizure susceptibility despite the underlying excitatory insufficiency." That passage sat inside a chapter titled *The Chronic Excitatory Insufficiency Hypothesis* and presented as an exposition of Moosmann and Sohre.

Why it is wrong as attribution. CEI is a hypothesis about a network tipped *toward* inhibition — in Moosmann's own formulation, "excitatory insufficiency (meaning too much inhibition)." His account of the seizure excess is explicitly *not* disinhibition. It is compensatory over-amplification of small excitatory signals against an over-inhibited background, the silent- and absence-seizure model, and he derives it that way precisely to pre-empt the disinhibition misreading. The corpus took a hypothesis whose author had guarded against exactly one error and committed exactly that error on the single phenomenon he used as the guard.

Why the error's type is the interesting part. The obvious reading is that the corpus held a wrong mechanism. It did not. A five-cluster primary-literature fact-check found that the disinhibition account the corpus was smuggling in is the **consensus-aligned and better-supported** one, and that Moosmann's over-amplification account is the **minority** reading of Alzheimer seizures. So the corpus flipped a rival's sign *toward the better answer*. This makes the fault a **fidelity** failure rather than a mechanistic one — and fidelity failures are the more dangerous species, for three reasons.

First, they are invisible to the ordinary check. A reader who knows the literature reads that passage, finds it correct, and moves on; the error is only visible to a reader who has read *Moosmann*, and nobody had. Second, they destroy the thing a rival is for. A corpus keeps a rival framework in order to be argued with. A rival silently rewritten into agreement cannot argue, and its retention becomes decorative — the appearance of pluralism with none of the cost. Third, and most consequentially here, the rewrite concealed a genuine open problem. Because CEI had been rendered as disinhibition, and the corpus's own terminal model is disinhibition, the two appeared to agree, and nobody noticed that **two opposite-signed mechanisms in the same corpus**

were claiming the same datum. The error did not merely misstate a hypothesis; it hid an unreconciled contradiction inside the corpus's own account.

What has been done, and what this volume adds. The §4 subsection has been rewritten: it now states Moosmann's actual mechanism, flags it as the minority reading, relocates the interneuron-loss clause to the corpus's own terminal account, and adds a line noting that the two accounts claim the same datum and are resolved by epoch. That is the correction. But a one-sentence assertion that two models "are resolved by epoch" is not a resolution; it is a promissory note. **This volume is the argument the note promised** — and it is written as a standalone adjudication rather than a longer footnote because a claim asserted in five places and argued in none is precisely the structural failure the corpus has diagnosed elsewhere in itself.

III. Chronic Excitatory Insufficiency, Stated at Its Strongest

An adjudication is worthless if the losing model is presented in a form its author would not recognise. What follows is CEI as Moosmann and Sohre argue it, with the corpus's objections withheld until Section XII.

The core claim. The established risk factors for Alzheimer's disease look heterogeneous — a lipoprotein allele, a γ-secretase mutation, a supernumerary chromosome, a blow to the head, the loss of an ovarian hormone — and CEI proposes that they are heterogeneous only at the surface. Each of them, by a different route, reduces glutamatergic, NMDA-type excitatory neurotransmission. The disease is what a cortex does when its excitatory throughput has been chronically insufficient for decades. On this reading the primary lesion is a *deficit of signal*, and everything the pathologist stains for is downstream of it.

The genetic arm, which is the strongest leg. ApoE4's mechanism, in CEI's account, is not lipid handling in general but a specific receptor lesion: ApoE4 impairs the recycling of ApoE receptors, reducing neuronal surface expression of ApoER2 together with NMDA and AMPA receptors by sequestering them in intracellular compartments, and thereby reducing reelin's ability to enhance synaptic glutamate receptor activity (Chen and colleagues, 2010). This is a real, reproducible, Herz-laboratory mechanism, and it is worth the corpus noticing that it is the *same receptor* on which four volumes of the reelin series are built. Where CEI is strongest it is not a rival to this corpus at all; it is a differently-motivated description of the corpus's own axis.

The developmental arm. For the amyloid-dosage factors — *APP* duplication, trisomy 21 — CEI proposes a prenatal origin, a "glutamatergic hypogenesis," in which the excitatory apparatus is under-built before birth and the lifetime of amyloid overproduction is compensation for the shortfall. The logic is the same whole-chromosome, time-compressed logic the corpus's own Down-syndrome volume uses, arrived at independently.

The valence claim, which is the load-bearing one. Aβ and phospho-tau are, in this account, not toxins but instruments. Aβ at physiological concentrations is a glutamatergic sensitiser; tau phosphorylation is an adaptive modification of the cytoskeleton under energetic stress. The disease is the *allostatic cost* of running these compensations for forty years — a bill that comes due, rather than a poison that accumulates.

The phenomenological reach, which is genuinely impressive. CEI predicts, and the clinic supplies, a set of non-cognitive features that the amyloid cascade explains poorly: sleep fragmentation and shortened slow-wave sleep; early high-frequency sensorineural hearing loss; apathy as a motivational rather than a mnemonic deficit. Whether or not the mechanism is right, a framework that generates the non-memory phenotype from one premise has done something the field's dominant framework has not, and the corpus should say so.

And the seizure claim, in its own terms. Given an inhibition-dominant network, why should it seize at all? Moosmann's answer is that a network in which excitation is sparse and inhibition is high does not fall silent — it becomes unstable in a particular way, over-amplifying the few excitatory events that get through. This is not an invention for the occasion. It is the established physiology of a specific class of epilepsy, and Section V takes it seriously enough to look the precedent up.

IV. Terminal Disinhibition, Stated at Its Strongest

The corpus's own model deserves the same treatment, including the parts of it that are weaker than the corpus has admitted.

The core claim. The fast-spiking, parvalbumin-expressing basket cell is the cortex's brake. It fires above two hundred hertz, carries the highest oxidative load of any cortical neuron, and is wrapped in a perineuronal net that serves as structural scaffold, ion-exchange buffer and antioxidant shield. In Alzheimer's disease the net is digested by the proteases of an ungoverned microglion; the interneuron, stripped, is poisoned by its own activity; inhibition fails; pyramidal cells fire without restraint; and the resulting excitotoxicity is the terminus. The seizures are the audible edge of that process.

The direct evidence for network disinhibition. Aberrant excitatory neuronal activity, non-convulsive seizure activity and compensatory inhibitory remodelling are documented in amyloid models (Palop and colleagues, 2007). Interneuron dysfunction is the organising theme of the field's most-cited review of the topic (Palop and Mucke, 2016). And the causal experiment has been done: restoring Nav1.1 levels in parvalbumin interneurons — that is, restoring the interneuron's *ability to fire* — corrects network dysfunction and cognitive deficits in an amyloid model (Verret and colleagues, 2012). Restoring inhibition rescues. That single result is the model's spine.

The imaging-level evidence that the lesion is local, not global. In amyloid-depositing mice, cortical layer 2/3 contains both silenced and hyperactive neurons in the same field; the hyperactive ones are found exclusively near plaques; and the hyperactivity is attributable to a relative *decrease* in synaptic inhibition (Busche and colleagues, 2008). This is the observation on which the two-axis reading of Section IX ultimately rests, and it deserves emphasis: **the same tissue contains cells that are firing too little and cells that are firing too much, metres apart in cortical terms, at the same moment.** Any model that assigns the brain one number for "excitability" has already lost information the microscope supplies.

The mechanism that makes it self-sustaining. Zott and colleagues (2019) closed the loop: hyperactivation is initiated by suppression of glutamate reuptake, occurs preferentially in neurons that already had high baseline activity while inactive neurons resist, and is sustained by Aβ dimers purified from Alzheimer brain extracts. Hyperactivity begets amyloid — synaptic activity regulates interstitial-fluid Aβ directly (Cirrito and colleagues, 2005) — and amyloid begets hyperactivity. It is a vicious cycle in the strict sense, and it is the reason the disinhibition model does

not require a continuing external driver once it has begun.

Where the corpus's model is weaker than it has admitted. Three places. First, the perineuronal-net arm is largely animal and post-mortem, and the corpus's own errata found that the key human resilience study reports net *remodelling*, not preservation — the resilient brain has *fewer* nets, not more. Second, "terminal" is doing quiet work: the model is a model of the endgame, and its authors have sometimes written as though it described the whole disease. Third, and most awkwardly, the model does not on its own explain the global hypometabolism, which is early, large, and in the wrong direction for a brain whose problem is too much firing. Section IX is where that awkwardness is paid off rather than deferred.

V. The Precedent Moosmann Invokes Is Real

It would be convenient for the corpus if the seizure claim were an invention. It is not, and this section exists to say so before the adjudication turns against it.

Absence epilepsy runs on excess inhibition. Cope and colleagues (2009) found that extrasynaptic GABA-A-receptor-dependent *tonic* inhibition is increased in thalamocortical neurons across diverse genetic and pharmacological models of typical absence seizures, and that in the genetic models the increase arises from compromised GABA uptake by the transporter GAT-1. The paper's title states the finding without hedging: enhanced tonic GABA-A inhibition in typical absence epilepsy. This is a seizure disorder in which the inhibitory tone is too high, and the therapeutic logic that follows — that inverse agonists at extrasynaptic GABA-A receptors should help — has been pursued on exactly that reasoning.

And a second, independent precedent. Klaassen and colleagues (2006) engineered two mouse models of human autosomal dominant nocturnal frontal lobe epilepsy, carrying mutations in the α4 or β2 neuronal nicotinic acetylcholine receptor subunits. The animals showed persistent abnormal cortical electroencephalograms and frequent spontaneous seizures — alongside *enhanced* cortical GABAergic inhibition. Two different genes, two different receptor families, the same counter-intuitive result: a cortex can seize because it is too inhibited.

What this establishes, precisely. It establishes that "the network is over-inhibited" and "the network seizes" are compatible propositions, and that anyone who dismisses CEI's seizure account as physiologically incoherent is simply wrong. The mechanism exists. It has been measured. It has a therapeutic direction of its own. Moosmann did not reach for a convenient hand-wave; he reached for the correct precedent for the topology he was proposing.

What it does not establish. That Alzheimer's disease is such a case. The existence of a mechanism is not evidence that a given disease uses it, and the two precedent disorders are of a very particular kind — thalamocortical oscillatory disorders with paroxysmal, generalised, brief events, and a frontal-lobe nocturnal epilepsy with a channel lesion. Alzheimer's disease is a progressive neurodegeneration whose epileptiform activity is focal, mesial-temporal, and superimposed on cell loss. Nothing about the anatomy invites the comparison, and the burden is on CEI to show that the topology transfers.

Which brings the adjudication to a testable form. Two mechanisms, both real, both capable of producing seizures, with opposite signs. There is a way to tell them apart that requires no new theory and no new anatomy, and the literature has already run it many times without noticing what it was running. It is the subject of the next section.

VI. The Discriminating Test — The Sign of the Drug Tells You the Sign of the Lesion

This is the analytical core of the volume and it can be stated in one sentence: **if a network seizes because it is over-inhibited, the treatment is to reduce inhibition; if it seizes because it is under-inhibited, the treatment is to reduce excitation — and the two are not confusable, because a drug that helps in one case must worsen the other.**

The test in the control case. In absence epilepsy, where excess tonic inhibition is measured rather than inferred (Cope and colleagues, 2009), the therapeutic direction follows the physiology: reduce the tonic inhibitory current. The proposal to use inverse agonists at extrasynaptic GABA-A receptors is derived directly from the mechanism, and it runs in the opposite direction from ordinary antiepileptic practice, which is to *increase* inhibition. Absence epilepsy is thus a worked example of the test's discriminating power: the mechanism predicted an unusual drug direction, and the unusual direction is the one the field pursues.

The test in Alzheimer's disease, run four times. Every intervention that has modified the Alzheimer network phenotype runs in the disinhibition direction.

- **Restore inhibition, and it helps.** Nav1.1 restoration in parvalbumin interneurons — increasing the brake's capacity to act — corrects network dysfunction and cognitive deficits (Verret and colleagues, 2012).
- **Reduce excitability, and it helps.** Levetiracetam, uniquely among a panel of antiepileptic drugs tested, suppressed abnormal spike activity and reversed hippocampal remodelling, synaptic dysfunction and learning and memory deficits in hAPP mice (Sanchez and colleagues, 2012).
- **Reduce excitability in humans, and it helps.** A low dose of levetiracetam reduced elevated dentate-gyrus/CA3 activation in amnestic mild cognitive impairment to a level indistinguishable from healthy controls, and improved memory performance relative to placebo (Bakker and colleagues, 2012).
- **Measure the cause of the hyperactivity directly, and it is loss of inhibition.** The hyperactive neurons near plaques are hyperactive because synaptic inhibition onto them is relatively reduced (Busche and colleagues, 2008).

What CEI predicts instead, and why the prediction is not idle. If the Alzheimer network were over-inhibited in the absence-epilepsy sense, then reducing inhibition should improve it and reducing excitation should worsen it. CEI is admirably explicit about this and does not flinch: its therapeutic programme calls for NMDA enhancement, glutamate-receptor trafficking enhancement, and *selective GABAergic antagonism* — deliberately lifting the brake. That is a real, falsifiable, opposite-signed prescription. It has not been tested in Alzheimer's disease, and the results above make it a prediction the corpus expects to fail.

The strength and the limit of the test. Its strength is that it is not an argument from authority or from the number of papers on each side; it is a structural constraint. A drug cannot help by lowering excitability in a network whose problem is insufficient excitation. Its limit is that all four Alzheimer results above are model systems or small human studies, and that the human levetiracetam story does not stop at Bakker. Section VII follows it to the end, because the end is where the honest complication lives.

VII. The Levetiracetam Ladder — Four Trials, Read in Order

The temptation, having established that anti-excitability treatment helps, is to stop. The clinical literature does not permit it, and reading the trials in sequence produces a more useful conclusion than reading any one of them.

Rung one — the model system, and the strongest result. Sanchez and colleagues (2012) tested a panel of antiepileptic drugs in hAPP mice. Only levetiracetam reduced the abnormal electroencephalographic spike activity; chronic treatment then reversed hippocampal remodelling, behavioural abnormalities, synaptic dysfunction, and deficits in learning and memory. This is close to an ideal preclinical result: a mechanistic readout and a functional readout, moving together, with a within-experiment negative control in the other drugs.

Rung two — the small human study, and the cleanest human evidence. Bakker and colleagues (2012) gave a low dose of levetiracetam to patients with amnestic mild cognitive impairment. Hippocampal activation in the dentate gyrus and CA3, elevated under placebo relative to healthy controls, fell under drug to a level that no longer differed from controls; and memory performance in the scanning task improved. The design's power lies in its coupling: it did not merely show a cognitive change, it showed the *targeted physiological abnormality* normalising alongside it.

Rung three — the randomised trial, and the first restriction. The LEV-AD study, a phase 2a randomised double-blind placebo-controlled crossover trial in thirty-four adults with Alzheimer's disease, found that levetiracetam improved cognition in patients **with** epileptiform activity, and did not do so in the group as a whole (Vossel and colleagues, 2021). The signal survived, but it had acquired a boundary: the drug worked where the abnormality it targets was present.

Rung four — the pilot that missed, and the second restriction. The ILiAD study, a double-blind placebo-controlled crossover pilot in mild-to-moderate Alzheimer's disease, **did not meet its primary endpoint**: no overall difference in cognitive outcomes against placebo. A small improvement on a Stroop subscale and a virtual route-learning test appeared at group level among the nine participants who had detectable epileptiform activity on baseline electroencephalography, suggesting a domain-specific effect on executive skills (Sen and colleagues, 2024).

What the ladder actually shows. Not that the disinhibition model is wrong — the physiological target normalises whenever it is measured, at every rung. What it shows is that

hyperexcitability is a feature of a subset of patients, not of the disease, and that treating the disease as though every patient had it produces a null. Two rungs with unselected populations produced restricted or absent effects; every rung that selected on the abnormality produced an effect. This is what a real but non-universal mechanism looks like when it is tested without a biomarker.

Why this is the volume's most useful practical result. It converts a sign war into an inclusion criterion. The question "is the Alzheimer brain over- or under-inhibited?" has no single answer because it is not a single population, and the appropriate clinical question is not which model is right but *which patients have the phenotype* — which, given Lam's foramen ovale recordings, is a larger and less identifiable group than scalp electroencephalography suggests. The corpus should stop asking the first question and start asking the second, and Section XVI states what that requires.

VIII. The Memantine Anomaly, Re-read

CEI's most rhetorically effective move is its account of memantine, and it deserves a direct answer rather than the silence the corpus has previously given it.

The anomaly. Memantine is an uncompetitive NMDA-receptor antagonist with a modest but reproducible benefit in moderate-to-severe Alzheimer's disease (Reisberg and colleagues, 2003). On a naive reading of any model in which the disease involves too little excitation, blocking the principal excitatory receptor should make things worse. On a naive reading of the amyloid cascade it should do nothing at all. That a glutamate antagonist is one of only two approved symptomatic drug classes is a fact both dominant frameworks have tended to mention and move past.

CEI's explanation. In an excitation-insufficient network, parvalbumin interneurons are said to become hyperexcitable by compensation while depending on tonic, extrasynaptic NMDA drive. Memantine, blocking tonic current while sparing phasic transmission, preferentially removes the interneurons' tonic excitation, lowers inhibitory tone, and releases the pyramidal cell. The drug works, on this account, by *disinhibiting* — which is coherent, ingenious, and exactly the shape one wants from a framework that claims the network is over-inhibited.

The straightforward alternative, which requires no inversion. Memantine's pharmacology is voltage-dependent, use-dependent, low-affinity channel block. It preferentially occludes tonically activated, extrasynaptic NMDA receptors — the ones carrying pathological standing current — while relinquishing the channel during high-concentration synaptic glutamate transients. In a network with excess extrasynaptic glutamate and impaired reuptake, that is a drug that removes noise and spares signal. This is precisely the profile Zott and colleagues' (2019) mechanism calls for, in which hyperactivation is initiated by suppression of glutamate reuptake. The drug does not need to disinhibit anything; it needs only to clip a pathological tonic current.

How to tell the two readings apart. They make different predictions about interneurons. CEI's account requires memantine's benefit to depend on intact parvalbumin populations, since the drug is said to act by removing tonic drive *from those cells*; benefit should therefore fall as interneurons are lost. The alternative requires benefit to depend on the presence of pathological extrasynaptic glutamate, and predicts that benefit should track markers of excitotoxic tone and impaired reuptake, not interneuron count. The experiment is straightforward in a model system and has not been done.

The adjudication's verdict on this point, stated with its uncertainty. The alternative reading is more parsimonious, is directly supported by the reuptake mechanism, and does not require the network to be over-inhibited. But CEI's version is not refuted, and the corpus should

stop treating the memantine anomaly as though it were already answered. **This is the one place in the adjudication where CEI's account remains live**, and it is live precisely because nobody has run the interneuron-dependence experiment. That is recorded in Section XV as a prediction rather than papered over.

IX. Two Axes, Not One

Everything so far has treated the question as though the brain had a single dial marked *excitation* with a single correct setting. The remainder of the volume argues that this framing is the actual source of the disagreement, and that it is false.

The two variables. *Throughput* is the amount of meaningful excitatory signal a circuit can carry — the information-bearing traffic between neurons that are supposed to be talking. *Excitability* is the propensity of neurons to fire in the absence of, or out of proportion to, that traffic. These are different quantities. They are usually correlated in healthy tissue, which is why one dial seems sufficient. In damaged tissue they decouple, and they decouple in a characteristic direction: **as the meaningful traffic falls, the noise rises**, because the restraints that shaped the traffic fail at the same time as the traffic does.

The Alzheimer brain demonstrably has both. Global temporoparietal hypometabolism, the most reproducible functional imaging finding in the disease, is a throughput measurement: less glucose consumed means less synaptic work done. Silent hippocampal seizures (Lam and colleagues, 2017) are an excitability measurement. They are found in the same disease, in overlapping stages, and there is no contradiction in that, any more than there is a contradiction between a failing radio that receives less programme and more static. The single-dial framing forces a choice between two facts that do not compete.

The decisive observation is at cellular scale. Busche and colleagues (2008) found, in one cortical field, that 29% of layer 2/3 neurons showed *decreased* activity and 21% showed *increased* activity — the hyperactive ones exclusively near plaques. That is not a network with a setting. It is a network with a *distribution*, whose mean is uninformative and whose variance is the pathology. Averaging it produces a number that describes no cell in the tissue, and the sign war between CEI and terminal disinhibition is, in large part, a disagreement about which tail of that distribution to name the disease.

What this does to each model. It does not vindicate them equally. CEI's throughput claim survives — a cortex with declining glutamatergic capacity is a real description of the Alzheimer brain, and the risk-factor convergence on NMDA function is not nothing. What does *not* survive is CEI's derivation of the seizures from that claim, because the seizures are demonstrably a decoupling phenomenon with the opposite local sign, and their cause has been measured as reduced synaptic inhibition rather than over-amplification. Terminal disinhibition, meanwhile, is vindicated on excitability and remains silent on throughput, which the corpus should now say plainly rather than allowing the model to be read as an account of the whole disease.

The phrase, and what it commits the corpus to. Low signal, high noise. The formulation is not a compromise between two models; it is a claim that they were measuring different things and reporting the results as though they were the same measurement. It commits the corpus to reporting both variables whenever it speaks about excitatory–inhibitory balance, to never inferring one from the other, and to treating any statement of the form "the Alzheimer network is hyperexcitable" or "the Alzheimer network is hypoactive" as incomplete without a specified variable, scale and epoch.

X. Resolution by Epoch

The two-axis reading dissolves most of the contradiction. Epoch dissolves the rest, and this section states what C-016 asserted in one sentence and left unargued.

The proposal. The two models describe different halves of the natural history, and their electrophysiological claims are not rivals but sequential.

Epoch	Throughput	Excitability	What is measurable	Whose account fits
Preclinical, decades	Falling slowly; risk factors converge on glutamatergic capacity	Locally rising near plaques; silent mesial-temporal spikes	FDG hypometabolism; foramen ovale spikes; DG/CA3 over-activation	CEI on throughput; disinhibition on the local excess
Prodromal / MCI	Falling; compensation visible as over-activation	Elevated and now measurable at scalp in a subset	Hippocampal hyperactivation normalised by low-dose levetiracetam	Disinhibition, with a treatable window
Mild–moderate	Substantially reduced	Heterogeneous; epileptiform in a subset only	LEV benefit confined to the epileptiform subgroup	Disinhibition, restricted by phenotype
Terminal	Collapsed	Frankly disinhibited where interneurons are lost	Interneuron loss; net digestion; clinical seizures	Terminal disinhibition

The two accounts, placed. Note that no row is occupied by a single model, which is the point.

Why the compensated prodrome belongs to CEI's variable and not its mechanism. In the long silent phase, throughput is falling and the brain is compensating for it — and Moosmann is right that compensation is what is happening, right that it is largely successful for decades, and right that its cumulative cost is part of the disease. What he is not right about is the *identity of the compensator*. The corpus's evidence is that amyloid is not the compensation but the injury the compensation is failing to contain, and Section XIII takes that up as the one genuinely unresolved fork.

Why the terminus belongs to disinhibition without qualification. By the time interneurons are lost and the perineuronal nets digested, there is no serious dispute: the network is under-inhibited where the brake has gone. Nothing in CEI addresses the terminus, and it should not be asked to.

The honest weak point of this resolution. It is a schema, not a measurement. No study has measured throughput and excitability in the same human brains across the staging series, and the epochs above are assembled from populations of different mean age studied with different instruments. The resolution is therefore graded in Section XIV as synthesis, and Section XV names the co-registration that would test it. It should not be quoted as though it had been observed.

What it nonetheless buys immediately. It explains why the field's literature reads as contradictory: papers reporting hypoactivity and papers reporting hyperactivity are sampling different epochs, different variables and different spatial scales, and are all correct. That is a more useful thing to be able to say than a verdict for one side.

XI. What CEI Gets Right and the Corpus Should Adopt

A rival is only worth keeping if it can change the corpus's mind about something. Four things.

One — the physiological role of amyloid at physiological concentrations. Picomolar Aβ42 markedly increases hippocampal long-term potentiation and enhances both reference and contextual fear memory (Puzzo and colleagues, 2008). This is a real result from a mainstream laboratory, it has been replicated, and it is conceded even by the cascade's principal defenders. The corpus has generally written about amyloid as though its only concentration were the pathological one, and that is a misdescription of the molecule. Aβ has a physiological function in synaptic plasticity, and a framework that cannot say so is not describing the biology.

Two — the risk-factor convergence on glutamatergic signalling is a real node, though not the only one. ApoE4 does impair ApoE-receptor recycling and thereby reduce surface glutamate receptors and reelin-evoked signalling (Chen and colleagues, 2010). That is one of the better-characterised mechanisms attaching the field's largest genetic risk factor to a synaptic consequence, and it is the *same receptor axis* the corpus's reelin series is built on. CEI arrived at ApoER2 from excitability; the corpus arrived at it from resilience and matrix staging. Two independent routes to one receptor is corroboration, and the corpus should record it as such rather than treating CEI's genetic arm as an opponent's claim.

Three — the non-cognitive phenotype demands an account, and the corpus has not given one. Sleep fragmentation, early high-frequency hearing loss and apathy are early, common and poorly explained by protein deposition. The corpus has volumes on sleep and on the noradrenergic system that bear on the first, and effectively nothing on the second and third. CEI generates all three from one premise. Whether or not the premise is right, the *demand* is legitimate, and the corpus's silence on hearing loss in particular is a coverage gap this adjudication surfaces and does not fill.

Four — the discipline of asking what compensation costs. CEI's allostatic framing — that a long series of individually adaptive responses can produce a disease by their cumulative price rather than by any one of them failing — is a genuinely useful analytical instrument, and it is one the corpus already uses without naming: the compensatory hyperactivity of the surviving locus coeruleus, the sprouting, the rising output, all of which are adaptive locally and destructive cumulatively. CEI states the principle better than the corpus has.

What adopting these does not concede. None of the four requires the valence inversion. One can hold that Aβ has a physiological role at picomolar concentrations *and* that it is pathogenic at the concentrations the disease produces; that glutamatergic signalling is a convergence node *and* not the only one; that compensation is costly *and* that the thing being compensated for is amyloid rather than the reverse. The four adoptions are separable from the claim the next section rejects, and stating them separately is what distinguishes an adjudication from a defeat.

XII. What CEI Gets Wrong, Graded

The inversion is the load-bearing claim and it fails on three independent lines, in increasing order of decisiveness.

First — the concentration argument does not survive human material. CEI's position requires that Aβ's toxicity at micromolar concentrations is an artefact of experiments that used more peptide than a brain contains. That argument was answerable in 2005 and is not answerable now. Amyloid-β protein dimers isolated *directly from Alzheimer brains* potently inhibit long-term potentiation, enhance long-term depression, reduce dendritic spine density in normal rodent hippocampus, and disrupt the memory of a learned behaviour in normal rats — at low-nanomolar concentrations, from human tissue, without synthetic preparation (Shankar and colleagues, 2008). The species that the disease actually makes, at the concentrations the disease actually reaches, is synaptotoxic.

Second — the causal arrow between activity and amyloid runs the wrong way for CEI. CEI requires that loss of excitation induces amyloid production, as compensation. The measured relation is the opposite: synaptic activity regulates interstitial-fluid amyloid-β levels in vivo, with more activity producing more Aβ (Cirrito and colleagues, 2005). Zott and colleagues (2019) then closed it into a cycle running in the disinhibition direction — Aβ suppresses glutamate reuptake, hyperactivity follows preferentially in already-active neurons, and Aβ dimers sustain it. For CEI to be right, this entire relation would have to be an epiphenomenon of a deeper compensation, and no evidence is offered that it is.

Third, and decisively — the direction of therapeutic benefit. If amyloid were a compensatory sensitiser propping up a failing excitatory system, removing it should accelerate decline. Two large trials removed it and slowed decline: lecanemab and donanemab both produced statistically significant reductions in clinical progression against placebo. One may hold any view of the size or clinical meaningfulness of those effects — the corpus's own steelman volume takes that argument seriously in both directions — but the *sign* is not in dispute, and the sign is what CEI predicts wrongly. A compensator whose removal helps was not compensating.

A fourth line, on the developmental arm. CEI's strong form — that elevated Aβ causes disease only through prenatal or perinatal exposure — is refuted by inducible adult-onset APP models, in which switching on the transgene in adulthood produces the pathology. The cleanest counterexample is human: individuals with partial trisomy 21 who have the Down syndrome phenotype but only two copies of *APP* reach their seventies without dementia and without amyloid. Amyloid dose, not developmental excitatory under-building, is what determines Alzheimer's disease

in Down syndrome.

The grading, stated fairly. CEI is **citation-honest at the level of individual findings** — the fact-check found its sources faithfully reported — and **built by directional selection at the level of argument**, retaining the findings that support an inversion and omitting the ones that refute it. Its physiological core is real; its mechanism as stated is contradicted wherever it has been tested. The underlying preprint was never peer-reviewed and is essentially uncited. The honest description is not "refuted rival" but **dormant minority provocation**: a hypothesis that never entered the field's argument, whose value to this corpus is diagnostic rather than competitive — it marks the places where the corpus's own account is thin.

XIII. The One Fork That Remains Open

Strip away the seizure question, which Section VI settles, and the throughput question, which Section IX dissolves, and one genuine disagreement remains. It is not about excitability at all.

The fork. Is amyloid a cause or a response? CEI says response — the brain makes it because excitatory transmission is failing. The corpus says cause. And the honest statement is that the *evidence adjudicating this* is the anti-amyloid trial data and the genetics, not anything in the electrophysiology. The E/I literature cannot settle it, which is why an adjudication conducted on the seizure datum can be decisive about the seizure datum and must stop at the edge of the valence question.

Why the corpus's answer is nonetheless the better-supported one. Three reasons, none of them electrophysiological. The autosomal-dominant mutations are in the amyloid pathway and they cause the disease with near-complete penetrance. The protective *APP* A673T variant reduces Aβ production and protects. And removal slows decline. Against these, CEI offers a reinterpretation of each — the mutations act by disturbing development, the trials are confounded — but reinterpretation is what a framework does when the primary data run against it.

Why the corpus should not treat the fork as closed. Because the corpus already maintains, deliberately, a self-contained steelman of amyloid primacy written to close a fairness gap, and that document's §VI forward-references a rebuttal to the permissive-trigger reading that **does not exist anywhere in the corpus**. The valence question is therefore live inside the corpus's own architecture, independently of Moosmann, and it is logged as an open structural gap. This volume does not attempt to close it; it notes that CEI's inversion and the steelman's permissive-trigger argument are two attacks on the same joint, arriving from opposite directions, and that the corpus owes an answer to the joint rather than to either attacker.

What this means for how CEI should be filed. Not as a refuted framework to be recorded and forgotten, and not as a live rival to be reconciled as a peer. As a **standing provocation on the valence question**, retained because the corpus's answer to that question is currently asserted in five documents and argued in none. When that argument is written, CEI's status should be revisited; until then, a hypothesis that attacks an unargued claim is doing the corpus a service.

XIV. The Validity Ledger

Tier I — Established

Alzheimer's disease carries a large excess of epileptiform activity, including clinically silent mesial-temporal seizures invisible to scalp recording. Directly measured with intracranial foramen ovale electrodes (Lam and colleagues, 2017).

Cortical neurons near amyloid plaques are hyperactive, and the hyperactivity is attributable to relatively reduced synaptic inhibition; silenced and hyperactive neurons coexist in the same field. (Busche and colleagues, 2008.)

Restoring parvalbumin-interneuron function rescues network dysfunction and cognition in an amyloid model. (Verret and colleagues, 2012.)

Levetiracetam suppresses abnormal spike activity and reverses synaptic and cognitive deficits in hAPP mice, uniquely among the antiepileptics tested; low-dose levetiracetam normalises dentate/CA3 hyperactivation and improves memory in amnestic MCI. (Sanchez and colleagues, 2012; Bakker and colleagues, 2012.)

Seizures can arise from excess inhibition — the mechanism CEI invokes is real. Enhanced tonic GABA-A inhibition in typical absence epilepsy (Cope and colleagues, 2009); seizures with enhanced cortical GABAergic inhibition in two ADNFLE models (Klaassen and colleagues, 2006).

Picomolar Aβ42 enhances hippocampal LTP and memory. (Puzzo and colleagues, 2008.)

Human-brain-derived Aβ dimers impair LTP, enhance LTD, reduce spine density and disrupt learned behaviour at low-nanomolar concentrations. (Shankar and colleagues, 2008.)

Synaptic activity raises interstitial-fluid Aβ. (Cirrito and colleagues, 2005.)

Tier II — Well-supported inference

Alzheimer's epileptiform activity is of the disinhibition type and not the absence type. The inference rests on the convergent direction of four independent interventions and one direct measurement of the cause of hyperactivity. It is an inference from therapeutic and mechanistic sign rather than a direct demonstration that inhibitory tone is globally reduced.

Anti-excitability treatment benefits the subset with documented epileptiform activity and not the unselected population. Supported by the restriction of benefit to the epileptiform subgroup in a randomised crossover trial (Vossel and colleagues, 2021) and by a missed primary endpoint with a subgroup-confined executive signal in a pilot (Sen and colleagues, 2024). Both are small; neither was powered for the subgroup as a primary analysis.

CEI's genetic arm and the corpus's reelin axis describe the same receptor lesion. (Chen and colleagues, 2010.) Two independent derivations converging on ApoER2 — corroboration, not proof of either framework.

Tier III — Synthesis, and graded as such

The two-axis reading: throughput and excitability are distinct variables that decouple in the diseased brain, and the models were measuring different ones. This is the volume's organising claim. It is supported by the coexistence of hypometabolism and epileptiform activity and by the within-field coexistence of silenced and hyperactive neurons, but no study has measured both variables in the same human brains across the staging series.

The epoch table of Section X. Assembled across populations, species and instruments that were never co-registered. A schema for organising a contradictory literature, not an observation.

That memantine's benefit is better explained by clipping pathological extrasynaptic current than by CEI's disinhibition account. More parsimonious and mechanistically connected to the reuptake lesion, but the discriminating experiment — interneuron-dependence of benefit — has not been performed. **CEI's reading is not refuted here.**

Not established, and not required

That inhibitory tone is globally elevated or globally reduced in the Alzheimer brain. The argument requires only that it is locally reduced where hyperactivity is measured, which is established. Global statements about a distribution with two tails are avoided throughout.

That CEI's therapeutic programme would fail. Predicted, on the sign argument of Section VI. Never tested. Recorded as a prediction rather than a result.

XV. Predictions and Falsification

On the sign of the lesion. Selective GABAergic antagonism or NMDA enhancement, given in an amyloid or tauopathy model, will worsen network dysfunction and cognition rather than improving them. *If lifting the inhibitory brake improves the phenotype, the disinhibition reading of the Alzheimer seizure datum is refuted and CEI's is supported — this is the cleanest single experiment in the volume, and CEI's own therapeutic programme specifies it.*

On the memantine mechanism. Memantine's benefit will track markers of extrasynaptic glutamate and impaired reuptake rather than parvalbumin-interneuron density, and will not fall in proportion to interneuron loss. *If benefit is abolished by selective interneuron depletion with reuptake intact, CEI's account of memantine is right and the corpus's is wrong.*

On the two axes. In a single cohort measured longitudinally, throughput (FDG-PET or an equivalent) and excitability (high-density or intracranial electrophysiology) will diverge rather than covary across the preclinical-to-prodromal transition. *If the two track together in the same individuals, the two-axis claim is unnecessary and the single-dial framing was adequate all along.*

On the epoch schema. Co-registering both variables across the staging series in one cohort will place declining throughput earlier than rising local excitability, with an interval of overlap. *If rising excitability precedes any measurable throughput decline, the compensated-prodrome row of Section X's table is wrong and CEI loses the one epoch it was granted.*

On patient selection, which is the testable clinical claim. A trial of levetiracetam enrolling only patients with epileptiform activity documented by an adequately sensitive method — extended or overnight recording rather than a routine scalp electroencephalogram — will show a benefit that unselected trials do not. *If a properly powered, biomarker-selected trial is also null, the anti-excitability programme fails and the disinhibition model loses its therapeutic corollary even if its physiology stands.*

On the prevalence of the phenotype. Applying the sensitivity of intracranial or overnight recording to an unselected Alzheimer cohort will find epileptiform activity in a substantially larger fraction than routine electroencephalography reports. *If sensitive recording finds no more than routine recording, then the epileptiform subgroup is genuinely small, and the anti-excitability approach is a niche therapy rather than an under-recognised one.*

XVI. Therapeutic Consequences

Three consequences follow, and they are unusually concrete for a document about a theoretical disagreement.

One — stop running unselected anti-excitability trials. The ladder of Section VII is a sequence of increasingly rigorous studies producing decreasingly visible effects, and the pattern is fully explained by dilution: the effect is real in a subgroup and the subgroup was not enriched. The corpus's recommendation is that no further trial of an anti-excitability agent in Alzheimer's disease should enrol without documented epileptiform activity as an inclusion criterion, and that the documentation should use a method sensitive enough to find it — which routine scalp electroencephalography, on Lam's evidence, is not.

Two — the enabling technology is a detection method, not a drug. The therapeutic question here is bottlenecked on identification. Levetiracetam is generic, cheap, well tolerated in this population — ILiAD reported no withdrawals and no significant difference in side effects against placebo — and has a plausible mechanism and a normalised imaging target. What is missing is a practical way to know which patients have the phenotype. Extended ambulatory recording, overnight recording, and magnetoencephalography are the candidates; none is currently standard in a memory clinic. **The corpus's position is that a screening protocol for occult epileptiform activity is a higher-value target than any new compound in this space.**

Three — treat the two axes with different agents and do not expect one drug to do both. Nothing that lowers excitability restores throughput, and nothing that restores throughput lowers noise. The volume's framing implies a combination logic — an anti-excitability agent for the noise in patients who have it, and whatever preserves synaptic and metabolic capacity for the signal — and implies that trials measuring only cognition will continue to produce ambiguous results because cognition is a summary of both variables.

What this volume does not recommend. It does not recommend CEI's programme — NMDA enhancement, glutamate-trafficking enhancement, selective GABAergic antagonism. On the evidence assembled in Section VI, lifting the inhibitory brake in a network whose brake is already failing is the wrong direction, and the prediction that it would worsen the phenotype is stated in Section XV so that it can be held against this document if it is wrong.

A caution on timing that applies to all three. Every intervention discussed here is symptomatic. Lowering excitability normalises a physiological abnormality and improves a memory score; it does not address whatever produced the abnormality. The corpus's position elsewhere —

that the disease's decisive events are preclinical — is not modified by this volume, and the anti-excitability programme should be understood as a way of removing a source of ongoing injury in an already-established disease, not as a disease-modifying strategy.

XVII. Coda — Two Instruments, One Reading

The disagreement this volume adjudicates lasted as long as it did because both sides were right about what they had measured and wrong about what they had measured it with.

One instrument was pointed at the whole brain and reported that it was doing less: less glucose, less traffic, less of the work that thinking consists of. From that reading came a hypothesis that the disease is a failure of excitation, and that everything else — the plaques, the tangles — is the brain's long attempt to shout over its own deafness. It is a humane hypothesis and it is wrong in an interesting way.

The other instrument was pointed at a few hundred neurons beside a plaque and reported that they were doing too much: firing without being asked, in a tissue whose restraints had been eaten. From that reading came a hypothesis that the disease ends in an unbraked network consuming itself. That one is right, about the terminus, and about the seizures, and it has the drugs to show for it.

Neither instrument was lying. A brain is not a dial. It is a distribution, and in this disease the distribution spreads at both ends at once — the mean falling while the tail rises, the signal thinning while the noise grows, until the tissue is simultaneously too quiet to think with and too loud to be safe in.

What the corpus got wrong was smaller and worse than either hypothesis. It took a rival's argument, silently reversed its sign, and printed it as exposition — not out of carelessness about the facts, which were right, but out of carelessness about the *source*, which was not. The result looked like agreement and was actually an unnoticed contradiction sitting inside the corpus's own pages for a year. That is the failure this volume was written to repair, and it is worth naming plainly, because a synthesis that cannot quote its opponents accurately has no way of finding out that it is wrong.

Low signal, high noise. It is a description of a diseased cortex. For a while it was also a description of the literature about one.

A Note on Provenance and Verification

Every citation new to this volume was verified against the publisher or PubMed Central record during its preparation, not reconstructed from memory: Palop and Mucke 2016 (PMID 27829687); Verret 2012 (PMID 22541439); Bakker 2012 (PMID 22578498); Sanchez 2012 (PMID 22869752);

Vossel 2021; Sen 2024 (ILiAD); Lam 2017 (PMID 28459436); Busche 2008 (PMID 18802001); Zott 2019 (PMID 31395777); Puzzo 2008 (PMID 19118188); Shankar 2008; Cope 2009 (PMID 19966779); Klaassen 2006; Chen 2010 (PMID 20547867); Reisberg 2003 (PMID 12672860). Citations imported from prior volumes of this corpus — Palop 2007, Cirrito 2005 — are reproduced in the form in which they were verified there.

Two limitations are recorded. The PubMed connector was unavailable in the session in which this volume was prepared, so verification proceeded through publisher and PubMed Central records retrieved directly; author lists for Vossel 2021 and Sen 2024 are given in truncated form because the full lists could not be confirmed from the sources retrieved, and are marked accordingly. And the lecanemab and donanemab trial results invoked in Section XII are cited from the corpus's existing verified records rather than re-verified here; the claim they support in this volume concerns only the *direction* of effect, which is not in dispute.

This volume supersedes the treatment of chronic excitatory insufficiency in *Convergent Synaptic Collapse* §4 as the corpus's account of the excitatory–inhibitory question. That chapter remains as the framework-by-framework survey entry it was written to be, corrected under C-016; the adjudication it gestured at is here.

References

Bakker A, Krauss GL, Albert MS, Speck CL, Jones LR, Stark CE, Yassa MA, Bassett SS, Shelton AL, Gallagher M. Reduction of hippocampal hyperactivity improves cognition in amnestic mild cognitive impairment. *Neuron*. 2012;74(3):467–474. PMID 22578498.

Busche MA, Eichhoff G, Adelsberger H, Abramowski D, Wiederhold KH, Haass C, Staufenbiel M, Konnerth A, Garaschuk O. Clusters of hyperactive neurons near amyloid plaques in a mouse model of Alzheimer's disease. *Science*. 2008;321(5896):1686–1689. PMID 18802001.

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

Cirrito JR, Yamada KA, Finn MB, Sloviter RS, Bales KR, May PC, Schoepp DD, Paul SM, Mennerick S, Holtzman DM. Synaptic activity regulates interstitial fluid amyloid-β levels in vivo. *Neuron*. 2005;48(6):913–922. DOI: 10.1016/j.neuron.2005.10.028

Cope DW, Di Giovanni G, Fyson SJ, Orbán G, Errington AC, Lőrincz ML, Gould TM, Carter DA, Crunelli V. Enhanced tonic GABA-A inhibition in typical absence epilepsy. *Nature Medicine*. 2009;15(12):1392–1398. PMID 19966779.

Klaassen A, Glykys J, Maguire J, Labarca C, Mody I, Boulter J. Seizures and enhanced cortical GABAergic inhibition in two mouse models of human autosomal dominant nocturnal frontal lobe epilepsy. *Proceedings of the National Academy of Sciences USA*. 2006;103(50):19152–19157.

Lam AD, Deck G, Goldman A, Eskandar EN, Noebels J, Cole AJ. Silent hippocampal seizures and spikes identified by foramen ovale electrodes in Alzheimer's disease. *Nature Medicine*. 2017;23(6):678–680. PMID 28459436.

Moosmann B, Sohre S. Chronic excitatory insufficiency as proximate cause of Alzheimer's disease. Oskar Fischer Prize 2020 submission #145; experimental companion preprint, *bioRxiv* 418566 (2018). *Neither peer-reviewed nor substantively cited; see Section XII.*

Palop JJ, Chin J, Roberson ED, Wang J, Thwin MT, Bien-Ly N, Yoo J, Ho KO, Yu GQ, Kreitzer A, Finkbeiner S, Noebels JL, Mucke L. Aberrant excitatory neuronal activity and compensatory remodeling of inhibitory hippocampal circuits in mouse models of Alzheimer's disease. *Neuron*. 2007;55(5):697–711. DOI: 10.1016/j.neuron.2007.07.025

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

Puzzo D, Privitera L, Leznik E, Fà M, Staniszewski A, Palmeri A, Arancio O. Picomolar amyloid-β positively modulates synaptic plasticity and memory in hippocampus. *Journal of Neuroscience*. 2008;28(53):14537–14545. PMID 19118188.

Reisberg B, Doody R, Stöffler A, Schmitt F, Ferris S, Möbius HJ; Memantine Study Group. Memantine in moderate-to-severe Alzheimer's disease. *New England Journal of Medicine*. 2003;348(14):1333–1341. PMID 12672860.

Sanchez PE, Zhu L, Verret L, Vossel KA, Orr AG, Cirrito JR, Devidze N, Ho K, Yu GQ, Palop JJ, Mucke L. Levetiracetam suppresses neuronal network dysfunction and reverses synaptic and cognitive deficits in an Alzheimer's disease model. *Proceedings of the National Academy of Sciences USA*. 2012;109(42):E2895–E2903. PMID 22869752.

Sen A, et al. Safety, tolerability, and efficacy outcomes of the Investigation of Levetiracetam in Alzheimer's disease (ILiAD) study: a pilot, double-blind placebo-controlled crossover trial. *Epilepsia Open*. 2024;9(6):2353–2364. *Author list truncated — full list not confirmed in the preparing session.*

Shankar GM, Li S, Mehta TH, Garcia-Munoz A, Shepardson NE, Smith I, Brett FM, Farrell MA, Rowan MJ, Lemere CA, Regan CM, Walsh DM, Sabatini BL, Selkoe DJ. Amyloid-β protein dimers isolated directly from Alzheimer's brains impair synaptic plasticity and memory. *Nature Medicine*. 2008;14(8):837–842. DOI: 10.1038/nm1782

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

Vossel K, Ranasinghe KG, Beagle AJ, La A, Ah Pook K, Castro M, Mizuiri D, Honma SM, Chan D, Bonasera SJ, Kramer JH, Miller BL, Nagarajan SS. Effect of levetiracetam on cognition in patients with Alzheimer disease with and without epileptiform activity: a randomized clinical trial. *JAMA Neurology*. 2021;78(11):1345–1354. *Author list truncated — full list not confirmed in the preparing session.*

Zott B, Simon MM, Hong W, Unger F, Chen-Engerer HJ, Frosch MP, Sakmann B, Walsh DM, Konnerth A. A vicious cycle of β amyloid-dependent neuronal hyperactivation. *Science*. 2019;365(6453):559–565. PMID 31395777.