THE UNNETTED GOVERNOR

What the Somatostatin Neuron Does in Alzheimer's Disease — the Cell That Is Its Circuit's Own Self-Regulation, Whose Peptide Licenses the Disposal of Amyloid in the Very Layer Its Axon Occupies, and Which Wears No Perineuronal Armour

The Peptide and the Cell · The Somatostatin Loop · The Coincidence of Address
The Cell Without Armour · The Grid That Fails Around It

A Dissertation Submitted in Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy in Neuroscience

Prepared under the Organic Network Synthesis methodology

AdultCognitiveDisease.com

Benjamin Aaron Gustafsson

July 2026

The inhibitory-cell volume the corpus had not written. Its companions took the excitatory neuron, the microglion, the matrix, and the first neuron to tangle. This one takes the interneuron — and finds that the corpus's recurring motif, the neuron built without a net, has a second and much larger instance than the locus coeruleus.

"A governor is not a brake. A brake is worked from outside; a governor is driven by the very speed it limits. Break a brake and the engine is merely unguarded. Break a governor and the engine is unguarded and the fault compounds — because the faster it runs, the less able the thing that would have slowed it becomes."

— ONS Methodology Working Notes, 2026

For Peter Davies, Robert Katzman and Robert Terry, who in 1980 measured a peptide and found the second-oldest fact in this disease — and for the forty-six years in which nobody could say what it meant.

THE COLLAPSE FRAMEWORK THE CELLULAR VOLUMES

The excitatory neuron — Convergent Synaptic Collapse

The microglion — The Corruption of the Gardener · The Gardener's Restraint

The matrix — The Architect's Scaffold · The Coerulean Shears

The first neuron to tangle — The First Ember · The Coerulean Pincer

The inhibitory neuron ← The Unnetted Governor — this volume

The corpus had written the cell that fires, the cell that eats, the lattice that shields, and the nucleus that ignites. It had not written the cell that holds the rest in check. This volume supplies it, and finds that the omission was consequential: the somatostatin interneuron is not merely another casualty but a term in the disease's own progression, because its peptide sets the clearance rate of the substance whose accumulation defines the illness. It is also the corpus's net-less neuron at cortical scale — the same structural argument that explained why the locus coeruleus tangles first, applied to a third of the cortical inhibitory population, and predicting the exact order in which the human single-cell atlases find the inhibitory classes fall.

Contents

Abstract

The oldest neurochemical finding in Alzheimer's disease that is not cholinergic is somatostatinergic. In 1980, four years after the cholinergic deficit was described, Davies, Katzman and Terry reported that somatostatin-like immunoreactivity is reduced in the Alzheimer cortex; Beal and colleagues confirmed it as a widespread cortical loss in 1986, most severe in the temporal lobe; and Morrison and colleagues found somatostatin immunoreactivity inside the neuritic plaque itself in 1985. Forty-six years later, the finding has never been overturned and has never been placed. This dissertation exists to place it.

Its first move is to separate two things the literature has run together under one word. *Somatostatin* names a peptide and it names a cell, and the two have different biologies, different failure modes, and different timetables. The peptide is a secreted inhibitory neuromodulator that augments the M-current and damps the excitability of the neurons it reaches. The cell — the somatostatin-expressing GABAergic interneuron, about thirty per cent of the cortical inhibitory population — is a dendrite-targeting cell that supplies *feedback* inhibition: it is driven by the very pyramidal cells it inhibits, and it inhibits the distal dendrites where those cells' long-range inputs arrive. That architecture makes the somatostatin neuron the physical substrate of a circuit's self-regulation. It is a governor in the mechanical sense — a device whose output is a function of the speed it is meant to limit.

The dissertation's second move is to state a loop. The peptide is not merely a brake on firing; it licenses the disposal of amyloid. Somatostatin up-regulates neprilysin, the brain's principal Aβ42-degrading protease; somatostatin deficiency raises Aβ42 in a manner the original report likened to a presenilin mutation; the pathway runs through α-endosulfine and the KATP channel; and the receptor subtypes that carry it, SST1 and SST4, are redundantly required. Somatostatin also binds Aβ directly and diverts its assembly, slowing plaque deposition by a route independent of neprilysin. So the peptide is, in two separate senses, an anti-amyloid agent. And amyloid, reciprocally, is what injures the cell that makes it: the largest human single-nucleus atlas of the disease finds loss of somatostatin interneurons in the *early*, cognitively silent phase of pseudoprogression — before excitatory neurons, before parvalbumin and VIP interneurons, and, tellingly, before tangle deposition in the same tissue. The loop closes on itself: the cell that licenses the disposal of the poison is the cell the poison disables first, and each turn of the loop lowers the licence.

The dissertation's third move is a claim about address, and it is the volume's signature contribution. The somatostatin-receptor-dependent neprilysin that clears Aβ is reduced, when those receptors are deleted, *specifically in the stratum lacunosum-moleculare of the hippocampus*. That layer is not an arbitrary compartment. It is where the entorhinal perforant path terminates on the distal dendrites of CA1 pyramidal cells; it is where the somatostatin-expressing oriens–lacunosum-moleculare interneuron sends its axon; and it is, by the same finding, where

amyloid disposal depends on somatostatin receptors. One layer, three functions — the entorhinal input, the dendritic brake upon it, and the enzymatic clearance of the peptide the input's activity releases. The somatostatin neuron's axonal territory *is* the amyloid disposal site. No paper in either literature states this, because the two literatures — circuit physiology and amyloid catabolism — do not cite each other.

The fourth move explains selectivity, and it recovers a motif this corpus has met before. The parvalbumin interneuron, the somatostatin cell's sibling, is the archetypal wearer of the perineuronal net; the somatostatin cell, with rare exceptions, is not. The corpus has already established that net-bearing neurons resist tau while net-less nuclei succumb, and that microglia must degrade the net before they reach what it covers. The somatostatin neuron presents the microglion with no such wall — and microglia in Alzheimer models preferentially engulf *inhibitory* synapses, and are transcriptionally tuned to GABA. The prediction that falls out is exactly the order the human atlas reports: the unnetted inhibitory cell goes first, the netted one goes late. The somatostatin neuron is the corpus's net-less neuron at cortical scale.

Around this spine the dissertation assembles the mechanisms the question demands — the cell's calcium problem (a dendrite built for supralinear NMDA-receptor-dependent integration, in a cell that does not carry its sibling's namesake buffer, in a circuit running hot); p38α mitogen-activated protein kinase as the transducer that carries both the microglial and the amyloid insult into the entorhinal synapse, with the honest complication that the best-controlled study of the amyloid arm finds it calcium-*flux*-independent; and the four distinct senses in which the disease is a loss of self-regulation. It then separates the trajectory into the dysfunctional phase, in which the cell is silent in output while hyperactive in body, and the degenerative phase, in which it is simply gone; and walks the anatomical line — entorhinal cortex, CA1, prefrontal cortex — that the corpus's temporal architecture predicts and the postmortem gradient confirms.

Finally, it corrects the question it was asked. Grid-cell failure is real, early, and entorhinal; it appears at the first fibrillar amyloid, it tracks path-integration failure, and it is measurable in young human ε4 carriers decades before disease. But it is not the somatostatin cell's direct doing, and the dissertation says so. When somatostatin interneurons of the medial entorhinal cortex are silenced, grid cells are unaffected; parvalbumin silencing is what degrades the hexagonal code. The anatomy explains the physiology: entorhinal somatostatin cells preferentially inhibit layer III–V pyramidal cells and largely spare layer II, which is where grid cells live. The somatostatin cell is therefore not the proximate cause of grid failure — it is the cause of the *aperiodic* spatial code's failure, of the deep-layer output stage, and of the chronic excitatory-inhibitory derangement in which the layer II grid cells themselves are eventually killed. The grid fails *around* the somatostatin neuron, not at its hands. Stating that clearly is worth more than a tidier story, and it yields the sharper prediction: rescue the somatostatin cell and one should recover the deep-layer aperiodic code and the CA1 hand-off first, and the hexagonal code only later and only by way of keeping layer II alive.

Every joint of the argument is graded in a three-tier validity ledger that is explicit about which claims are measured, which are synthesised across literatures that do not cite one another, and which are inference. Six falsifiable predictions are stated. The therapeutic corollary is double-edged and the dissertation says so: the peptide arm argues for restoring somatostatin signalling, and the circuit arm has been read by some as arguing for *suppressing* somatostatin-cell activity — two opposite instructions that are only reconcilable by phase, and a chronic suppression experiment has already failed behaviourally. The governor must be repaired, not merely turned down.

I. The Question, and the Two Things Called Somatostatin

The question this dissertation was set is deceptively simple: *what is the role of the somatostatin neuron?* It is worth saying at the outset why the question has proved so hard to answer well, and the reason is a piece of vocabulary. The word "somatostatin" designates two entirely different biological objects, and almost every confusion in this literature can be traced to their conflation.

The peptide and the cell. Somatostatin — somatotropin release-inhibiting factor, in its original endocrine name — is a small cyclic peptide, existing in the brain principally as the fourteen- and twenty-eight-residue forms, secreted from dense-core vesicles, acting at five G-protein-coupled receptors, and functioning as an inhibitory neuromodulator. The somatostatin *neuron* is a GABAergic inhibitory interneuron that happens to express that peptide, and that constitutes roughly thirty per cent of the cortical interneuron population. The peptide can be depleted from a living cell; the cell can be lost with its peptide; the peptide can be present in a cell that has stopped delivering it usefully; and the peptide can be found, as Morrison and colleagues found it in 1985, deposited in a plaque, far from any cell at all (Morrison and colleagues, 1985). These are four distinct states, and the classical measurements — radioimmunoassay of tissue somatostatin-like immunoreactivity — cannot distinguish among them.

The oldest finding, and its misreading. In 1980, Davies, Katzman and Terry reported that somatostatin-like immunoreactivity is reduced in the cerebral cortex of Alzheimer cases (Davies and colleagues, 1980). Six years later, Beal and colleagues extended this to a widespread cortical reduction — most severe in the temporal lobe, significant in frontal and occipital cortex, absent in subcortical regions — and drew the conclusion the field then adopted: that "the reduction in somatostatin immunoreactivity in Alzheimer's disease may be caused by degeneration of intrinsic somatostatin cortical neurons" (Beal and colleagues, 1986). That conclusion was a hypothesis about cells inferred from a measurement of peptide, and it was not unreasonable; but it foreclosed the alternative — that the peptide might fall *before* the cell does, and that its fall might itself be pathogenic — for roughly three decades. The modern single-cell atlases have finally made the

distinction empirically, and they show that both things happen, at different times, in a definite order. Restoring the distinction is the first work of this dissertation, because the peptide's biology and the cell's biology enter the disease by different doors.

Why the answer matters beyond bookkeeping. The stakes of the distinction are not taxonomic. If the lesion is loss of the cell, the therapeutic question is one of neuroprotection or replacement, and the window is early. If the lesion is loss of the peptide from a surviving cell, the therapeutic question is one of receptor pharmacology, and the window is wider — one can supply an agonist to a receptor whose cell is silent but alive. And if, as this dissertation will argue, the two are joined in a loop, then the window closes at a rate the loop determines, and the intervention must break the loop rather than treat either arm. The corpus's recurring instruction — that timing is the therapy — applies with unusual force here, because the somatostatin lesion is one of the very few in Alzheimer's disease that is documented, in human tissue, to occur in the cognitively silent phase.

The shape of the answer. This dissertation will argue that the somatostatin neuron's role in Alzheimer's disease is that of a *governor* — the mechanical sense of the word, a device whose output is a negative function of the quantity it limits — and that the disease is, in this cell's part of the story, the failure of a governor in four separate senses that arrive in sequence. The peptide's brake on excitability fails. The cell's own governors — cholinergic, disinhibitory — fail. The circuit's feedback loop, of which the cell is the return limb, fails. And the disposal system that the peptide licenses fails, so that the substance that injures the cell accumulates in the very layer the cell's axon occupies. Each failure feeds the next, and the whole assembly runs down along a definite anatomical line: entorhinal cortex first, CA1 next, prefrontal cortex last. We begin with what the cell is.

II. The Governor What the Somatostatin Neuron Is Built to Do

Before any pathology, the cell must be described as a piece of circuitry, because everything the disease does to it is intelligible only against what it was for.

Feedback, not feedforward. Cortical and hippocampal inhibition divides, to a first approximation, into two functional classes distinguished by what they target and by what drives them. The parvalbumin-expressing basket cell is a *perisomatic* inhibitor: it clamps the soma and axon initial segment, it is fast, it is driven substantially by feedforward input, and it enforces the timing of the output spike — it is a clock. The somatostatin-expressing interneuron is a *dendrite-targeting* inhibitor: its axon ascends to the distal dendritic tuft, it is slower, and — the decisive property — it is

driven by the local pyramidal cells themselves. It receives its excitation from the very population it inhibits. This is the canonical architecture of a negative-feedback controller, and it is why the somatostatin cell can be said, without metaphorical strain, to *be* the circuit's self-regulation. The parvalbumin cell tells the circuit *when*; the somatostatin cell tells the circuit *how much*.

What dendritic inhibition actually controls. The functional consequence of targeting the dendrite rather than the soma was established directly in hippocampal CA1: dendritic inhibition, not perisomatic inhibition, is the dominant regulator of the input–output transformation of the pyramidal cell, and it works by gating the dendritic electrogenesis that drives burst firing (Lovett-Barron and colleagues, 2012). The somatostatin cell does not merely subtract from the pyramidal cell's output; it determines whether the dendrite is permitted to generate the regenerative events that convert synaptic input into bursts. Remove it, and the pyramidal dendrite becomes a permissive rather than a selective integrator. This is the physiological definition of a lost gain control.

Which inputs the somatostatin cell gates. In CA1, the somatostatin-expressing oriens–lacunosum-moleculare (OLM) cell has an unusually specific job. Its soma lies in stratum oriens; its axon ascends to stratum lacunosum-moleculare, the layer in which the perforant path from the entorhinal cortex terminates on CA1 distal dendrites. Optogenetic dissection showed that OLM cells therefore *differentially* weight the two great inputs to CA1: they facilitate transmission of intrahippocampal information from CA3 while reducing the influence of extrahippocampal input from the entorhinal cortex, and they are directly driven by subcortical cholinergic afferents (Leão and colleagues, 2012). The OLM cell is a switch on the entorhinal-to-hippocampal channel. Hold that fact; the entire amyloid argument of Sections IV and V will return to that one layer.

What the somatostatin cell does in the entorhinal cortex itself. In the medial entorhinal cortex the somatostatin interneuron has a different and equally specific target profile, and it is the fact on which this dissertation's correction to the grid-cell question turns. Somatostatin-positive interneurons of the medial entorhinal cortex preferentially inhibit layer III–V pyramidal cells and leave layer II neurons largely unaffected; their inhibition is dual, carried by GABA and by the peptide itself, and it is prolonged rather than fast — sustained suppression rather than the rapid clamp the parvalbumin cell delivers; and behavioural work assigns this motif to the selective modulation of working-memory formation rather than to the retrieval of already-learned spatial routes (Kecskés and colleagues, 2020). Two things follow. First, the entorhinal somatostatin cell is a *deep-layer* controller — it governs the output stages of the entorhinal cortex, which project onward to the hippocampus and back to the cortex. Second, it does not, anatomically, have its hands on layer II, where the grid cells are densest. Section XV will show that the physiology agrees with the anatomy, and that this is the honest answer to the grid-cell question.

The peptide is part of the machine, not a marker. It is easy to read "somatostatin-expressing" as a labelling convenience, as though the peptide were a convenient antigen and the cell's real business were GABA. It is not. Somatostatin, applied to hippocampal CA1 pyramidal neurons, elicits a steady outward current and selectively augments the M-current — the non-inactivating, voltage-dependent potassium conductance that sets a neuron's willingness to fire repetitively — and, in the same experiments, muscarinic cholinergic agonists antagonise that current (Moore and colleagues, 1988). The peptide is, in other words, a slow excitability brake acting on the same conductance that acetylcholine releases, and the two transmitters are reciprocal at that conductance. When the disease reduces cortical somatostatin, it is not removing a stain. It is removing a tonic hyperpolarising influence on the pyramidal cells of the very circuits that will later become hyperexcitable. That the same interneuron delivers both a fast ionotropic brake (GABA) and a slow modulatory brake (somatostatin) to the same dendritic territory is what makes it a governor rather than merely an inhibitor.

The scale of the population. Somatostatin interneurons constitute on the order of thirty per cent of cortical GABAergic interneurons and populate layers II through VI; they are morphologically and molecularly diverse, comprising the Martinotti cells of neocortex, the OLM cells of hippocampus, and the deep-layer-targeting cells of entorhinal cortex, among others (reviewed in Sandoval and Witt, 2024). This is not a boutique population. It is one of the two great inhibitory classes of the mammalian cortex, and the one charged with gain rather than timing.

III. The First Brake Somatostatin as a Regulator of Excitability

The dissertation's first mechanistic claim is the simplest, and it is the one the 1980 finding should have made obvious: the peptide whose loss defines the somatostatinergic deficit is itself an anti-excitatory agent, so that its loss is not a marker of damage but a *removal of restraint*.

The M-current, and what augmenting it means. The M-current is the slowly activating, non-inactivating potassium conductance that opposes sustained depolarisation; a neuron with a large M-current adapts quickly and fires sparsely, and a neuron with a suppressed M-current fires repetitively and is prone to bursting. Somatostatin, in both its fourteen- and twenty-eight-residue forms, augments this current in CA1 pyramidal neurons and produces a steady outward current besides (Moore and colleagues, 1988). The peptide is therefore a *tonic* damping influence supplied by the same cell that supplies the phasic dendritic GABA. Where the somatostatin neuron is active, the pyramidal cells around it are less willing to fire repetitively; where it is silent, they are more willing.

Reciprocity with acetylcholine, and why it matters in this disease. The same study found that muscarinic agonists antagonise the M-current — that acetylcholine and somatostatin push the same conductance in opposite directions (Moore and colleagues, 1988). In a disease defined for half a century by cholinergic denervation, that reciprocity is worth stating precisely, because it cuts against a naive reading. The cholinergic deficit *removes* an M-current suppressor, which by itself would reduce excitability; the somatostatinergic deficit *removes* an M-current augmenter, which increases it. These are not two versions of the same lesion. The network-level consequence of the somatostatin loss is disinhibition; the network-level consequence of cholinergic loss on this particular conductance is not. The corpus should not fold the somatostatinergic deficit into the cholinergic one, as the 1980s literature often did on the grounds that both were "neurotransmitter deficits of the Alzheimer cortex." They act on the same knob from opposite sides.

The brake is the first thing lost. The peptide falls early and it falls hard. The reductions are widespread across cortex, most severe in temporal lobe, and of substantial magnitude — on the order of a seventy per cent reduction in frontal cortex and roughly thirty per cent in temporal cortex across the assembled literature (Sandoval and Witt, 2024; Beal and colleagues, 1986). Cerebrospinal-fluid somatostatin is correspondingly reduced, and the reduction appears at preclinical and prodromal stages (Sandoval and Witt, 2024). Whatever else the disease does to this system, it begins by lowering the concentration of a peptide whose demonstrated action is to make cortical pyramidal neurons less excitable. Any account of Alzheimer's network hyperexcitability that does not include this is incomplete.

A caution on the direction of inference. It does not follow from the above that the somatostatinergic deficit *causes* the network hyperexcitability of Alzheimer's disease; the disease removes many restraints at once, and the parvalbumin/Nav1.1 lesion, in which reduced levels of the interneuron-predominant sodium channel subunit drive network hypersynchrony and premature mortality, is documented with a directness the somatostatin arm has not matched (Verret and colleagues, 2012). What follows is narrower and defensible: the somatostatin peptide is a demonstrated excitability brake; it is demonstrably reduced early; therefore its reduction is a plausible contributor to the excitability phenotype, and it is one that acts through a different conductance and a different compartment than the parvalbumin lesion does. The ledger in Section XVI grades this accordingly.

IV. The Licence How the Peptide Governs Amyloid's Disposal

Here the dissertation leaves the physiology of excitability and enters the biochemistry of amyloid, because the somatostatin peptide has a second job that has nothing obvious to do with membrane potential, and it is the job that ties this cell to the disease's defining lesion.

Somatostatin up-regulates neprilysin. Neprilysin is the principal Aβ-degrading protease of the brain, and its capacity is the rate-limiting term in the clearance side of the amyloid balance. In a screen of candidate effectors, only somatostatin up-regulated neprilysin activity in primary cortical neurons; genetic deficiency of somatostatin altered hippocampal neprilysin activity and its subcellular localisation and increased the quantity of the hydrophobic forty-two-residue form of Aβ — and the original report drew the comparison explicitly, noting that the effect resembled that of the presenilin mutations that cause familial Alzheimer's disease (Saito and colleagues, 2005). This is the load-bearing finding of the section, and it is worth pausing on how strange it is. A neuropeptide released by an inhibitory interneuron sets the activity of the enzyme that disposes of the peptide whose accumulation defines the disease. Nothing in the classical picture of interneuron function anticipates this.

The pathway from receptor to protease. The mechanism was subsequently resolved. Somatostatin-evoked Aβ catabolism runs through α-endosulfine, which acts as a negative regulator of neprilysin downstream of somatostatin signalling and is itself degraded by neprilysin; α-endosulfine expression rises in Alzheimer models and in patients; and the pathway is executed at ATP-sensitive potassium channels, such that pharmacological intervention at specific KATP channel subtypes activates neprilysin, reduces Aβ deposition, and rescues memory in transgenic mice (Watamura and colleagues, 2022). The chain is therefore: somatostatin → somatostatin receptor → α-endosulfine → KATP channel → neprilysin → Aβ degradation. It is a defined, druggable pathway, and it begins at the interneuron's secretory vesicle.

Which receptors carry it, and where. The receptor arm has now been resolved as well. Neprilysin is regulated by the SST1 and SST4 receptor subtypes in a redundant manner; deletion of both reduced presynaptic neprilysin, and — this is the detail on which the next section turns — the reduction was localised *specifically to the hippocampal stratum lacunosum-moleculare*; crossing the double deletion onto an amyloid-precursor-protein transgenic background worsened Aβ pathology; and treatment with an SST1/SST4 agonist ameliorated it (Nilsson and colleagues, 2026). The receptor pharmacology of amyloid clearance is thus established both by loss and by gain of function, and it has an address.

A second, receptor-independent anti-amyloid action. The peptide also acts on Aβ directly. Somatostatin binds the human Aβ peptide and favours the formation of distinct oligomeric species, forming mixed assemblies with Aβ42 that interfere with fibrillisation (Wang and colleagues, 2017). In vivo, somatostatin deficiency produced a subtle but significant increase in the density of cortical Aβ amyloid plaques in aged APP knock-in mice by a mechanism *independent* of neprilysin —

neprilysin transcript, protein, and activity were unchanged across genotypes — with the peptide instead interfering with early stages of Aβ assembly and manifesting as high-molecular-weight oligomers in brain extracts (Williams and colleagues, 2023). Two independent routes, then: an enzymatic licence and a direct chaperone-like interference. The honest caveat, stated by the reviewers of this literature, is that the in vitro binding work used high peptide concentrations and did not account for normal proteolytic processing, so the physiological weight of the direct-binding arm is less secure than that of the neprilysin arm (Sandoval and Witt, 2024).

And the peptide is found in the plaque. The oldest anatomical observation in this literature belongs here rather than in a footnote. Somatostatin immunoreactivity is present in the neuritic plaques of Alzheimer patients (Morrison and colleagues, 1985). Read against the biochemistry just assembled, that observation acquires a meaning it could not have had in 1985: the peptide that binds Aβ and diverts its assembly is found *inside the deposit*, which is what one would expect of a molecule that engages the aggregate and is sequestered by it. This dissertation does not claim that plaque-bound somatostatin is the mechanism of the cortical deficit — the deficit is far too large for sequestration alone to explain, and the cell loss documented in Section V is real. But the co-localisation is consistent with an interaction rather than a coincidence, and it is the first hint in the historical record that the peptide and the plaque are in a relationship.

The claim, stated. Somatostatin is an anti-amyloid peptide. It licenses the enzymatic disposal of Aβ42 through a defined receptor–endosulfine–KATP–neprilysin chain, and it interferes with Aβ assembly directly. Whatever the somatostatin neuron is doing for cognition, it is simultaneously operating one of the brain's principal amyloid-clearance controls. That is a role no other interneuron class is known to hold.

V. The Loop The Cell That Disposes of the Poison Is the Cell the Poison Disables First

If Section IV established that the peptide governs amyloid's disposal, this section establishes the return arm, and with it the loop that organises the whole dissertation.

Amyloid injures the somatostatin cell, and does so first. The strongest human evidence is the Seattle Alzheimer's Disease Brain Cell Atlas. Across eighty-four donors spanning the pathological range, multiomic and spatial profiling of the middle temporal gyrus, ordered along a pseudoprogression axis, resolved the disease into two phases: an early phase of slowly increasing pathology, inflammatory microglia, reactive astrocytes, *loss of somatostatin-positive inhibitory neurons*, and

an oligodendrocyte-precursor remyelination response; and a later phase of exponentially increasing pathology with loss of excitatory neurons and of parvalbumin- and VIP-expressing inhibitory subtypes (Gabitto and colleagues, 2024). Donors in the early phase have no cognitive deficits. Two features of this result are decisive for the present argument. First, the somatostatin interneuron is lost *before* the excitatory neuron and *before* its parvalbumin sibling — it is, on this evidence, among the earliest cellular casualties of the disease. Second, the atlas notes that somatostatin neurons are not known to accumulate neurofibrillary tangles, and that their early loss precedes tangle deposition in the same tissue, which points the aetiology of their loss toward Aβ rather than toward tau.

The cell's dysfunction is documented before its death. In the APP/PS1 model, two-photon imaging of the O-LM interneuron found severely impaired synaptic rewiring at both its input and its output, with learning-dependent remodelling of its connections disrupted, and identified reduced cholinergic drive from the septo-hippocampal pathway as a critical mechanism linking that dysfunction to memory impairment (Schmid and colleagues, 2016). The dendrite-targeting interneuron is thus not merely lost late; it is functionally deranged early, and the derangement is traceable to the loss of one of its own governing inputs.

The cell is hyperactive while its circuit is disinhibited. The activity phenotype is, at first sight, paradoxical, and the paradox is informative. In awake APP/PS1 mice, calcium imaging shows somatostatin-expressing interneurons to be *hyperactive* while parvalbumin interneurons are *hypoactive*, and only the somatostatin hyperactivity correlates with proximity to amyloid plaque; excitatory neuron activity is reduced overall (Algamal and colleagues, 2022). A cell can be hyperactive in its soma and hypofunctional in its effect, and there are at least two established routes by which somatostatin-cell hyperactivity produces net *disinhibition* rather than net inhibition. The first is circuit topology: in neocortex, somatostatin cells inhibit other interneurons as well as pyramidal cells, so over-driving them can lift inhibition from the principal cell. This is not speculation — in a TDP-43 model of amyotrophic lateral sclerosis and frontotemporal dementia, over-active somatostatin interneurons *reduced* inhibitory signalling onto layer 5 pyramidal neurons and thereby contributed to neuronal damage, and focal ablation of those interneurons restored normal pyramidal activity and diminished neurodegeneration (Zhang and colleagues, 2016). The second is compartmental: a cell firing more action potentials while its peptide stores are depleted and its axonal arbour is disrupted delivers less of its two brakes to the dendrite even as it spikes more. The reading this dissertation adopts is that the hyperactivity is a *compensatory or dysregulated* state of a governor whose setpoint has been broken — the phenotype of a controller with a failed sensor, not of a controller working harder to good effect.

The human cell loss, and its regional signature. Immunohistochemical characterisation of interneurons in Alzheimer cases found significantly fewer somatostatin interneurons in the temporal cortex, with parvalbumin-expressing cells unchanged, and noted the regional variation in interneuron density that any such comparison must control for (Waller and colleagues, 2020). This

is a small, focused study rather than an atlas, but it agrees in sign and in cell type with the atlas result, and it agrees with the classical peptide measurements' regional gradient — temporal worst (Beal and colleagues, 1986).

The loop, stated formally. Assemble the arms. The somatostatin peptide licenses neprilysin-mediated Aβ42 degradation and interferes with Aβ assembly (Saito and colleagues, 2005; Watamura and colleagues, 2022; Nilsson and colleagues, 2026; Wang and colleagues, 2017; Williams and colleagues, 2023). Aβ pathology is associated with early dysfunction and early loss of the somatostatin interneuron, ahead of the excitatory neuron and ahead of tangle formation (Gabitto and colleagues, 2024; Schmid and colleagues, 2016; Algamal and colleagues, 2022). Therefore: rising Aβ reduces the population and the peptide output of the cells whose peptide licenses Aβ's disposal, which lowers disposal capacity, which raises Aβ. The loop has positive gain. It is, in the strict sense, self-amplifying, and it is self-amplifying at the very earliest, cognitively silent phase of the disease — which is precisely where a self-amplifying loop does the most damage to the eventual outcome and where it is least visible to any clinical measure.

And a second turn of the same loop, through activity. There is a further multiplier the corpus has met before. Interstitial-fluid Aβ is dynamically regulated by synaptic activity on a timescale of minutes to hours, an effect traceable primarily to synaptic vesicle exocytosis (Cirrito and colleagues, 2005). The somatostatin neuron's job is to limit the activity of the pyramidal cells it inhibits. Lose the governor, and those cells fire more; and cells that fire more release more Aβ. So the loop closes twice: once biochemically, through the clearance licence, and once physiologically, through the production rate. The governor's failure raises Aβ from both ends of the balance simultaneously. This dissertation regards the double closure as its second contribution, and grades it as synthesis — each arm is established, the joining is not measured in a single experiment.

VI. The Coincidence of Address One Layer, Three Functions

Here the dissertation makes its signature claim, and it arises from putting two facts side by side that no paper puts side by side, because they belong to literatures that do not cite each other.

The first fact: the somatostatin cell's axon occupies stratum lacunosum-moleculare. The CA1 somatostatin-expressing OLM interneuron has its soma in stratum oriens and sends its axon to stratum lacunosum-moleculare, where it inhibits the distal apical dendrites of CA1 pyramidal cells — the compartment in which the perforant path from the

entorhinal cortex terminates. That placement is the physical basis of its documented function: it weights the entorhinal input down relative to the CA3 input (Leão and colleagues, 2012), and dendritic inhibition of exactly this kind is the dominant regulator of the CA1 pyramidal cell's input–output transformation (Lovett-Barron and colleagues, 2012).

The second fact: somatostatin-receptor-dependent neprilysin is concentrated in the same layer. When the SST1 and SST4 receptors are deleted together, the resulting reduction in presynaptic neprilysin is found *specifically in the hippocampal lacunosum-moleculare layer* (Nilsson and colleagues, 2026). The amyloid-degrading capacity that depends on somatostatin receptors is, in other words, localised to the same thin lamina.

The joined claim. Stratum lacunosum-moleculare is therefore the site of three things at once: the arrival of the entorhinal cortex's long-range input onto CA1; the axonal terminal field of the somatostatin interneuron that gates that input; and the somatostatin-receptor-dependent neprilysin capacity that disposes of the Aβ which that input's activity helps release. The somatostatin neuron's axonal territory *is* the amyloid disposal site. Its peptide is released where its clearance licence is redeemed, and both coincide with the synapse whose activity supplies the substrate. This is an unusually tight piece of biological economy, and it has a corollary that is the point of stating it: an injury to the somatostatin cell's axon in this layer is simultaneously an injury to the entorhinal-to-hippocampal gate and to the local amyloid-clearance capacity. There is no need to postulate two lesions. One lesion, in one lamina, produces both.

Why the corpus should care. This dissertation's companion volumes have repeatedly found that the disease's most consequential mechanisms are ones in which a single structure serves several functions, so that its failure is multiply expressed — the sulfated matrix that both shields the neuron and stages the reelin signal; the mitochondria-associated membrane at which the lipid and energetic disorders are one buckling; the locus coeruleus whose transmitter is its poison. The lacunosum-moleculare coincidence is the same kind of object, and it belongs on the corpus's list. It also supplies the anatomical hinge for the trajectory of Section XIV, because a lesion at the entorhinal input layer of CA1 is precisely a lesion at the junction between the disease's first cortical region and its second.

Grading the claim honestly. The two facts are each measured. The joining is a synthesis, and it carries one genuine uncertainty: the Nilsson finding localises the *receptor-dependent* neprilysin deficit to the layer, and neprilysin in this layer is described as presynaptic — that is, on axon terminals — but the experiment does not establish that the somatostatin released by the OLM cell is the ligand that normally maintains it, nor that the somatostatin source and the neprilysin-bearing terminal are the same synapses. The alternative is that entorhinal terminals bear the neprilysin and some other somatostatin source maintains it. Section XVII states the experiment that would settle this, and Section XVI grades the claim as a strong synthesis with one unmeasured joint.

VII. The Cell Without Armour Why the Somatostatin Neuron and Not Its Sibling

The order of loss in the human atlas — somatostatin early, parvalbumin late (Gabitto and colleagues, 2024) — demands an explanation, because the two cell classes are neighbours in the same tissue, exposed to the same amyloid, and derive from the same embryonic eminence. This dissertation's explanation recovers a structure the corpus has already established for a different cell.

Parvalbumin cells wear the net; somatostatin cells do not. The perineuronal net is a condensed lattice of chondroitin- and heparan-sulfate proteoglycans that ensheathes certain neurons, and its principal cortical clientele has been known for three decades: Wisteria floribunda agglutinin-labelled nets surround parvalbumin-containing neurons (Härtig and colleagues, 1992). Somatostatin interneurons, with the exception of a small subset, are not so ensheathed. The corpus has already established what the net does for the neurons that have it: the subcortical nuclei attacked earliest and hardest by tau are precisely those devoid of an aggrecan-based net, while net-ensheathed neurons even in the thick of the pathology rarely tangle (Morawski and colleagues, 2010), and the configuration and integrity of perineuronal nets tracks with resistance to Alzheimer's disease across human cases (de Vries and colleagues, 2024).

The motif, at cortical scale. The companion volume *The First Ember* built its account of why the locus coeruleus tangles first on exactly this asymmetry: the blue nucleus is the archetype of the net-less neuron, unshielded and unstaged. That argument concerned a few thousand brainstem cells. The present observation is the same argument applied to roughly a third of the cortical interneuron population. The somatostatin interneuron is the net-less neuron at cortical scale — and the corpus's own logic then predicts that it should be reached earlier than its netted sibling by whatever the net protects against. The human atlas reports exactly that order (Gabitto and colleagues, 2024). This dissertation regards the agreement as one of the stronger structural corroborations the corpus has produced, because the prediction and the observation come from entirely independent literatures.

The cell supplies the armour it does not wear. There is a second asymmetry running through the same two cell classes, and it sharpens the first. Reelin — the developmental architect that persists in the adult brain as a brake on tau phosphorylation, signalling through ApoER2 and VLDLR to Disabled-1 — is not made by cortical neurons generally. In the adult cortex it is expressed preferentially in GABAergic neurons (Pesold and colleagues, 1998), and the subpopulation was then characterised precisely: many interneurons expressing neuropeptide Y or

somatostatin are reelin-immunopositive, and *none* of those expressing parvalbumin are. Morphologically, the reelin-storing cells are the layer I horizontal cells, the layer II–V bitufted neurons, and the deep-layer Martinotti cells — and they secrete reelin *into the perineuronal net*, where it acts extrasynaptically on the neurons it surrounds (Pesold and colleagues, 1999). The same study reports that Dab1, the adaptor through which reelin's message is read, is expressed predominantly in pyramidal neurons.

Set that beside this section's asymmetry and the two turn out to be one asymmetry read from opposite ends. The somatostatin class is the class that does not wear the net and does make the reelin; the parvalbumin class wears the net and makes none. The cell that fails first is therefore the cell that manufactures a protective ligand, deposits it into a lattice that does not cover the manufacturer, and addresses it to a third cell that reads it. Its loss is not only the withdrawal of two brakes on excitability, which Sections III and V describe; it is also the withdrawal of a source of the tau brake from the cortical compartment in which that brake is staged — and it occurs in the preclinical phase, ahead of the excitatory neuron and ahead of the netted sibling (Gabitto and colleagues, 2024). The dissertation grades the anatomy as established (Pesold and colleagues, 1998, 1999; Härtig and colleagues, 1992) and the consequence as an inference not yet measured: no study has shown that somatostatin-cell loss lowers reelin availability in the cortex those cells innervate. The measurement is straightforward in principle — reelin output at somatostatin and Martinotti cells against intact and depleted populations in staged tissue — and it would convert this paragraph from a joining of two literatures into a result.

What the net actually protects against, in this instance. The mechanism is not merely mechanical shielding. Microglia must degrade the perineuronal net before they reach what it covers, and they do so in the Alzheimer brain — microglia facilitate the loss of perineuronal nets in Alzheimer's disease (Crapser and colleagues, 2020). That is a demolition step: a delay, a rate-limiter, a requirement for enzymatic work before the cell surface beneath is accessible. The somatostatin neuron presents no such step. Whatever the microglion is doing to the netted parvalbumin cell only after it has stripped the lattice, it can do to the somatostatin cell immediately. The net is not only a shield; it is a *latency*. Section IX develops what the microglion then does.

A caution about the exception. A subset of somatostatin interneurons — those expressing Kv3 potassium channels — do carry perineuronal nets, and cortical net coverage is not perfectly partitioned between the two classes. The claim this dissertation makes is therefore statistical rather than absolute: the parvalbumin class is the predominant netted population and the somatostatin class is predominantly unnetted, so that as populations they differ in exposure. Whether the netted somatostatin subset is spared in the disease at the rate this argument predicts is, so far as the present search of the literature found, unmeasured — and it is one of the cleanest tests this dissertation can propose. Section XVII states it.

VIII. Calcium A Dendrite Built to Integrate, in a Cell That Carries No Buffer in Its Name

The corpus's bioenergetic and calcium volumes have made calcium dyshomeostasis a recurring upstream term. Here it is asked to do something narrower and more specific: to explain why *this* cell's dendrite, in *this* circuit, under *this* excitatory load, is a poor risk.

The dendrite is built for supralinearity. Dendrite-targeting interneurons of the hippocampus — the OLM cell among them — exhibit robust NMDA-receptor-dependent supralinear integration of dendritic inputs; in the companion neurogliaform population, the supralinear dendritic calcium transients depend on L-type voltage-gated calcium channels and on release from intracellular stores, such that blocking the sarco/endoplasmic-reticulum calcium ATPase abolishes the non-linear calcium summation while preserving the voltage supralinearity; and calcium-permeable AMPA receptors have been implicated in dendritic calcium transients in oriens interneurons (Griesius and colleagues, 2025). The functional point is that this dendrite does not sum its inputs passively. It amplifies coincident input non-linearly, and the amplification is carried substantially by calcium — through NMDA receptors, through L-type channels, and through store release. A dendrite designed to amplify is a dendrite designed to accumulate calcium when its inputs are excessive, and the somatostatin cell's inputs are, by construction, the pyramidal cells whose runaway it exists to prevent.

The buffer that is in the sibling's name and not in this cell's. The nomenclature of cortical interneurons encodes a real asymmetry that is worth making explicit because it is so easily overlooked. The parvalbumin cell is named for a calcium-binding protein expressed in it at high concentration — a slow-onset, high-capacity cytosolic buffer that shapes calcium transients and, by that route, the cell's firing and its tolerance of load. The somatostatin cell is named for a secreted peptide message. The parvalbumin cell carries a buffer in its name; the somatostatin cell carries a signal. This is not an argument that somatostatin cells lack all calcium buffering — subsets express calbindin, and buffering capacity is not exhausted by the named proteins — and this dissertation grades the claim as suggestive rather than established. But the asymmetry is real, it runs in the same direction as the net asymmetry of Section VII, and it belongs in an account of why one class of interneuron is reached before the other.

The cross-disease anchor: what happens to this cell when a circuit runs hot. The clearest demonstration that the somatostatin interneuron is the cell that fails when excitatory drive is excessive comes from outside Alzheimer's disease altogether, and its independence is its value. In

experimental epilepsy, the prediction was that GABA-containing inhibitory neurons would be damaged by seizure activity; instead they survived, and what was found was a nearly complete loss of the adjacent somatostatin-containing interneurons and mossy cells that normally activate them (Sloviter, 1987). Sustained excitatory load kills the somatostatin cell while sparing the basket cell. That is a forty-year-old result in a different disease, and it says that the vulnerability described in this dissertation is a property of the cell type rather than an artefact of amyloid.

The Alzheimer-specific load. Two features of the disease raise this cell's calcium burden specifically. The first is that somatostatin interneurons near amyloid plaques are hyperactive, and only their hyperactivity correlates with plaque proximity (Algamal and colleagues, 2022) — the cell is being driven harder precisely where the pathology is densest. The second is the loop of Section V: as the governor's output falls, the pyramidal cells it fails to restrain fire more, and those cells are the somatostatin cell's own excitatory drive. The cell is therefore caught in a feedback arrangement that converts its own failure into more input to itself. A dendrite built for supralinear calcium amplification, in a cell class without its sibling's namesake buffer, receiving progressively more drive because it is progressively less able to limit that drive, is a defensible account of why this cell dies early. Section XVI grades the individual links.

IX. The Transducer The p38 Mitogen-Activated Protein Kinase

If the previous sections describe insult and exposure, this one asks what carries the insult into the synapse, and the answer the literature best supports for the entorhinal circuit is p38α mitogen-activated protein kinase. The section also contains this dissertation's most important internal correction, because the naive chain — calcium overload activates p38 — is not what the best-controlled experiment on the amyloid arm actually shows.

p38α carries the amyloid insult into the synapse. Aβ increases p38 MAPK activity, and inhibiting p38 abolishes Aβ-induced loss of dendritic spines. The mechanism was dissected carefully: the synaptic loss and the p38 activation both required *glutamate binding* to NMDA receptors — the NMDA-receptor antagonist APV prevented both — but neither required calcium flux, since the calcium chelator BAPTA and open-channel blockers failed to prevent them, and the pathway proved independent of G proteins as well (Birnbaum and colleagues, 2015). This is a metabotropic-like action of the NMDA receptor, and it is the cleanest available account of how Aβ reaches p38.

The correction this forces. The outline this dissertation was asked to develop pairs "calcium handling" with "MAPK − p38," and the natural reading is that disordered calcium activates the kinase. That reading is not supported for the amyloid arm, and this dissertation declines to assert it. What the evidence supports is a *two-arm* structure: calcium load is one insult to the somatostatin cell's amplifying dendrite (Section VIII), and p38α activation is a second, largely calcium-flux-independent insult reaching the same dendrite by way of glutamate binding at the NMDA receptor (Birnbaum and colleagues, 2015). The two arms converge on the same compartment without one causing the other. Saying so costs the argument a tidy causal chain and buys it accuracy; the ledger records the correction.

p38α carries the microglial insult too, and it does so in the entorhinal cortex. The second route into p38 is inflammatory, and the demonstration is regionally exact. Microglial receptor for advanced glycation end products drives Aβ-induced synaptic depression and long-term-depression impairment in the entorhinal cortex: blockade of RAGE prevented the Aβ-induced increase in phosphorylated p38 MAPK, and JNK inhibition rescued normal long-term depression (Origlia and colleagues, 2010). Here in one experiment are the microglion, the amyloid, the kinase, and the entorhinal cortex — the four terms this dissertation must join. The microglial arm of Section X and the kinase arm of this section are not separate stories; the microglion's RAGE signalling is one of the two established routes to neuronal p38 activation in the region where the disease begins cortically.

Inhibiting p38α rescues the entorhinal cortex specifically, and early. The therapeutic converse has been demonstrated in the same region and at the right time. An isoform-selective p38α inhibitor rescues *early* entorhinal cortex dysfunction in a mouse model of Alzheimer's disease (Rutigliano and colleagues, 2018). More broadly, the isoform-selective p38α inhibitor MW150 attenuates disease progression in Alzheimer models, with reduced synaptic loss, reduced tau phosphorylation, and partial normalisation of electrophysiology, and with benefit mediated primarily by rescue of neuronal function rather than by effects on the primary pathologies (Roy and colleagues, 2015). That the benefit is a *functional* rescue rather than a pathological one is important for this dissertation's two-phase framing: it is the signature of an intervention that acts on the dysfunctional phase.

The honest state of the somatostatin-specific link. What is not established is a somatostatin-cell-specific role for p38α. The kinase's documented effects here are on the pyramidal dendrite, on the entorhinal synapse, and on the microglion. There is a suggestive but distant observation that somatostatin receptors of the SST2 and SST4 subtypes promote intense and prolonged p38 activity — but the primary demonstrations of somatostatin-receptor-driven p38 signalling are in non-neural, largely oncological preparations, where sustained p38 activity serves the peptide's anti-proliferative programme. Whether the somatostatin receptor drives p38 in the Alzheimer cortex, and in which cell, is not shown. This dissertation therefore places p38α in the

argument as the *transducer of the insult that reaches the somatostatin cell's circuit*, not as a somatostatin-specific effector, and grades the receptor-to-kinase link as unestablished in brain. Section XVII proposes the cell-type-resolved experiment.

X. The Microglion The Cell Stripped From Outside

The somatostatin neuron is not only injured from within by the load its own failure creates. It is dismantled from outside, by a cell whose behaviour in this disease the corpus has already characterised at length, and which turns out to have a specific appetite for inhibitory synapses.

Complement and microglia mediate early synapse loss. The foundational demonstration is that complement and microglia mediate early synapse loss in Alzheimer mouse models, with C1q required and the loss preceding plaque deposition (Hong and colleagues, 2016). Synapse elimination in this disease is therefore not a passive consequence of neuronal death but an executed programme, and it runs early.

And the programme is biased toward inhibitory synapses. The bias is documented directly. In Alzheimer's disease mouse models, C1q-dependent synapse elimination is divided between two cell types with opposite preferences: microglia preferentially remove *inhibitory* synapses while astrocytes preferentially remove excitatory ones, with microglial lysosomes containing more inhibitory synaptic material; C1q deletion reduced glial engulfment and rescued synapse density (Dejanovic and colleagues, 2022). And microglia are equipped to recognise the inhibitory synapse as such: GABA-receptive microglia, acting through microglial GABA-B receptors, selectively interact with and sculpt inhibitory cortical synapses without affecting excitatory ones, running a distinct transcriptional remodelling programme when they do so (Favuzzi and colleagues, 2021). The machinery for a microglion to find, recognise, and eliminate an inhibitory synapse specifically is established.

The joint with Section VII. Now combine the appetite with the exposure. Microglia preferentially engulf inhibitory synapses; microglia must first degrade the perineuronal net to reach the neurons it covers (Crapser and colleagues, 2020); the parvalbumin interneuron is covered and the somatostatin interneuron is not (Härtig and colleagues, 1992). The prediction is immediate: of the two great inhibitory classes, the unnetted one should be reached sooner. The human atlas finds the unnetted one lost in the early phase and the netted one in the late phase (Gabitto and colleagues, 2024). This dissertation offers that convergence as its explanation for the order of inhibitory loss in Alzheimer's disease, and grades it as a strong inference assembled from measured parts rather than

as a measured finding — no experiment has yet compared microglial engulfment of somatostatin-cell synapses with that of parvalbumin-cell synapses in the same tissue while manipulating the net. That experiment is stated in Section XVII, and it is the single most decisive test this volume can propose.

The corpus's microglial verdict applies here without amendment. The companion atlas of microglial dysfunction ranked ageing and the TREM2/APOE genetic axis as the dominant drivers of the microglial turn, with proteinopathy and the NLRP3 inflammasome major, and the noradrenergic setpoint a threshold-setter distinguished by being earliest and reachable rather than largest. Nothing in the present volume disturbs that ranking. What the present volume adds is a *target*: whatever turns the microglion, the cell it reaches first among inhibitory neurons is the one without a wall, and that cell is the governor.

XI. Loss of Self-Regulation Four Senses, in Order

"Loss of self-regulation" is the phrase that best captures what this cell's failure does, and it is worth disaggregating, because the phrase is true in four distinct senses that arrive in a definite order and that have different therapeutic implications.

Sense one: the peptide's own brake is withdrawn. Somatostatin augments the M-current and damps repetitive firing in the neurons it reaches (Moore and colleagues, 1988). The peptide falls early and steeply — most severely in frontal cortex, next in temporal — and falls in cerebrospinal fluid at preclinical and prodromal stages (Davies and colleagues, 1980; Beal and colleagues, 1986; Sandoval and Witt, 2024). This is the first restraint removed, and it is removed at the level of a conductance rather than a synapse.

Sense two: the cell's own governors lapse. The somatostatin interneuron is itself governed, and its governors fail. The septo-hippocampal cholinergic input that drives the OLM cell is reduced, and that reduction is the mechanism the imaging study identified as linking O-LM dysfunction to memory impairment (Schmid and colleagues, 2016); OLM cells are among the cells that receive direct subcortical cholinergic drive and mediate nicotine's effects on plasticity (Leão and colleagues, 2012). A controller whose own input is degraded cannot hold its setpoint, and the phenotype of a controller with a degraded sensor is not silence but *erratic output* — which is what the hyperactivity near plaques looks like (Algamal and colleagues, 2022).

Sense three: the circuit's negative feedback opens. Because the somatostatin cell is the feedback limb — driven by the pyramidal cells it inhibits, inhibiting the dendrites where their

long-range input arrives (Lovett-Barron and colleagues, 2012; Leão and colleagues, 2012) — its failure converts a closed loop into an open one. The consequence is the well-documented network phenotype of the disease: aberrant excitatory activity and compensatory inhibitory remodelling, non-convulsive seizure activity, GABAergic sprouting, and plasticity deficits (Palop and colleagues, 2007). It is worth being precise about attribution here: the best-documented single molecular lesion behind that phenotype is the parvalbumin/Nav1.1 deficit, whose correction restores gamma oscillations, reduces hypersynchrony, and rescues memory and premature mortality (Verret and colleagues, 2012). The somatostatin contribution is to gain rather than to synchrony, and the two are complementary rather than competing.

Sense four: the disease's own loop closes. The fourth sense is the one this dissertation adds. The governor's failure raises Aβ from both ends — by lowering the somatostatin-licensed neprilysin capacity (Saito and colleagues, 2005; Watamura and colleagues, 2022; Nilsson and colleagues, 2026) and by permitting the activity-dependent release of Aβ from the pyramidal cells it no longer restrains (Cirrito and colleagues, 2005) — and rising Aβ injures the governor further (Gabitto and colleagues, 2024; Algamal and colleagues, 2022). The system loses the capacity to regulate the very substance whose accumulation defines it. This is self-regulation lost at the level of the disease process rather than at the level of the circuit, and it is the sense in which the somatostatin lesion is not merely a consequence of Alzheimer's disease but a term in its progression.

Why the order matters. The four senses are not simultaneous, and the ordering is the therapeutic content of the section. The peptide falls before the cell dies; the cell's governors lapse before the cell's output fails; the circuit opens before the cell is lost; and the disease's loop is running throughout, at a gain set by how much of the first three has already happened. An intervention aimed at sense one — receptor agonism — remains available for as long as the receptors are there, which is longer than the cells are. An intervention aimed at sense four must arrive before the loop's gain has done its compounding, which is very early indeed.

XII. The Dysfunctional Phase The Cell Falls Silent Before It Dies

The corpus's temporal architecture insists on separating dysfunction from degeneration, and the somatostatin literature supplies an unusually clean instance of that separation — clean enough that the human atlas resolved the two phases statistically.

The human two-phase result. Pseudoprogression analysis of the middle temporal gyrus across eighty-four donors resolved an early phase — slow accumulation of pathology, inflammatory microglia, reactive astrocytes, loss of somatostatin-positive inhibitory neurons, oligodendrocyte-precursor remyelination — from a late phase of exponentially increasing pathology with loss of excitatory neurons and of parvalbumin and VIP interneuron subtypes; and donors in the early phase have no cognitive deficits, which the authors read as a preclinical stage (Gabitto and colleagues, 2024). The somatostatin lesion is therefore, in human tissue, a *preclinical* lesion.

What "dysfunction" consists of, concretely. Before the cell is lost, four things are documented. Its peptide output falls, measurably, in tissue and in cerebrospinal fluid (Davies and colleagues, 1980; Beal and colleagues, 1986; Sandoval and Witt, 2024). Its synaptic connectivity fails to remodel with learning, at both input and output, with reduced cholinergic drive identified as the mechanism (Schmid and colleagues, 2016). Its somatic activity becomes abnormal — hyperactive, and specifically so near plaques (Algamal and colleagues, 2022). And the amyloid-clearance capacity it licenses declines in the lamina its axon occupies (Nilsson and colleagues, 2026, for the receptor dependence; Saito and colleagues, 2005, for the licence). None of these requires the cell to be dead. All of them are, in principle, reversible.

A negative result that matters: no tangles. The atlas notes that somatostatin neurons are not known to accumulate neurofibrillary tangles, and that their early loss precedes tangle deposition in the same tissue (Gabitto and colleagues, 2024). This is a genuine constraint on the aetiology, and this dissertation takes it seriously. Whatever injures the somatostatin cell in the early phase is not, on this evidence, its own tau pathology. The candidates the present volume assembles are amyloid acting through the routes of Sections IV, V, and IX; microglial elimination of its synapses (Section X); and the excitatory load its own failure creates (Section VIII). It is worth noting how sharply this distinguishes the somatostatin cell from the corpus's other early casualty, the locus coeruleus neuron, whose early lesion *is* a tau lesion. Two cells, both unnetted, both early, injured by different agents. The net explains exposure; it does not dictate the assailant.

The corollary. A preclinical, reversible, cell-type-specific lesion in a population that operates an amyloid-clearance control is close to the ideal target that the whole corpus has been arguing must exist. Section XVIII takes up what to do about it, and why the obvious thing is not obviously right.

XIII. The Degenerative Phase The Cell Is Lost, and Then the Rest

The late phase requires less argument, because it is the phase the classical literature described, but two features of it deserve statement.

The loss is real, selective, and regionally graded. Somatostatin interneuron numbers are reduced in the temporal cortex of Alzheimer cases while parvalbumin-expressing cells are not (Waller and colleagues, 2020); somatostatin-like immunoreactivity is reduced widely across cortex, most severely in the temporal lobe, with frontal and occipital involvement and subcortical sparing (Beal and colleagues, 1986); and the reduction in frontal cortex, across the assembled literature, is of the order of seventy per cent (Sandoval and Witt, 2024). The classical authors' inference — that the peptide loss reflected degeneration of intrinsic somatostatin cortical neurons — turns out to be correct *for this phase*, and wrong only in having been applied to the whole disease.

The order of the collapse. What the single-cell era adds is sequence. The somatostatin interneuron goes in the early, silent phase; the excitatory neuron and the parvalbumin and VIP interneurons go in the late, exponential phase (Gabitto and colleagues, 2024). That order has an interpretation this dissertation is willing to defend as a hypothesis and grades as such: the excitatory neuron's late death is partly a *consequence* of the inhibitory governor's early loss, by way of the excitotoxic route documented in the cross-disease anchor of Section VIII, where over-active or absent somatostatin interneurons produce net disinhibition and pyramidal degeneration, and focal removal of the deranged interneurons rescues the pyramidal cells (Zhang and colleagues, 2016). The atlas itself gestures in this direction, noting that somatostatin-neuron loss may create an excitatory-to-inhibitory imbalance leading to higher excitability, with implications for pathological tau propagation and for the epilepsy susceptibility of Alzheimer patients (Gabitto and colleagues, 2024). If that is right, the late phase's headline casualty — the pyramidal neuron — dies in a circuit whose brake was removed a decade earlier.

What the disease then loses, functionally. With the governor gone, the pyramidal dendrite becomes a permissive integrator (Lovett-Barron and colleagues, 2012); the entorhinal-to-CA1 channel loses its weighting relative to the CA3 channel (Leão and colleagues, 2012); the deep entorhinal output layers lose their sustained suppression and their working-memory modulation (Kecskés and colleagues, 2020); the network becomes hyperexcitable and prone to non-convulsive seizure activity (Palop and colleagues, 2007); and the local amyloid-clearance licence in stratum lacunosum-moleculare lapses (Nilsson and colleagues, 2026). These are not five deficits. They are five readings of one cell's absence.

XIV. The Anatomical Trajectory Entorhinal Cortex, CA1, Prefrontal Cortex

The disease walks a line, and the somatostatin lesion walks it too. This section assembles what is known at each station and is explicit about where the evidence is human and where it is model.

Station one: the entorhinal cortex. The entorhinal cortex is where Alzheimer's disease becomes a cortical disease, and it is where the somatostatin cell's regional job is best defined. Entorhinal somatostatin interneurons preferentially inhibit layer III–V pyramidal cells, deliver dual GABAergic and peptidergic inhibition, produce sustained rather than fast suppression, and serve working-memory formation rather than retrieval (Kecskés and colleagues, 2020). The disease reaches this region early and by more than one route: entorhinal tau pathology produces excitatory cell loss and network alteration (Fu and colleagues, 2017); microglial RAGE signalling drives Aβ-induced synaptic depression and impaired long-term depression here, through p38 (Origlia and colleagues, 2010); and a p38α inhibitor rescues *early* entorhinal dysfunction (Rutigliano and colleagues, 2018). The entorhinal cortex is therefore the station at which the transducer of Section IX, the microglion of Section X, and the somatostatin cell of Section II are all documented in the same tissue.

Station two: CA1, and the layer that joins the two stations. The entorhinal output arrives at CA1's distal dendrites in stratum lacunosum-moleculare, where the OLM cell's axon gates it (Leão and colleagues, 2012) and where somatostatin-receptor-dependent neprilysin capacity resides (Nilsson and colleagues, 2026). The OLM cell's dysfunction in amyloid models is documented at both its input and its output, with reduced cholinergic drive as mechanism (Schmid and colleagues, 2016), and dendritic inhibition of this kind is the dominant regulator of the CA1 pyramidal cell's input–output transformation (Lovett-Barron and colleagues, 2012). CA1 is thus not merely the second station of the trajectory; it is the station at which the first station's output is received, and the somatostatin cell is stationed at the junction. A lesion here degrades the hand-off itself. The consequence has been measured: in an APP knock-in model, hippocampal CA1 remapping is disrupted while spatial responsiveness is relatively preserved, entorhinal spatial tuning is severely lost, and the fast gamma oscillations that couple entorhinal cortex to CA1 are substantially impaired (Jun and colleagues, 2020). The coupling between the first two stations fails.

Station three: the prefrontal cortex. The prefrontal station is the least resolved at cell-type level and the most striking at peptide level. Somatostatin-like immunoreactivity is reduced in frontal cortex (Beal and colleagues, 1986), and across the assembled literature the frontal reduction is the largest reported — on the order of seventy per cent, against roughly thirty per cent in temporal

cortex (Sandoval and Witt, 2024). This dissertation notes, and does not resolve, the apparent tension between that gradient and the temporal-worst gradient of the classical immunoreactivity studies and of the cell-count study (Beal and colleagues, 1986; Waller and colleagues, 2020): different cohorts, different assays, different disease stages, and a peptide measure that cannot distinguish depletion from cell loss. What can be said is that the prefrontal somatostatin system is heavily involved by the late phase, that prefrontal involvement in the disease's natural history is late, and that the working-memory function which entorhinal somatostatin cells serve at their own station is the function prefrontal cortex is most identified with. The corpus's temporal architecture predicts the prefrontal station to be the last of the three, and the somatostatin data are consistent with that ordering without independently establishing it. The ledger grades this station as the weakest of the three.

The line, stated. Entorhinal cortex, where the governor controls the deep output layers and where the microglion, the amyloid, and the kinase are documented together. CA1, where that output is received in the one lamina that is simultaneously the governor's axonal territory and the amyloid disposal site. Prefrontal cortex, where the peptide's loss is greatest and the function it served is most characteristically expressed. The somatostatin lesion is not scattered across the cortex; it runs along the disease's own path, and at each station it is stationed at the junction rather than in the middle.

XV. Grid-Cell Failure What the Somatostatin Neuron Does Not Explain, and What It Does

This dissertation was asked to connect the somatostatin neuron to grid-cell failure. The honest answer is more interesting than the expected one, and it requires stating a negative result plainly before building on it.

Grid-cell failure in Alzheimer's disease is real, early, and well documented. In a mouse expressing mutant human tau predominantly in the entorhinal cortex, mature tangles in old animals were accompanied by excitatory cell loss, destabilised grid fields, reduced firing rates, altered network activity, and spatial memory deficits (Fu and colleagues, 2017). In the J20 amyloid model, medial entorhinal grid cells showed reduced spatial periodicity, spatial stability, and synchrony with interneurons and head-direction cells, while non-grid spatial coding in the entorhinal cortex and place cells in the hippocampus remained intact; the deficits emerged at the earliest incidence of Aβ fibril deposition and coincided with impaired path integration (Ying and colleagues, 2022). In an APP knock-in model, entorhinal grid cells were nearly absent by seven to

thirteen months, with mild disruption already present at three to five months when behaviour was still intact (Jun and colleagues, 2020). And in humans, young ε4 carriers show reduced grid-cell-like representations, altered navigational behaviour, and increased hippocampal activity consistent with compensation, decades before any disease (Kunz and colleagues, 2015). The phenomenon is not in doubt.

But silencing somatostatin interneurons does not disturb grid cells. The decisive experiment is a pharmacogenetic dissection in freely moving mice. Silencing parvalbumin cells antagonised the hexagonal spatial selectivity of grid cells, especially in layer II, and reduced speed modulation in co-localised speed cells. Silencing somatostatin cells had *no impact on grid cells or speed cells*; what it did instead was decrease the spatial selectivity of cells with discrete, aperiodic firing fields. Border cells and head-direction cells were unaffected by either manipulation (Miao and colleagues, 2017). The two interneuron classes control different space-coding networks, and the grid network is the parvalbumin cell's.

The anatomy explains the physiology. The negative result is not a puzzle once the anatomy is placed beside it. Somatostatin interneurons of the medial entorhinal cortex preferentially inhibit layer III–V pyramidal cells and leave layer II largely unaffected (Kecskés and colleagues, 2020). Layer II is where grid cells are densest — and it is exactly the layer in which parvalbumin silencing had its strongest effect on the hexagonal code (Miao and colleagues, 2017). Two papers from different laboratories, one anatomical and one behavioural-physiological, published three years apart and, so far as this dissertation's search found, not read against one another, give the same answer: the entorhinal somatostatin cell does not have its hands on the grid layer. This dissertation regards the joining of those two results as one of its contributions, and it is a *deflationary* contribution — it removes a connection the field's enthusiasm for somatostatin might otherwise have invented.

What the somatostatin cell does contribute to spatial failure, then. Three things, in ascending order of confidence's cost.

First, and most securely, it governs a *different* space-coding network. Somatostatin silencing degrades the spatial selectivity of aperiodic, discrete-field cells (Miao and colleagues, 2017) — a population whose contribution to navigation is less celebrated than the grid module's but which is no less part of the entorhinal spatial code. Somatostatin-cell failure should therefore degrade the aperiodic code, and the fact that the J20 model preserved non-grid entorhinal spatial coding while destroying grid coding (Ying and colleagues, 2022) is a genuine tension with that expectation which the ledger records and Section XVII proposes to resolve.

Second, it governs the deep-layer output stage. Entorhinal somatostatin cells control layer III–V pyramidal cells and serve working-memory formation (Kecskés and colleagues, 2020); layer III is the origin of the perforant-path input to CA1's stratum lacunosum-moleculare, and layer V is the

entorhinal output to neocortex. A somatostatin lesion therefore degrades what the entorhinal cortex *sends*, not what it *computes* in layer II. That is consistent with the measured failure of entorhinal-to-CA1 gamma coupling and of CA1 remapping in the presence of relatively preserved CA1 spatial responsiveness (Jun and colleagues, 2020) — a hand-off failure rather than a computation failure.

Third, and most inferentially, it contributes to the chronic excitatory–inhibitory derangement in which the layer II grid cells themselves eventually die. Grid-cell loss in the tau model is accompanied by outright excitatory cell loss (Fu and colleagues, 2017); somatostatin-cell loss precedes excitatory-neuron loss in human tissue (Gabitto and colleagues, 2024); and deranged somatostatin interneurons drive pyramidal excitotoxicity in a different neurodegenerative model (Zhang and colleagues, 2016). The proposal is that the somatostatin cell's failure does not perturb the grid computation acutely but degrades the conditions under which the grid cells survive chronically.

The claim, stated in one sentence. The grid fails *around* the somatostatin neuron, not at its hands: acutely, the hexagonal code is the parvalbumin cell's to break; chronically, the somatostatin cell's early failure removes the gain control under which layer II must survive, and takes with it the aperiodic code and the deep-layer hand-off to CA1. Section XVII turns this into a prediction that distinguishes it from the alternative.

XVI. The Validity Ledger

Every load-bearing joint of the argument is graded here, with the experiment that would settle the contested ones named in Section XVII. The corpus's convention is three tiers — established, moderate, inference — and this volume's contested joints are concentrated in two places: the coincidence of address, and the somatostatin-specific attribution of network and spatial phenotypes.

Established — the somatostatin peptide is reduced early and widely in the Alzheimer cortex. Directly measured, repeatedly, since 1980: reduced somatostatin-like immunoreactivity in Alzheimer cortex (Davies and colleagues, 1980); widespread cortical reduction, most severe temporally, with subcortical sparing (Beal and colleagues, 1986); somatostatin immunoreactivity present in neuritic plaques (Morrison and colleagues, 1985); cerebrospinal-fluid reductions at preclinical and prodromal stages (Sandoval and Witt, 2024). The oldest fact in this dissertation is also its most secure.

Established — the peptide is an excitability brake acting on the M-current. Directly demonstrated in hippocampal CA1 pyramidal neurons, with muscarinic agonists antagonising the same current (Moore and colleagues, 1988). The physiological identity of the peptide as a damping agent is not in question.

Established — somatostatin licenses neprilysin-mediated Aβ42 degradation, through a defined receptor pathway. Multiple independent demonstrations: somatostatin up-regulates neprilysin, and its genetic deficiency raises Aβ42 (Saito and colleagues, 2005); the pathway runs through α-endosulfine and KATP channels, and pharmacological intervention there reduces deposition and rescues memory (Watamura and colleagues, 2022); SST1 and SST4 redundantly regulate neprilysin, their double deletion worsens amyloid pathology on an APP background, and an SST1/4 agonist ameliorates it (Nilsson and colleagues, 2026). Loss- and gain-of-function both.

Established — somatostatin interneurons are lost early in human Alzheimer's disease, before excitatory and other inhibitory classes. Directly measured in the largest available multimodal human atlas, with replication across other major studies, and with the loss placed in a cognitively unimpaired phase (Gabitto and colleagues, 2024); corroborated at smaller scale by immunohistochemical cell counts showing somatostatin but not parvalbumin loss in temporal cortex (Waller and colleagues, 2020).

Established — the somatostatin cell is a dendrite-targeting feedback inhibitor whose axon occupies stratum lacunosum-moleculare and gates the entorhinal input to CA1. Directly demonstrated (Lovett-Barron and colleagues, 2012; Leão and colleagues, 2012), with the entorhinal counterpart's layer III–V targeting likewise directly demonstrated (Kecskés and colleagues, 2020).

Established — silencing somatostatin interneurons does not degrade grid cells; silencing parvalbumin interneurons does. Directly demonstrated by pharmacogenetic silencing during recording in freely moving mice (Miao and colleagues, 2017). This dissertation's deflationary correction rests on a clean positive-and-negative experiment.

Established — microglia preferentially eliminate inhibitory synapses, are GABA-receptive, and must degrade perineuronal nets to reach netted cells. Three separate direct demonstrations (Dejanovic and colleagues, 2022; Favuzzi and colleagues, 2021; Crapser and colleagues, 2020), on a foundation of complement-dependent early synapse loss (Hong and colleagues, 2016).

Established — p38α transduces both the microglial-RAGE and the amyloid insult in the entorhinal cortex, and its isoform-selective inhibition rescues early entorhinal dysfunction. Directly demonstrated (Origlia and colleagues, 2010; Rutigliano and colleagues,

2018; Roy and colleagues, 2015; Birnbaum and colleagues, 2015).

Established, with a correction the dissertation makes against its own outline — the amyloid-to-p38 route is calcium-flux-independent. APV prevented both synaptic loss and p38 activation; BAPTA and open-channel blockers did not (Birnbaum and colleagues, 2015). The intuitive chain "calcium overload activates p38" is *not* supported for this arm, and this volume does not assert it. Calcium load and p38 activation are treated as two convergent arms of the same insult.

Moderate — the somatostatin cell's calcium exposure is unusually high and its buffering unusually low. The dendritic supralinearity, its NMDA-receptor dependence, and the contributions of L-type channels, store release, and calcium-permeable AMPA receptors are directly measured in dendrite-targeting interneurons (Griesius and colleagues, 2025); the plaque-proximal hyperactivity is measured (Algamal and colleagues, 2022); the cross-disease vulnerability of this cell class to excitatory load is measured (Sloviter, 1987). What is *not* established, and is graded down accordingly, is the buffering asymmetry: that parvalbumin cells carry a high-capacity namesake buffer while somatostatin cells largely do not is true as stated, but no study cited here measures the two classes' calcium-handling capacity against Alzheimer-relevant load. The asymmetry is suggestive and directionally consistent; it is not a demonstration.

Moderate — the somatostatin interneuron is predominantly unnetted while its parvalbumin sibling is predominantly netted, and this explains the order of inhibitory loss. The net's clientele is established (Härtig and colleagues, 1992), as is the net's association with resistance to tau and to Alzheimer's disease (Morawski and colleagues, 2010; de Vries and colleagues, 2024) and the requirement for microglial net degradation (Crapser and colleagues, 2020). The *explanation* of the loss order is an inference assembled from measured parts, not a measured finding; no experiment has compared microglial engulfment at somatostatin- versus parvalbumin-cell synapses while manipulating the net. Additionally, a Kv3-expressing somatostatin subset does carry nets, so the claim is statistical rather than absolute.

Moderate — the somatostatin lesion contributes to Alzheimer network hyperexcitability. The network phenotype is established (Palop and colleagues, 2007), the peptide's damping action is established (Moore and colleagues, 1988), and the atlas itself proposes the excitatory-to-inhibitory imbalance reading (Gabitto and colleagues, 2024). But the best-documented molecular lesion behind the network phenotype is parvalbumin-specific (Verret and colleagues, 2012), and no equivalent somatostatin-specific rescue experiment exists. The contribution is plausible and mechanistically routed; its magnitude is unmeasured.

Synthesis, with one unmeasured joint — the coincidence of address. That stratum lacunosum-moleculare is the entorhinal terminal field, the OLM axonal territory, and the site of somatostatin-receptor-dependent neprilysin capacity rests on measured facts on both sides (Leão

and colleagues, 2012; Lovett-Barron and colleagues, 2012; Nilsson and colleagues, 2026). The joining is this dissertation's, and it carries one genuine gap: no experiment establishes that the somatostatin released by the OLM cell is the ligand maintaining that layer's neprilysin, nor that the somatostatin source and the neprilysin-bearing terminal are the same synapses. The alternative — entorhinal terminals bearing the neprilysin, maintained by a different somatostatin source — is not excluded. Named as the volume's most important untested claim.

Synthesis — the somatostatin loop has positive gain. Each arm is established; the closure is assembled here. The activity-dependent second closure through Aβ release (Cirrito and colleagues, 2005) is likewise an assembly of established parts. No experiment has measured loop gain, and this dissertation makes no quantitative claim about it.

Inference — the somatostatin cell's early loss contributes causally to the late excitatory-neuron loss. Supported by the human ordering (Gabitto and colleagues, 2024), by a cross-disease demonstration that deranged somatostatin interneurons drive pyramidal degeneration and that their focal ablation rescues it (Zhang and colleagues, 2016), and by the atlas's own suggestion. It is not demonstrated in Alzheimer's disease, and the transfer from a TDP-43 model to an amyloid disease is a real extrapolation.

Inference, and the volume's weakest station — the prefrontal arm. Frontal somatostatin loss is measured (Beal and colleagues, 1986; Sandoval and Witt, 2024), but no cell-type-resolved human study cited here establishes the prefrontal somatostatin interneuron's fate as a function of stage, and the frontal-worst versus temporal-worst gradients across the literature are unreconciled. The trajectory's third station is asserted on the corpus's temporal architecture and on peptide measurements, not on cell-resolved regional data.

Unestablished in brain — the somatostatin-receptor-to-p38 link. That SST2 and SST4 promote prolonged p38 activity is reported in non-neural, largely oncological systems. No demonstration in the Alzheimer cortex, in any cell type, was found. This volume uses p38α as the transducer of insults reaching the somatostatin cell's circuit, and explicitly does not claim a somatostatin-receptor-driven p38 pathway in brain.

A tension the ledger records rather than resolves. Silencing somatostatin cells degrades aperiodic entorhinal spatial coding (Miao and colleagues, 2017), yet in the J20 amyloid model non-grid entorhinal spatial coding was preserved while grid coding failed (Ying and colleagues, 2022). If somatostatin cells are lost early in amyloid pathology, the aperiodic code should have suffered. Possible reconciliations — species and model differences, partial rather than complete somatostatin-cell dysfunction at that stage, compensation, or the possibility that J20-stage somatostatin loss is milder than the human early phase — are all available and none is tested. This is a genuine anomaly for the account, and Section XVII proposes the experiment.

XVII. Predictions and Falsification

Six commitments, each stated so that a definite outcome would overturn it.

On the coincidence of address — the load-bearing test. If the somatostatin neuron's axonal territory is the amyloid disposal site, then conditional deletion of somatostatin from OLM interneurons alone, sparing all other somatostatin sources, should reduce neprilysin activity specifically in stratum lacunosum-moleculare and raise local Aβ42 there, without equivalent effect in stratum radiatum or oriens. Should OLM-specific somatostatin deletion leave lacunosum-moleculare neprilysin intact, the coincidence is a coincidence, and the dissertation's signature claim fails.

On the net, and the order of inhibitory loss. If the somatostatin cell is lost before the parvalbumin cell because it lacks the net that imposes a demolition step, then two things should hold. Microglial engulfment of somatostatin-cell synapses should exceed that of parvalbumin-cell synapses in the same tissue at the early stage; and the Kv3-expressing, *net-bearing* somatostatin subset should be relatively spared at the stage when the unnetted majority is lost. Finding equal engulfment of both classes, or finding the netted somatostatin subset lost at the same rate as the unnetted, would sever the net explanation from the loss order.

On the loop. If the somatostatin loop has positive gain, then supplying an SST1/SST4 agonist during the preclinical window should reduce Aβ accumulation *and* preserve somatostatin-interneuron number, and the two effects should be correlated across animals. If agonist treatment reduced amyloid without preserving the cells, the clearance arm would stand and the loop would not.

On the grid correction. If the grid code fails around the somatostatin cell rather than at its hands, then cell-type-specific rescue of entorhinal somatostatin interneurons in an amyloid model should restore the aperiodic spatial code and entorhinal-to-CA1 gamma coupling *before and more completely than* it restores grid periodicity, and any grid recovery should be attributable to preserved layer II cell survival rather than to restored inhibition of layer II. Finding that somatostatin rescue restores hexagonal periodicity directly and acutely would refute the correction and reinstate the simple story.

On the aperiodic anomaly. If the tension the ledger records is real, then in a model with documented early somatostatin-interneuron loss, aperiodic entorhinal spatial coding should be measurably degraded at the stage when somatostatin cells are first lost. Finding aperiodic coding

intact at a stage of demonstrated somatostatin-cell loss would indicate that the acute silencing result does not transfer to chronic degeneration, and would weaken the spatial arm of this dissertation to the deep-layer hand-off alone.

On the transducer. If p38α carries the insult to this circuit, then isoform-selective p38α inhibition during the preclinical window should preserve somatostatin-interneuron function — peptide content, dendritic output, plaque-proximal activity — and not merely pyramidal-cell physiology. Finding p38α inhibition to rescue pyramidal synapses while leaving somatostatin-cell dysfunction untouched would place the kinase downstream of, rather than upstream of, the governor's failure, and would reorder the therapeutic argument of Section XVIII.

XVIII. Therapeutic Corollaries and the Contradiction That Must Be Faced

The mechanism yields two therapeutic instructions that point in opposite directions, and this dissertation states the contradiction rather than choosing a side by silence.

The first instruction: restore the peptide's signalling. Everything in Sections III and IV argues for somatostatin-receptor agonism. SST1 and SST4 redundantly regulate neprilysin; their double deletion worsens amyloid pathology; and an SST1/4 agonist ameliorates it (Nilsson and colleagues, 2026). SST4 agonists increase cortical neprilysin activity and reduce Aβ42 oligomer species in models, and SST2 and SST4 are the review literature's primary therapeutic candidates (Sandoval and Witt, 2024). The KATP arm supplies a second handle, since pharmacological intervention at specific KATP subtypes activates neprilysin, reduces Aβ deposition, and rescues memory (Watamura and colleagues, 2022). The advantage of the receptor route is the one Section I anticipated: a receptor agonist can be supplied to a circuit whose somatostatin cells are silent but alive, which extends the therapeutic window past the point at which the peptide's own source has failed. Its limit is equally clear — it cannot help once the receptors' host cells are gone.

The second instruction, and why it is not simply wrong. The circuit literature has read the same cell in the opposite direction. Somatostatin interneurons are hyperactive near plaques while excitatory neurons are hypoactive (Algamal and colleagues, 2022), and over-active somatostatin interneurons drive pyramidal degeneration in a TDP-43 model, with focal ablation restoring pyramidal function (Zhang and colleagues, 2016). If somatostatin-cell over-activity disinhibits the principal cell, then the therapeutic instruction is to *suppress* the somatostatin cell, not to restore it. That instruction has been tested. Acute chemogenetic inhibition of somatostatin

interneurons enhanced excitatory-neuron function in APP/PS1 mice; but chronic, brain-wide somatostatin-interneuron inhibition failed to restore behavioural deficits — no improvement in locomotor activity, working memory, or fear-memory consolidation (Algamal and colleagues, 2025). This dissertation notes that the negative chronic result is reported at conference-abstract level and grades it accordingly; but its direction is informative, and it is what one would predict from the account developed here.

How the contradiction resolves — by phase, and by what is actually being fixed. The two instructions are not equally weighted, and the resolution is the volume's central therapeutic claim. Suppressing the somatostatin cell treats a *symptom of a broken governor* — erratic output from a controller whose sensor has failed — and does nothing about the four losses of Section XI: it does not restore the peptide, does not restore the cholinergic drive, does not close the feedback loop, and, crucially, *lowers the amyloid-clearance licence further* by reducing peptide release. It should therefore work acutely on the excitability read-out and fail chronically on the disease, which is what the model data show. Restoring somatostatin *signalling* — as opposed to restoring somatostatin-cell *firing* — addresses the peptide brake and the clearance licence simultaneously and does not depend on the cell's activity being well-regulated. The governor must be repaired or bypassed, not merely turned down; and a receptor agonist is a bypass.

The transducer as a third handle. Isoform-selective p38α inhibition is the corpus's kind of intervention: it rescues *early* entorhinal dysfunction (Rutigliano and colleagues, 2018), attenuates progression with benefit mediated primarily by rescue of neuronal function rather than by effects on the primary pathologies (Roy and colleagues, 2015), and blocks the amyloid route to spine loss at its documented step (Birnbaum and colleagues, 2015). Its virtue for this dissertation's argument is that it acts on the dysfunctional phase — on the reversible derangement of Section XII — rather than on deposits.

The microglial handle, and its known caution. Blocking the complement-dependent elimination that preferentially removes inhibitory synapses (Dejanovic and colleagues, 2022; Hong and colleagues, 2016) would, on this account, protect the governor's output synapses selectively. The corpus's microglial volumes have already established the caution that applies: microglial restraint is a resilience mechanism as well as a lesion, and blanket suppression forfeits the containment function that walls the deposit. The specific target here is not microglial activation in general but the inhibitory-synapse-directed engulfment programme, which the GABA-receptor-dependent recognition step makes at least conceptually separable (Favuzzi and colleagues, 2021).

And the timing, as always. The somatostatin lesion is documented, in human tissue, in the phase before cognitive impairment (Gabitto and colleagues, 2024). Every corollary above is a preclinical corollary. An SST4 agonist given to a brain whose somatostatin interneurons are already gone is an agonist without a physiological source to supplement and, quite possibly, without enough

surviving receptors to act on. The corpus's standing instruction applies here as sharply as anywhere it has been stated: the governor must be repaired while it is still a governor.

XIX. Coda The Governor and the Wall

There is a symmetry in this cell's fate that is worth stating plainly at the end, because it is the reason the somatostatin neuron deserved a volume rather than a section.

The brain builds two great classes of inhibitory interneuron. One it names for a calcium buffer and wraps in a lattice of sulfated sugar; that cell keeps time, and it survives into the disease's late phase. The other it names for a peptide message and leaves bare; that cell sets gain, and it is among the first things the disease takes. The unnetted cell is the one whose peptide damps the excitability of the cortex, and whose peptide licenses the enzyme that disposes of amyloid, and whose axon lies in the exact lamina where the entorhinal cortex hands its signal to the hippocampus and where that disposal happens. The brain concentrated three jobs in one bare cell, and then the disease found it.

What follows is a loop rather than a lesion. Amyloid injures the governor; the injured governor releases less of the peptide that clears amyloid and less of the brake that keeps the pyramidal cells from making more of it; the pyramidal cells fire more and make more; and the amyloid injures the governor further. That loop runs for years in a brain that tests normally, in tissue that shows no tangles in the cells that are dying, in a phase the atlases can now see and the clinic still cannot. By the time the disease is visible, the loop has been compounding for a decade, and the cell that might have broken it is largely gone.

This dissertation has declined one connection it was invited to make. The grid cells fail — early, measurably, in mice and in young human carriers of the risk allele — but they do not fail at the somatostatin cell's hands. That code belongs to the parvalbumin cell and to layer II, and the somatostatin cell's axon is not there. What the somatostatin cell holds instead is the deep-layer output, the aperiodic code, the hand-off to CA1, and the slow chronic conditions under which layer II must go on living. Losing it does not scramble the map at a stroke. It removes the restraint under which the map is maintained, and the map degrades afterwards, from causes the governor was there to prevent. That is a less dramatic story than the one the question invited, and it has the advantage of being what the experiments say.

To call the somatostatin neuron a governor is not a figure of speech. It is driven by what it limits; its output opposes the quantity that drives it; and its peptide sets, at one remove, the disposal rate of the substance whose accumulation is the disease. When a governor fails, the engine does not

stop. It runs faster, and it runs faster for a long time before anything visibly breaks. That interval — silent, compounding, and now, for the first time, measurable in human tissue — is where this cell's part of Alzheimer's disease is decided, and where anything that is going to help it must arrive.

References

All references below were retrieved and verified against the PubMed database during preparation; digital object identifiers are provided for each. One entry (reference 34) is a published conference abstract and is identified as such in the text and graded accordingly in the validity ledger.

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

2. Beal MF, Mazurek MF, Svendsen CN, Bird ED, Martin JB. Widespread reduction of somatostatin-like immunoreactivity in the cerebral cortex in Alzheimer's disease. *Annals of Neurology*. 1986;20(4):489–495. DOI: 10.1002/ana.410200408

3. Morrison JH, Rogers J, Scherr S, Benoit R, Bloom FE. Somatostatin immunoreactivity in neuritic plaques of Alzheimer's patients. *Nature*. 1985;314(6006):90–92. DOI: 10.1038/314090a0

4. Moore SD, Madamba SG, Joëls M, Siggins GR. Somatostatin augments the M-current in hippocampal neurons. *Science*. 1988;239(4837):278–280. DOI: 10.1126/science.2892268

5. Saito T, Iwata N, Tsubuki S, Takaki Y, Takano J, Huang SM, Suemoto T, Higuchi M, Saido TC. Somatostatin regulates brain amyloid β peptide Aβ42 through modulation of proteolytic degradation. *Nature Medicine*. 2005;11(4):434–439. DOI: 10.1038/nm1206

6. Watamura N, Kakiya N, Nilsson P, Tsubuki S, Kamano N, Takahashi M, Hashimoto S, Sasaguri H, Saito T, Saido TC. Somatostatin-evoked Aβ catabolism in the brain: mechanistic involvement of α-endosulfine-KATP channel pathway. *Molecular Psychiatry*. 2022;27(3):1816–1828. DOI: 10.1038/s41380-021-01368-8

7. Nilsson P, Sörgjerd K, Kakiya N, Sasaguri H, Watamura N, Johansson L, Shimozawa M, Tsubuki S, Zhou Z, Loera-Valencia R, Takamura R, Sekiguchi M, Pegel A, Schulz S, Saito T, Iwata N, Winblad B, Saido TC. Somatostatin receptor subtypes 1 and 4 regulate neprilysin, the major amyloid-β degrading enzyme in brain. *Journal of Alzheimer's Disease*. 2026;109(2):651–666. DOI: 10.1177/13872877251392782

8. Wang H, Muiznieks LD, Ghosh P, Williams D, Solarski M, Fang A, Ruiz-Riquelme A, Pomès R, Watts JC, Chakrabartty A, Wille H, Sharpe S, Schmitt-Ulms G. Somatostatin binds to the human amyloid β peptide and favors the formation of distinct oligomers. *eLife*. 2017;6:e28401. DOI: 10.7554/eLife.28401

9. Williams D, Yan BQ, Wang H, Negm L, Sackmann C, Verkuyl C, Rezai-Stevens V, Eid S, Vediya N, Sato C, Watts JC, Wille H, Schmitt-Ulms G. Somatostatin slows Aβ plaque deposition in aged APP^NL-F/NL-F mice by blocking Aβ aggregation. *Scientific Reports*. 2023;13:2337. DOI: 10.1038/s41598-023-29559-z

10. Gabitto MI, Travaglini KJ, Rachleff VM, Kaplan ES, Long B, Ariza J, Ding Y, Mahoney JT, Dee N, Goldy J, et al. Integrated multimodal cell atlas of Alzheimer's disease. *Nature Neuroscience*. 2024;27(12):2366–2383. DOI: 10.1038/s41593-024-01774-5

11. Waller R, Mandeya M, Viney E, Simpson JE, Wharton SB. Histological characterization of interneurons in Alzheimer's disease reveals a loss of somatostatin interneurons in the temporal cortex. *Neuropathology*. 2020;40(4):336–346. DOI: 10.1111/neup.12649

12. Schmid LC, Mittag M, Poll S, Steffen J, Wagner J, Geis HR, Schwarz I, Schmidt B, Schwarz MK, Remy S, Fuhrmann M. Dysfunction of somatostatin-positive interneurons associated with memory deficits in an Alzheimer's disease model. *Neuron*. 2016;92(1):114–125. DOI: 10.1016/j.neuron.2016.08.034

13. Algamal M, Russ AN, Miller MR, Hou SS, Maci M, Munting LP, Zhao Q, Gerashchenko D, Bacskai BJ, Kastanenka KV. Reduced excitatory neuron activity and interneuron-type-specific deficits in a mouse model of Alzheimer's disease. *Communications Biology*. 2022;5(1):1323. DOI: 10.1038/s42003-022-04268-x

14. Zhang W, Zhang L, Liang B, Schroeder D, Zhang ZW, Cox GA, Li Y, Lin DT. Hyperactive somatostatin interneurons contribute to excitotoxicity in neurodegenerative disorders. *Nature Neuroscience*. 2016;19(4):557–559. DOI: 10.1038/nn.4257

15. Lovett-Barron M, Turi GF, Kaifosh P, Lee PH, Bolze F, Sun XH, Nicoud JF, Zemelman BV, Sternson SM, Losonczy A. Regulation of neuronal input transformations by tunable dendritic inhibition. *Nature Neuroscience*. 2012;15(3):423–430. DOI: 10.1038/nn.3024

16. Leão RN, Mikulovic S, Leão KE, Munguba H, Gezelius H, Enjin A, Patra K, Eriksson A, Loew LM, Tort AB, Kullander K. OLM interneurons differentially modulate CA3 and entorhinal inputs to hippocampal CA1 neurons. *Nature Neuroscience*. 2012;15(11):1524–1530. DOI: 10.1038/nn.3235

17. Kecskés M, Henn-Mike N, Agócs-Laboda Á, Szőcs S, Petykó Z, Varga C. Somatostatin expressing GABAergic interneurons in the medial entorhinal cortex preferentially inhibit layer III–V pyramidal cells. *Communications Biology*. 2020;3(1):754. DOI: 10.1038/s42003-020-01496-x

18. Miao C, Cao Q, Moser MB, Moser EI. Parvalbumin and somatostatin interneurons control different space-coding networks in the medial entorhinal cortex. *Cell*. 2017;171(3):507–521.e17. DOI: 10.1016/j.cell.2017.08.050

19. Fu H, Rodriguez GA, Herman M, Emrani S, Nahmani E, Barrett G, Figueroa HY, Goldberg E, Hussaini SA, Duff KE. Tau pathology induces excitatory neuron loss, grid cell dysfunction, and spatial memory deficits reminiscent of early Alzheimer's disease. *Neuron*. 2017;93(3):533–541.e5. DOI: 10.1016/j.neuron.2016.12.023

20. Ying J, Keinath AT, Lavoie R, Vigneault E, El Mestikawy S, Brandon MP. Disruption of the grid cell network in a mouse model of early Alzheimer's disease. *Nature Communications*. 2022;13(1):886. DOI: 10.1038/s41467-022-28551-x

21. Jun H, Bramian A, Soma S, Saito T, Saido TC, Igarashi KM. Disrupted place cell remapping and impaired grid cells in a knockin model of Alzheimer's disease. *Neuron*. 2020;107(6):1095–1112.e6. DOI: 10.1016/j.neuron.2020.06.023

22. Kunz L, Schröder TN, Lee H, Montag C, Lachmann B, Sariyska R, Reuter M, Stirnberg R, Stöcker T, Messing-Floeter PC, Fell J, Doeller CF, Axmacher N. Reduced grid-cell-like representations in adults at genetic risk for Alzheimer's disease. *Science*. 2015;350(6259):430–433. DOI: 10.1126/science.aac8128

23. 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

24. 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. DOI: 10.1016/j.cell.2012.02.046

25. 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

26. Sloviter RS. Decreased hippocampal inhibition and a selective loss of interneurons in experimental epilepsy. *Science*. 1987;235(4784):73–76. DOI: 10.1126/science.2879352

27. Griesius S, Richardson A, Kullmann DM. Supralinear dendritic integration in murine dendrite-targeting interneurons. *eLife*. 2025;13:RP100268. DOI: 10.7554/eLife.100268

28. Origlia N, Bonadonna C, Rosellini A, Leznik E, Arancio O, Yan SS, Domenici L. Microglial receptor for advanced glycation end product-dependent signal pathway drives β-amyloid-induced synaptic depression and long-term depression impairment in entorhinal cortex. *Journal of Neuroscience*. 2010;30(34):11414–11425. DOI: 10.1523/JNEUROSCI.2127-10.2010

29. Birnbaum JH, Bali J, Rajendran L, Nitsch RM, Tackenberg C. Calcium flux-independent NMDA receptor activity is required for Aβ oligomer-induced synaptic loss. *Cell Death & Disease*. 2015;6:e1791. DOI: 10.1038/cddis.2015.160

30. Roy SM, Grum-Tokars VL, Schavocky JP, Saeed F, Staniszewski A, Teich AF, Arancio O, Bachstetter AD, Webster SJ, Van Eldik LJ, Minasov G, Anderson WF, Pelletier JC, Watterson DM. Targeting human central nervous system protein kinases: an isoform-selective p38αMAPK inhibitor that attenuates disease progression in Alzheimer's disease mouse models. *ACS Chemical Neuroscience*. 2015;6(4):666–680. DOI: 10.1021/acschemneuro.5b00002

31. Rutigliano G, Stazi M, Arancio O, Watterson DM, Origlia N. An isoform-selective p38α mitogen-activated protein kinase inhibitor rescues early entorhinal cortex dysfunctions in a mouse model of Alzheimer's disease. *Neurobiology of Aging*. 2018;70:86–91. DOI: 10.1016/j.neurobiolaging.2018.06.006

32. Hong S, Beja-Glasser VF, Nfonoyim BM, Frouin A, Li S, Ramakrishnan S, Merry KM, Shi Q, Rosenthal A, Barres BA, Lemere CA, Selkoe DJ, Stevens B. Complement and microglia mediate early synapse loss in Alzheimer mouse models. *Science*. 2016;352(6286):712–716. DOI: 10.1126/science.aad8373

33. Dejanovic B, Wu T, Tsai MC, Solanoy H, Reichert P, Zhang J, Rose CM, Hruska KS, Foreman O, Hanson JE, Sheng M. Complement C1q-dependent excitatory and inhibitory synapse elimination by astrocytes and microglia in Alzheimer's disease mouse models. *Nature Aging*. 2022;2(9):837–850. DOI: 10.1038/s43587-022-00281-1

34. Algamal M, Abdallah S, Ndambakuwa W, Kastanenka KV, Bacskai BJ. Somatostatin interneuron inhibition as a strategy to restore excitation-inhibition balance in Alzheimer's disease. *Alzheimer's & Dementia*. 2025;21(Suppl 1):e106574. DOI: 10.1002/alz.70855 *(published conference abstract)*

35. Favuzzi E, Ibrahim LA, Saldi GA, Cao Y, Fernández-Otero M, Zeine A, Sefah A, Xu Q, Dimidschstein J, Bonneau R, Datta SR, Stevens B, Fishell G. GABA-receptive microglia selectively sculpt developing inhibitory circuits. *Cell*. 2021;184(15):4048–4063.e32. DOI: 10.1016/j.cell.2021.06.018

36. Härtig W, Brauer K, Brückner G. Wisteria floribunda agglutinin-labelled nets surround parvalbumin-containing neurons. *Neuroreport*. 1992;3(10):869–872. DOI: 10.1097/00001756-199210000-00012

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

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

39. Crapser JD, Spangenberg EE, Barahona RA, Arreola MA, Hohsfield LA, Green KN. Microglia facilitate loss of perineuronal nets in the Alzheimer's disease brain. *EBioMedicine*. 2020;58:102919. DOI: 10.1016/j.ebiom.2020.102919

40. Sandoval KE, Witt KA. Somatostatin: linking cognition and Alzheimer disease to therapeutic targeting. *Pharmacological Reviews*. 2024;76(6):1291–1325. DOI: 10.1124/pharmrev.124.001117

41. Pesold C, Impagnatiello F, Pisu MG, Uzunov DP, Costa E, Guidotti A, Caruncho HJ. Reelin is preferentially expressed in neurons synthesizing gamma-aminobutyric acid in cortex and hippocampus of adult rats. *Proceedings of the National Academy of Sciences USA*. 1998;95(6):3221–3226. DOI: 10.1073/pnas.95.6.3221

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