THE PROTEOLYTIC TURN

Convergence of Iron Liberation,
Lipid Gridlock, and Complement Priming
on the Aggrecan–Brevican Sheath
of the Parvalbumin Interneuron
as the Phase II → Phase III Transition Mechanism

Crapser · Marschallinger · Stockwell · Conrad
Ayton · Bush · Snow · Bishop · Stevens
Shatz · Lemere · Tsai · Fawcett · de Vries

in dialogue with the Locus Coeruleus Bridge and the Collapse Trilogy theses

Prepared under the ONS Methodology

AdultCognitiveDisease.com

Dr. James Truchard & Benjamin Aaron Gustafsson

22 May 2026

Companion to The Locus Coeruleus Bridge; in dialogue with the Collapse Trilogy

Contents

Abstract

The three-phase framework of late-onset Alzheimer's disease pathogenesis that has emerged from the Oskar Fischer Prize corpus — Bioenergetic Ignition in the locus coeruleus and dorsal raphe during the third through fifth decades, Microglial Bridgehead in the hippocampus, oligodendrocyte, and microglial network during the sixth and seventh decades, and Structural Disintegration of the perineuronal-net sheaths around parvalbumin-positive interneurons in the eighth decade and beyond — partitions the natural history of the disease into temporally and anatomically distinct windows whose distinct therapeutic logics underwrite the *temporal pharmacology* claim of the parent framework. The companion paper on the locus-coeruleus bridge has resolved the mechanistically unspecified Phase I Phase II boundary by identifying the coupled withdrawal of tonic noradrenergic suppression and trans-synaptic templated propagation of pretangle tau along the ascending LC efferents as the dual mechanism by which a brainstem lesion executes a forebrain disease. The Phase II Phase III boundary has been described in the parent framework as the better-specified of the two boundaries on the grounds that the matrix metalloproteinase output of post-homeostatic microglia is the proximate executor of perineuronal-net digestion and therefore the boundary is connected by a single effector mechanism. This paper argues that the apparent specification is in fact only a gesture, that the mechanism by which Phase II homeostatic collapse executes the Phase III matrix disintegration is no more rigorously specified than the Phase I Phase II boundary was before the LC bridge synthesis, and that a proper account requires the identification of the molecular logic by which the terminal state of Phase II microglia produces the specific proteolytic, redox, and complement signatures that converge on the aggrecan–brevican sheath of the parvalbumin-positive interneuron. We propose that the Phase II Phase III transition is tripartite and that its three arms converge on a single cellular substrate. The first arm is the proteolytic switch of post-homeostatic microglia from a phagocytic-clearance phenotype into a matrix-degrading effector phenotype, executed through lipid-droplet accumulation (LDAM) and the consequent metabolic reconfiguration that licenses MMP-9 and ADAMTS-4/5 transcription through the NLRP3-IL-1β axis. The second arm is the liberation of redox-active free iron from ferroptotic oligodendrocytes whose membranes have been catalytically dismantled by the lipid peroxidation chemistry that Stockwell, Conrad, and Hambright have established as the executioner of Phase II oligodendroglial death, with the liberated iron loading onto the chondroitin-sulfate proteoglycan backbone of nearby perineuronal nets and converting the net itself into a Fenton-active substrate that catalyzes its own degradation through the action of the released microglial proteases. The third arm is the complement priming of the synaptic interface by post-homeostatic microglial C1q and astrocytic C3, with C4d deposition on aggrecan and brevican opsonizing the perineuronal sheath for CR3-mediated stripping by the same proteolytic microglial population that is producing the matrix metalloproteinase output. The three arms converge on a single cellular substrate — the parvalbumin-positive interneuron whose perineuronal net is the only structure in the parenchyma that simultaneously serves as an

iron-chelation barrier, a diffusion barrier protecting GABAergic synapses from extracellular insults, and a structural scaffold for the OTX2 and Sema3A signaling that maintains the cell's high-firing identity — and the convergence is consequential because the prior Phase I bioenergetic collapse has already narrowed the substrate margin of these neurons to the point at which the simultaneous loss of iron buffering, lipid-peroxidation protection, and synaptic envelope integrity collapses their gamma-frequency drive within a clinically rapid window. The PV+ interneuron is the Phase III bridgehead in the same sense that the hippocampus is the Phase II bridgehead: it is the cellular environment in which the converging upstream pressures consolidate into a self-sustaining pathology that no longer requires further input from earlier phases to progress. The integrated transition model has direct therapeutic implications for the Phase II III window: an MMP-9 inhibitor strategy alone does not address the iron-mediated catalytic substrate of perineuronal-net dissolution, and an iron-chelation strategy alone does not address the proteolytic enzymes whose secretion proceeds from a metabolic state that iron chelation alone cannot reverse; the appropriate intervention is therefore combinatorial, with broad MMP inhibition (minocycline, doxycycline) or selective MMP-9 inhibition (JNJ0966) deployed alongside iron chelation (deferiprone, deferoxamine) and complement modulation (C1q/C3 blockade) during a narrow window in which the homeostatic signature has been lost but the perineuronal-net architecture has not yet been catalytically dismantled. The paper closes with the implications of the integrated transition model for the structure of the three-phase framework as a whole and with the explicit acknowledgment of the questions that remain unresolved.

1. Introduction: The Better-Specified Boundary Is Still Underspecified

The locus-coeruleus bridge paper that accompanies this synthesis observed in passing that the Phase II Phase III boundary of the three-phase framework was the better-specified of the two boundaries because the matrix metalloproteinase program of post-homeostatic microglia is the direct executor of perineuronal-net degradation, so that the Phase II microglial collapse and the Phase III matrix disintegration are connected by a single effector mechanism. The observation was true at the level of generality at which it was made. It is, however, insufficient at the level of specificity that a mechanistic transition account requires. The fact that MMP-9 is one of the secreted effectors of disease-associated microglia and that MMP-9 digests aggrecan and brevican does not yet constitute a specification of the transition mechanism. It is a statement about the existence of an effector connection; it is not a statement about the molecular conditions under which the effector connection becomes active, the redox and metabolic substrate on which the connection executes its function, the cellular target on which the convergence consolidates, or the temporal window within which the transition is open to interdiction. The Phase II Phase III boundary, on

closer examination, has been undertheorized in precisely the same manner that the Phase I Phase II boundary was undertheorized before the LC bridge specification was developed.

The lacuna matters because the framework's claim is not merely that three distinct pathologies follow each other in sequence but that the transition between them is itself a mechanistic event whose interdiction defines a therapeutic window distinct from interventions on either of the bracketing phases. A Phase II ferroptosis-inhibitor intervention may slow the rate at which oligodendrocyte cell death proceeds, but it does not by itself prevent the conversion of post-homeostatic microglia from a phagocytic-clearance phenotype into a matrix-degrading effector phenotype, because that conversion has its own metabolic and transcriptional logic that is downstream of, but not identical with, the ferroptotic substrate. A Phase III MMP-9 inhibitor intervention preserves the residual perineuronal-net architecture once the proteolytic program has begun, but it does not address the redox catalysis by which iron released from upstream ferroptotic events accelerates matrix dissolution independent of enzymatic action, and it does not address the complement opsonization of the synaptic interface that drives CR3-mediated stripping in parallel to MMP-9 cleavage. The transition itself is the locus of an intervention that neither bracketing-phase intervention captures, and the failure to specify the transition mechanism is therefore a failure to specify the most disease-modifying of the available windows.

This paper develops the argument in nine stages. We first review the terminal state of Phase II — the hippocampal bridgehead at the moment of its consolidation — and identify the four cellular populations whose configuration at that moment specifies the substrate from which the Phase III transition will proceed. We then develop the first arm of the transition mechanism, the proteolytic switch of post-homeostatic microglia from phagocytic clearance to matrix degradation through lipid-droplet accumulation and the consequent metabolic reconfiguration of NLRP3-IL-1β-MMP-9 transcription, drawing on the Marschallinger LDAM program, the Heneka inflammasome program, and the broader literature on the metabolic reprogramming of microglia in late-stage Alzheimer pathology. We then develop the second arm, the catalytic licensing of perineuronal-net degradation by iron liberated from ferroptotic oligodendrocytes, drawing on the Stockwell, Conrad, and Hambright programs on ferroptosis biochemistry, on the Ayton and Bush programs on iron in Alzheimer's disease, and on the Snow and Bishop programs on glycosaminoglycan iron binding. We then develop the third arm, the complement priming of the synaptic interface, drawing on the Stevens and Shatz programs on developmental and pathological complement-mediated synaptic pruning and on the Crapser program on complement involvement in perineuronal-net stripping. We then identify the structural convergence of the three arms — that they converge on a single cellular substrate, the parvalbumin-positive interneuron whose perineuronal net is the only structure in the parenchyma that performs the three functions the three arms simultaneously attack — and argue that the PV+ interneuron is the Phase III bridgehead in the same anatomical-functional sense in which the hippocampus is the Phase II bridgehead. We then address why the prior bioenergetic collapse of Phase I has already narrowed the substrate margin of these cells to the point at which the

convergence is acute, and we close with the therapeutic implications and with the questions that remain unresolved.

2. The Hippocampal Bridgehead at Its Terminal State

The terminal state of Phase II as the framework has been developed is not a quiescent endpoint but a metastable configuration of four cellular populations whose continuing interaction is what will produce the Phase III transition. The Phase II bridgehead, in the formulation of the parent thesis, is characterized by oligodendrocyte ferroptosis, by the homeostatic-to-disease-associated transition of microglia, by lysosomal acidification failure in neurons of the PANTHOS type, and by the consolidation of the hippocampal substrate as the locus of forebrain disease. By the end of Phase II — that is, at approximately the seventh decade of life in the framework's chronology — these four processes have not stopped occurring; they have reached a configuration in which their continuing operation produces a specific output, and that output is the substrate from which the Phase III transition will proceed.

The first relevant population is the oligodendrocyte. The Phase II oligodendrocyte ferroptosis described in the parent framework, drawing on the Stockwell and Conrad programs and on the more recent application of those programs to the Alzheimer's hippocampus by Ayton, Bush, Lei, and colleagues, is not a synchronous event in which all hippocampal oligodendrocytes die at the same moment. It is a stochastic, asynchronous process in which individual oligodendrocytes whose GPX4 buffer has been most thoroughly depleted by the cumulative lipid-peroxidation flux of Phase II metabolism cross the ferroptotic threshold first and execute their characteristic lipid-radical chain reaction, releasing into the surrounding parenchyma the contents of their cytoplasm and the iron-rich myelin debris of their disassembled processes. The proximate effect of each ferroptotic event is the release of a substantial bolus of redox-active iron — both the iron held within mitochondrial and lysosomal compartments and the iron stably bound to ferritin until its degradation in the dying cell — and the release of a corresponding quantity of lipid peroxidation products including 4-hydroxynonenal and malondialdehyde, both of which themselves serve as further oxidative signals to adjacent cells. The result is that the hippocampal parenchyma in late Phase II is characterized by a chronic, regionally elevated load of free iron and lipid-peroxidation byproducts whose magnitude is determined by the cumulative rate of oligodendrocyte ferroptosis up to that moment and whose distribution is biased toward the regions of densest myelinated axonal projection.

The second relevant population is the microglial network. The Phase II microglial state is not a uniform population but a continuous distribution that the disease-associated-microglia literature has divided, somewhat artificially, into discrete states designated DAM stage 1, DAM stage 2, lipid-droplet-accumulating microglia (LDAM), and a less well-defined dystrophic population. The Homeostatic Microglial Collapse thesis has argued that this apparent multiplicity is best understood as a continuous trajectory whose branches reflect the differential availability of phagocytic substrate, mitochondrial quality-control competence, and the local oxidative environment. By the end of Phase II, the dominant microglial population in the hippocampus has lost the homeostatic signature characterized by Butovsky and colleagues, has accumulated lipid droplets through repeated cycles of myelin-debris phagocytosis, and has undergone the metabolic reconfiguration that the Ulland and Colonna programs on TREM2-mediated metabolic licensing have shown to be a precondition for sustained effector output. The state is unstable in two senses. It is metabolically unstable, because the lipid load that has accumulated through phagocytic activity is gradually exceeding the cellular machinery available to process it, and the cell is approaching the lipid-gridlock condition that Marschallinger has identified as the LDAM phenotype. It is also functionally unstable, because the secretory profile of a microglial cell at this stage is no longer the cytokine-only profile of an early DAM state but is beginning to include the matrix metalloproteinases and other proteolytic enzymes whose transcription becomes licensed once the inflammasome program has reached a threshold of activation.

The third relevant population is the neuronal population whose autophagy-lysosomal pipeline has begun to fail in the PANTHOS pattern that Nixon and colleagues have characterized. By the end of Phase II, a substantial fraction of hippocampal pyramidal neurons in the CA1 and dentate gyrus regions are accumulating intraneuronal amyloid-beta within autolysosomes that have lost the capacity to acidify sufficiently to degrade their cargo, and the surrounding neuropil is beginning to receive the first wave of plaque deposition produced by the death and extrusion of the earliest-failing neurons. The relevance of this population to the Phase II Phase III transition is not that it is itself the source of the proteolytic effectors that will execute Phase III, but that it is the source of the antigenic and chemotactic signals that recruit and sustain microglial engagement with the parenchyma at high density, and that it is the source of the synaptic loss and complement-tagging that establishes the substrate on which the Phase III complement priming arm will operate. The PANTHOS neurons are, in functional terms, the inflammation-sustaining substrate that keeps the post-homeostatic microglial population continuously engaged at the level of activity at which the proteolytic switch becomes thermodynamically permissible.

The fourth relevant population is the parvalbumin-positive interneuron itself, which is at the end of Phase II still functionally intact but whose substrate margin has been narrowed by the convergent prior pressures of Phase I and Phase II. The PV+ interneuron is among the highest-firing neurons in the cortex and hippocampus, with sustained firing rates that exceed those of pyramidal cells by approximately an order of magnitude, and its high firing rate is supported by

an unusual cellular architecture: a high mitochondrial density, an unusually high expression of the cytochrome c oxidase subunits required for sustained OXPHOS, an elevated baseline glycolytic capacity, and a perineuronal-net sheath whose function is to provide simultaneously a diffusion barrier against extracellular insults, a chelation barrier against free iron at the cell surface, and a structural scaffold for the OTX2 and Sema3A signaling that maintains the cell's high-firing identity. The Phase I bioenergetic collapse has already taxed the mitochondrial substrate of these cells, and the Phase II oxidative environment has begun to load their PNN with iron at rates exceeding the chelation capacity of the aggrecan and brevican backbone. By the end of Phase II, the PV+ interneuron is in a state of latent vulnerability: its function is preserved, but the margin between its substrate availability and the demand placed on it by its sustained firing has narrowed substantially, and a relatively small additional perturbation will tip its capacity below threshold.

These four populations together specify the substrate from which the Phase III transition will proceed. The transition is not a single event with a single mechanism but the convergent consequence of the continuing operation of these four populations on each other, executed through the three arms developed in the sections that follow.

3. The First Arm: Lipid Gridlock and the Proteolytic Switch

The first transmission mechanism of the Phase II Phase III transition is the conversion of the post-homeostatic microglial population from a phagocytic-clearance phenotype, in which its dominant output is the internalization of debris and the secretion of cytokines, into a matrix-degrading effector phenotype, in which its dominant output includes the matrix metalloproteinases and the ADAMTS-family proteases whose specific cleavage products are the aggrecan and brevican fragments characteristic of perineuronal-net dissolution. The conversion is not a transcriptional novelty introduced in late Phase II; the genes for MMP-9, MMP-2, ADAMTS-4, and ADAMTS-5 are accessible to the microglial transcriptional program throughout the DAM trajectory, and their expression at low levels is detectable in DAM stage 1 cells. What changes in the late-Phase-II window is the rate of their transcription and the post-translational handling of their products, and the changes are licensed by a specific metabolic and inflammasome configuration that the parent framework has called the LDAM-to-effector switch.

The Marschallinger program has established that lipid-droplet-accumulating microglia constitute a distinct phenotypic state characterized by the accumulation of large neutral-lipid droplets in the cytoplasm, by a transcriptional signature that overlaps with but is not identical to the DAM signature, and by a functional profile that includes elevated ROS production, impaired

phagocytic capacity, and increased pro-inflammatory cytokine output. The state was originally characterized in aging mouse brain and has subsequently been identified in human Alzheimer's tissue, and its functional significance is the demonstration that microglial accumulation of lipid material does not merely impair their phagocytic capacity but actively reprograms their secretory output toward inflammatory and tissue-destructive products. The mechanism of this reprogramming has been clarified by subsequent work showing that the lipid droplets are themselves a signaling substrate, with the perilipin-coated droplet surface serving as a platform for NLRP3 inflammasome assembly and with the lipid metabolic flux through beta-oxidation and through cholesterol esterification serving as the substrate for transcription factor regulation of inflammatory programs.

The connection of the LDAM state to the proteolytic switch operates through the NLRP3-IL-1β axis. The Heneka program has established that NLRP3 inflammasome activation in microglia is gated by mitochondrial reactive oxygen species output and by lysosomal destabilization, both of which are elevated in the LDAM state because the lipid load has compromised both mitochondrial fatty-acid oxidation and lysosomal membrane integrity. NLRP3 activation produces the caspase-1-mediated cleavage of pro-IL-1β to mature IL-1β, and IL-1β acts in autocrine and paracrine fashion on microglia and adjacent cells to upregulate the transcription of MMP-9, MMP-3, and ADAMTS-family proteases through a combination of NF-κB and AP-1 transcription factor activation. The matrix metalloproteinase program is therefore not a novel transcriptional state but the predictable downstream consequence of sustained NLRP3 activation, and the sustained NLRP3 activation is the predictable consequence of LDAM lipid loading.

The functional consequence is that the dominant secretory output of late-Phase-II microglia in the immediate vicinity of perineuronal nets transitions from cytokines to proteases over a window of months to years that corresponds, in the framework's chronology, to the late seventh decade of life. The temporal coincidence is not accidental: it reflects the cumulative time required for the lipid load to reach the threshold at which the NLRP3 program enters sustained activation. The proteolytic switch is also irreversible in the relevant sense, in that lipid droplets in microglia do not spontaneously regress under the metabolic conditions of the aged forebrain, and the inflammasome activation does not spontaneously resolve in the absence of an intervention that addresses the upstream lipid burden. Once the switch has occurred, the microglial population in the vicinity of perineuronal nets is engaged in proteolytic effector output continuously, and the matrix metalloproteinase load on the local extracellular environment is sustained at a level substantially above the homeostatic baseline.

The Crapser program has documented the proximate consequence of this proteolytic load on the perineuronal-net architecture itself. Joshua Crapser and colleagues, working in the Green laboratory, demonstrated in 2020 that microglial depletion through CSF1R inhibition preserves perineuronal-net integrity in 5xFAD mice, that the perineuronal-net loss observed in untreated

5xFAD animals is mediated by microglial proteolytic output rather than by neuronal-intrinsic processes, and that the proteolytic output responsible for PNN loss includes MMP-9 alongside the ADAMTS family enzymes. The Crapser paper is the cleanest available experimental demonstration that the connection between Phase II microglial collapse and Phase III matrix disintegration runs through microglial-derived proteases, and the result has been confirmed and extended by subsequent work showing that the proteolytic output is regionally heterogeneous in a manner that matches the regional heterogeneity of PNN loss in human Alzheimer's tissue.

The first arm therefore specifies the proximate mechanism by which Phase II microglial post-homeostatic transition produces the Phase III matrix-degrading effector output. The specification, however, is incomplete in two respects. First, it does not yet account for the catalytic acceleration of matrix dissolution that proceeds independently of, but in parallel with, the enzymatic action of the secreted proteases — an acceleration mediated by free iron and discussed in the next section. Second, it does not yet account for the targeting of the proteolytic output to the perineuronal-net substrate specifically rather than to other extracellular substrates available in the parenchyma — a targeting that is mediated by the complement priming discussed in the section after that.

4. The Second Arm: Iron Liberation from Ferroptotic Oligodendrocytes Catalyzes Matrix Degradation

The second transmission mechanism operates in parallel with the first arm but through a non-enzymatic chemistry that the existing framework has not yet specified at the level of detail that its therapeutic implications require. The mechanism is the liberation of redox-active free iron from oligodendrocytes that have crossed the ferroptotic threshold in Phase II, the diffusion of this liberated iron to the perineuronal-net architecture surrounding nearby parvalbumin-positive interneurons, the binding of the iron to the chondroitin-sulfate side chains of aggrecan and brevican, and the consequent conversion of the perineuronal net itself into a Fenton-active catalytic substrate that drives its own degradation through hydroxyl-radical generation in the immediate vicinity of its core proteoglycan scaffold.

The biochemistry of the mechanism proceeds in three steps. The first step is the iron-release event itself. Oligodendrocytes are the highest-iron cell population in the central nervous system, with iron content per cell exceeding that of neurons by approximately tenfold and that of microglia by approximately threefold. The iron is held in three main pools: ferritin-bound iron within the cytoplasm, iron incorporated into the heme prosthetic groups of mitochondrial cytochromes and

other heme proteins, and iron incorporated into the iron-sulfur clusters of metabolic enzymes and of the myelin synthesis machinery. When an oligodendrocyte undergoes ferroptotic death, the lipid peroxidation chain reaction characteristic of ferroptosis dismantles the membranes of the cell, including the lysosomal and mitochondrial membranes that ordinarily segregate the iron pools from the cytoplasm and from the extracellular space. The result is the release of a substantial fraction of the cell's iron content into the local extracellular environment, both as free iron and as iron bound to released proteins and to small chelators. The Stockwell and Conrad programs have established that the released iron is overwhelmingly in the ferrous (Fe^{2}) oxidation state at the moment of release, both because ferroptosis depends on Fe^{2} as the catalyst of the chain reaction and because the cellular environment within which the chain reaction proceeds is reductively biased. Ferrous iron is the form that is redox-active in Fenton chemistry, and its release into the extracellular space therefore constitutes a release of catalytic potential, not merely a release of nutritional iron.

The second step is the binding of the liberated iron to the perineuronal-net glycosaminoglycan backbone. The Snow program, developed over more than three decades at ProteoTech and subsequently at the University of Washington, has established that chondroitin-sulfate and heparan-sulfate glycosaminoglycans bind metal ions with substantial affinity, that the binding is mediated principally by the sulfate groups on the side chains and secondarily by the carboxylate groups on the uronic acid residues, and that the binding of iron specifically converts the proteoglycan from an inert structural component into a redox-catalyst substrate. The Bishop program has extended this work to the cognate heparan-sulfate proteoglycans and has demonstrated that the iron-binding affinity of glycosaminoglycans depends on their sulfation pattern, with the 4-sulfate and 6-sulfate forms of chondroitin sulfate binding iron more avidly than the unsulfated form. The relevance to the perineuronal net is direct: the chondroitin-sulfate proteoglycan backbone of the PNN, dominated by aggrecan and brevican with secondary contributions from versican and neurocan, presents an extensive negatively charged surface to the extracellular space that is configured to bind cations, and the iron released from nearby ferroptotic oligodendrocytes binds to this surface preferentially because the local concentration of carboxylate and sulfate binding sites exceeds that of any other extracellular substrate in the immediate parenchymal environment.

The third step is the Fenton catalysis itself. Ferrous iron bound to a substrate that is in the immediate vicinity of hydrogen peroxide — and the hippocampal parenchyma of late-Phase-II Alzheimer's brain is rich in hydrogen peroxide produced by mitochondrial dysfunction in adjacent cells and by NADPH-oxidase activity in adjacent post-homeostatic microglia — catalyzes the production of hydroxyl radicals through the Fenton reaction. The hydroxyl radical is the most reactive of the biological oxidants and has a diffusion radius of approximately one nanometer before it reacts with the nearest available substrate. The substrate it reacts with most readily in this context is the proteoglycan core protein and the glycosaminoglycan side chain that has bound it, with the result that the iron-loaded perineuronal-net surface becomes a self-degrading substrate that

hydroxylates and fragments itself at the rate at which hydrogen peroxide is delivered to its bound iron. The Ayton and Bush programs have documented the generality of this mechanism in Alzheimer's disease as the *labile iron* hypothesis, in which the elevated free iron of the affected parenchyma is a catalytic accelerator of multiple downstream damage modalities; the application to perineuronal-net biology specifically follows directly from the convergence of the Ayton-Bush iron biology with the Snow-Bishop proteoglycan biology and with the established geography of oligodendrocyte ferroptosis in late Phase II.

The functional consequence of the iron arm is twofold. First, it provides a degradation pathway for the perineuronal net that operates in parallel to enzymatic proteolysis and that is not interdicted by matrix-metalloproteinase inhibition. An MMP-9 inhibitor strategy applied to late-Phase-II tissue blocks the enzymatic cleavage of aggrecan and brevican but does not block the Fenton-catalyzed fragmentation of those same molecules at the side-chain glycosaminoglycan attachment sites, and the parallel degradation pathway continues to dismantle the net even in the presence of complete proteolytic inhibition. Second, the iron arm cooperates with the proteolytic arm in a manner that amplifies both. The hydroxylated and oxidatively fragmented forms of aggrecan and brevican produced by Fenton catalysis are themselves more susceptible to enzymatic cleavage by MMP-9 and the ADAMTS proteases than the native proteoglycan, because the oxidative damage exposes cleavage sites that the native folded protein had protected; and conversely, the smaller fragments produced by initial enzymatic cleavage present more accessible iron-binding sites than the intact net, so that the rate of iron loading accelerates as the matrix begins to disassemble. The two arms therefore operate in positive feedback once the convergence has begun, and the rate of perineuronal-net dissolution accelerates rather than decays as the process proceeds.

The cellular consequence for the parvalbumin-positive interneuron is acute. The intact perineuronal net is, among its other functions, an iron-chelation barrier that holds free iron away from the cell surface and protects the lipid bilayer of the neuron from direct Fenton-catalyzed lipid peroxidation. As the net is dismantled, the chelation function is lost progressively, and the free iron that was previously sequestered in the extracellular matrix becomes available at the neuronal surface and within the perisomatic space. The PV+ interneuron, whose mitochondrial demand is already at the highest end of the cortical and hippocampal range and whose substrate margin has been narrowed by the prior bioenergetic pressures of Phase I, has the smallest tolerance of any neuron in the parenchyma for the additional oxidative load delivered by the loss of its chelation barrier. The consequence is that PV+ interneurons enter their own variant of the oxidative-stress-substrate-failure trajectory that LC neurons entered in Phase I, but they enter it within a clinically rapid window because the substrate has already been narrowed and because the additional pressure is delivered as a step function rather than as the slow accumulation of Phase I. The gamma-frequency drive that the PV+ interneuron population provides to the cortical and hippocampal circuits collapses on a timescale of months to years rather than the decades of the prior phases, and this is the clinical signature of Phase III.

5. The Third Arm: Complement Priming Tags the Synaptic Interface

The third transmission mechanism operates through the complement cascade and directs the proteolytic and oxidative output of the first two arms specifically to the synaptic interface at which the perineuronal net most consequentially constrains circuit function. Where the first two arms specify the mechanism by which the perineuronal-net matrix is dismantled in general, the third arm specifies the mechanism by which the dismantling is targeted to the structures whose loss most rapidly degrades cognitive function: the GABAergic perisomatic terminals of the parvalbumin-positive interneuron and the surrounding glutamatergic synapses on the PV+ cell's dendrites and soma.

The biological basis of the arm is the developmental and pathological repurposing of the complement cascade as a synaptic-pruning signal, established principally by the Stevens program in collaboration with the Shatz, Barres, and Lemere laboratories. Beth Stevens and colleagues demonstrated in a 2007 *Cell* paper, and confirmed in subsequent work, that C1q and the downstream C3 component of the classical complement cascade are deposited on synapses during developmental refinement of neural circuits and that complement-receptor-3 (CR3)-mediated phagocytosis by microglia is the proximate mechanism by which the tagged synapses are removed. The Stevens program has subsequently extended this finding to pathological contexts and has demonstrated that the same complement-tagging mechanism is activated in Alzheimer's disease, with C1q deposition on synapses in the affected hippocampus and cortex preceding the loss of those synapses by an interval consistent with a tagging-then-removal mechanism. The Hong, Beja-Glasser, and Lemere extension of this work in 2016 demonstrated that C1q deficiency or C3 deficiency in mouse models of Alzheimer's disease protects against synaptic loss in a manner that is otherwise refractory to amyloid-targeting interventions, and the protection is specific to synaptic loss rather than to amyloid burden, indicating that the complement-tagging arm is mechanistically distinct from amyloid-driven synaptotoxicity.

The relevance to the perineuronal net is direct but has been clarified only recently. The PNN itself is a substrate for complement tagging, in that the aggrecan and brevican core proteins present epitopes accessible to C1q binding, and C1q deposition on perineuronal-net structures has been documented in human Alzheimer's tissue and in mouse models of the disease. The deposited C1q activates the classical complement cascade through C1r and C1s, which proteolytically cleave C2 and C4 to generate C4d fragments that opsonize the underlying substrate, and the C4d-opsonized

PNN structures present an explicit *eat-me* signal to CR3-expressing microglia. The Crapser program has documented the consequence: post-homeostatic microglia in the vicinity of complement-tagged perineuronal nets engage CR3-mediated stripping of the matrix, and the stripping proceeds in parallel with the proteolytic and Fenton-catalyzed degradation specified in the first two arms.

The third arm therefore does three things simultaneously. It localizes the convergent proteolytic and oxidative output of the first two arms to the specific perineuronal-net structures whose loss most acutely degrades cognitive function. It accelerates the removal of those structures through a CR3-mediated phagocytic pathway that is mechanistically distinct from, but cooperative with, the enzymatic and catalytic pathways. And it explains the regional and cellular specificity of the Phase III degradation: only the perineuronal nets that have received complement tagging are subject to active microglial stripping, and the pattern of complement tagging is determined upstream by the regional pattern of synaptic-loss signaling that establishes the local complement microenvironment. The result is that the Phase III degradation is not a uniform dissolution of all perineuronal nets but a topographically structured loss that follows the geography of complement deposition, which in turn follows the geography of prior synaptic stress, which in turn follows the geography of Phase II microglial engagement with the PANTHOS neuronal substrate. The cascade is therefore self-consistent and self-targeting, with each upstream signal positioning the next downstream effector at the location at which it will be most damaging.

The Stevens program has continued to develop the implications of the complement arm and has identified, in recent work, that the C1q-targeting therapeutic strategy under development by Annexon Biosciences and others is operative in Alzheimer's disease models at a stage consistent with the late-Phase-II to early-Phase-III window of the three-phase framework. The therapeutic implication, developed below, is that the Phase II Phase III transition is open to interdiction by complement modulation in a manner that is mechanistically distinct from MMP-9 inhibition and from iron chelation but that operates on the same substrate and within the same window. The three-arm specification therefore licenses a three-target therapeutic strategy, and the strategy is more powerful than any of its components used alone because the three arms cooperate in positive feedback once the convergence has begun.

6. The Convergence: Three Arms on One Cellular Substrate

The structural claim of this paper is that the three arms of the Phase II Phase III transition mechanism are not independent processes that happen to occur in the same temporal window but are three aspects of a single convergent attack on a single cellular substrate. The substrate is the

parvalbumin-positive interneuron, and more specifically the perineuronal-net sheath that surrounds it, and more specifically still the chondroitin-sulfate glycosaminoglycan side chains of the aggrecan and brevican core proteins that constitute the structural backbone of the sheath. The three arms attack three distinct functions of this substrate simultaneously: the proteolytic arm attacks the structural integrity of the proteoglycan backbone, the iron arm attacks the chelation function of the glycosaminoglycan side chains, and the complement arm attacks the targeting and timing of the attack by directing the proteolytic and phagocytic effectors specifically to the synaptic interface at which the substrate most consequentially constrains circuit function.

The convergence is not a coincidence of three independent attacks but the predictable consequence of the cellular biology of the substrate itself. The perineuronal net is unusual among extracellular matrix structures in that it performs three biologically distinct functions in the same structural element. The first function is mechanical: the net provides a physical envelope that organizes the synaptic terminals impinging on the surface of the PV+ interneuron and that stabilizes the synaptic architecture against rearrangement once the developmental critical period has closed. The second function is biochemical: the chondroitin-sulfate side chains chelate divalent cations including iron and zinc and hold these reactive species away from the lipid bilayer of the underlying neuron, and they also bind growth factors and morphogens including OTX2 and Sema3A in a manner that maintains the high-firing identity of the underlying cell. The third function is signal-isolating: the dense polyanionic surface of the net constitutes a diffusion barrier that slows the access of extracellular molecules to the underlying neuron, with the result that the PV+ interneuron's perisomatic membrane is partially insulated from the ambient signaling environment of the parenchyma. The three functions are performed by overlapping but distinct molecular features of the same structure, and the three attack arms target the three functions correspondingly: proteolysis dismantles the mechanical function, iron-mediated Fenton catalysis dismantles the chelation function, and complement-mediated CR3 stripping dismantles the signal-isolating function.

The substrate vulnerability is therefore not a coincidence of three independent vulnerabilities but a structural property of the perineuronal net itself: a single biological structure performs three distinct protective functions, and the loss of any one of them substantially compromises the underlying neuron's capacity to sustain its high-firing identity, but the simultaneous loss of all three is what produces the rapid clinical degradation characteristic of Phase III. The PV+ interneuron is the Phase III bridgehead in the same sense that the hippocampus is the Phase II bridgehead: it is the cellular environment in which the converging upstream pressures consolidate into a self-sustaining pathology that no longer requires further input from earlier phases to progress, because the loss of the protective envelope creates conditions under which the neuron's own metabolic demand exceeds its substrate availability and the neuron enters a sustained oxidative-substrate-failure trajectory that does not reverse spontaneously.

The geography of Phase III is therefore the geography of the perineuronal-net population, and this geography is well-characterized. PV+ interneurons with perineuronal nets are present in highest density in the cortical layers III–IV of the somatosensory, motor, and prefrontal cortices, in the hippocampal CA1 and CA3 stratum oriens and stratum pyramidale, and in the basolateral amygdala. The hippocampal CA1 region is the site at which the Phase II microglial bridgehead has consolidated, and it is therefore the site at which the convergent attack on the PNN-bearing PV+ population begins. The progression of Phase III to cortical territories proceeds through the same kind of network-anatomical logic that the Phase I Phase II locus-coeruleus bridge established for the brainstem-to-hippocampal transition: the hippocampal projection geometry to cortical targets, combined with the local complement microenvironment established by ongoing PANTHOS pathology and microglial engagement, determines which cortical territories receive the Phase III pressure first and which are spared longest. The clinical phenotype of Phase III — the loss of gamma-frequency drive, the disinhibition of cortical and hippocampal circuits, the emergence of seizure activity, and the rapid degradation of cognitive function — follows this geography, with the medial temporal lobe and the prefrontal cortex showing the earliest functional manifestations and the primary sensory cortices showing the latest.

The bridgehead framing of the PV+ interneuron in Phase III also clarifies why no intervention applied after the bridgehead has consolidated can achieve more than symptomatic preservation of residual function. Once the perineuronal net has been substantially dismantled, the underlying PV+ interneuron has entered the substrate-margin-failure trajectory that LC neurons entered in Phase I, and the trajectory does not spontaneously reverse on any clinically relevant timescale. The intervention window for Phase III therefore opens at the moment of the Phase II Phase III transition and closes at the moment at which the perineuronal-net architecture has been catalytically dismantled, and the window is correspondingly narrow — measured in years rather than decades — and demands a combinatorial therapeutic strategy that addresses all three arms of the convergent attack simultaneously.

7. Why Now: The Substrate Margin and the Prior Bioenergetic Collapse

The transition account developed above raises a question that the three-phase framework has not previously addressed at the level of mechanistic specificity that the therapeutic argument requires. Why does the Phase II Phase III convergence produce clinical decline within a window of years, when each of the three arms — proteolytic, oxidative, and complement-mediated — has

presumably been operating at some baseline level throughout adulthood and could in principle have been expected to produce comparable damage earlier? The answer is that the substrate margin of the parvalbumin-positive interneuron has been narrowed by the cumulative pressures of Phase I and Phase II to a point at which the additional load delivered by the Phase II Phase III transition exceeds the cell's compensatory capacity. The acute decompensation of Phase III is therefore not a function of the magnitude of the Phase II III convergence considered in isolation but a function of the magnitude of the convergence relative to the residual substrate margin that the prior phases have left available.

The mechanism of the prior narrowing is bioenergetic. The PV+ interneuron is among the highest-firing neurons in the cortex and hippocampus, with sustained firing rates that can exceed 50 Hz under physiological conditions and that exceed 200 Hz during gamma-frequency entrainment. The metabolic cost of this firing rate is substantial: each action potential consumes ATP through Na/K-ATPase restoration of the membrane potential, and the cumulative ATP demand of a high-firing PV+ interneuron exceeds that of a low-firing pyramidal cell by approximately an order of magnitude. The cellular machinery that meets this demand is the mitochondrial OXPHOS pipeline, with a particularly heavy reliance on cytochrome c oxidase (Complex IV) activity, and the integrity of this pipeline depends on the quality-control machinery — mitophagy, mitochondrial biogenesis, and the broader autophagy-lysosomal pipeline — that the Bioenergetic Collapse thesis has identified as the substrate of Phase I vulnerability.

The Bioenergetic Collapse thesis has argued that the age-dependent decline of cellular quality-control competence proceeds across all cell types but manifests first in the cell types with the highest baseline metabolic demand, and that the locus coeruleus is the cell type with the highest such demand in the brainstem. The hippocampal and cortical PV+ interneurons are the cell types with the highest such demand in the forebrain. They have therefore been accumulating the consequences of declining quality-control competence throughout Phase I and Phase II in a manner analogous to but slower than the LC: their mitochondria have been accumulating damage that mitophagy has been less and less able to clear, their cytochrome c oxidase activity has been declining at the cellular level, and their substrate margin — the ratio of metabolic demand to substrate availability — has been narrowing throughout the prior decades.

By the moment of the Phase II Phase III transition, the PV+ interneuron's substrate margin has been narrowed to the point at which the cell's high-firing identity is sustained only by the combination of preserved mitochondrial competence and preserved perineuronal-net protection. The PNN provides the iron chelation that protects the mitochondrial machinery from Fenton-catalyzed damage at the cell surface, and it provides the diffusion-barrier function that protects the cell from the ambient oxidative environment of the late-Phase-II parenchyma. The cell can tolerate the loss of either the mitochondrial competence or the PNN protection in isolation for a period of time, but the simultaneous attack on the PNN through the three arms of the Phase II III

convergence — at a moment when the mitochondrial competence has already been narrowed by Phase I — produces an acute substrate insufficiency that the cell cannot compensate.

This framing also explains why the resilience phenotype documented by de Vries, Auer, and Crapser in cognitively intact individuals carrying symptomatic-level amyloid burden depends so heavily on the preservation of perineuronal-net coverage. De Vries and colleagues have documented in human postmortem tissue that cognitively intact individuals with high amyloid burden have substantially preserved PNN coverage relative to cognitively impaired individuals with comparable amyloid, and that the preserved PNN coverage correlates with preserved gamma-frequency function and with preserved PV+ interneuron density. The resilience phenotype is therefore not a function of reduced amyloid generation or improved amyloid clearance but of the preservation of the protective envelope around the most metabolically demanding cell population. The implication for the transition account is direct: the resilient individuals are those in whom the Phase II Phase III convergence has been delayed or arrested, either by genetic factors that reduce one of the three arms (lower MMP-9 activity, lower iron labile pool, lower complement deposition) or by environmental and lifestyle factors that preserve the substrate margin (preserved mitochondrial function through exercise, preserved metabolic fitness through caloric restriction or ketogenic substrate availability, preserved cardiovascular and metabolic health). The resilience phenotype is what the transition account predicts: when the convergent attack on the perineuronal net is delayed or attenuated, the substrate margin of the PV+ interneuron is preserved, and the clinical phenotype of Phase III is correspondingly delayed regardless of the amyloid burden.

The substrate-margin framing also has implications for the design of the Phase III intervention strategy that go beyond the immediate proteolytic, oxidative, and complement targets. Any intervention that improves the mitochondrial competence of the PV+ interneuron — that is, any intervention drawn from the Phase I bioenergetic-restoration class, including NAD-sparing strategies, sirtuin activators, ketogenic substrate provision, and mitophagy inducers — operates synergistically with the Phase III intervention class because it widens the substrate margin against which the residual proteolytic, oxidative, and complement load will operate. The integrated therapeutic logic is therefore not a sequence of phase-specific interventions but a combinatorial logic in which the Phase III intervention is most effective when it is layered on a sustained Phase I substrate-preservation program. The temporal-pharmacology principle that the parent framework has articulated is therefore not merely a sequential ordering but a layered ordering, with the substrate of each earlier phase serving as the persistent backdrop on which the next phase's intervention operates.

8. Therapeutic Implications: The Combinatorial Phase II III Window

The integrated transition model has direct therapeutic implications for the framework's Phase II III window, and the implications are more demanding than the single-target therapeutic logic that the parent framework has articulated for the bracketing phases. Where the Phase I window admits a single-target intervention focused on the bioenergetic substrate of LC vulnerability, and where the Phase II window admits a single-target intervention focused on the ferroptotic biology of oligodendrocyte death, the Phase II III window admits no single-target intervention because the three arms of the convergent attack operate through three distinct molecular mechanisms that are not interdicted by any single drug. The therapeutic logic of the transition window is therefore necessarily combinatorial, and the combinatorial logic has a specific structure that follows from the molecular biology of the three arms.

The proteolytic arm is addressed by matrix-metalloproteinase inhibition. The pharmacological options are well-characterized and span two classes. The broad-spectrum tetracycline antibiotics minocycline and doxycycline are non-selective MMP inhibitors with decades of clinical safety data, with established CNS penetration, and with secondary anti-inflammatory activity through additional mechanisms including microglial suppression. The selective MMP-9 inhibitor JNJ0966 is a newer chemical entity with greater selectivity but less established clinical experience, and intermediate options include the broader hydroxamate-class MMP inhibitors that have been developed for oncology indications. The choice among these options is a tradeoff between selectivity and clinical maturity, with minocycline favored for near-term repurposing on the basis of its established safety profile and doxycycline favored on the basis of its slightly broader MMP spectrum and its long history of use in CNS indications.

The iron arm is addressed by iron chelation. The pharmacological options are again well-characterized. Deferiprone is an orally bioavailable iron chelator with established CNS penetration, currently in clinical use for thalassemia and in active trial for Alzheimer's disease and Parkinson's disease in the programs developed by Bush, Ayton, and colleagues. Deferoxamine is the older parenteral chelator with less convenient administration but extensive clinical experience. Deferasirox is the third option, with broader iron specificity but less CNS penetration. The choice among these is again a tradeoff between convenience, CNS penetration, and side-effect profile, with deferiprone favored on the basis of its oral availability and its specific clinical experience in Alzheimer's contexts.

The complement arm is addressed by complement modulation. The pharmacological options are less mature than the proteolytic and iron classes but are advancing rapidly. The Annexon Biosciences C1q-targeting antibody program is the most advanced clinical effort, with anti-C1q

therapy ANX005 having completed early-stage trials in Guillain-Barré syndrome and other indications and with related programs targeting C1q in Alzheimer's disease and other CNS contexts under development. C5 inhibition through eculizumab and related antibodies is a more downstream option that has been used in non-CNS indications and that may have repurposing potential. The CR3 phagocytosis-blocking strategies are the most upstream option but are at earlier development stages. The choice among these is determined by the development timeline of the available agents, with the C1q-targeting class likely to be the first available for the Phase II III window.

The combinatorial strategy that the transition model recommends is therefore a three-target intervention deployed during a narrow temporal window that opens at the moment of the late-Phase-II homeostatic-to-effector switch and closes at the moment at which the perineuronal-net architecture has been substantially dismantled. The strategy combines MMP inhibition (minocycline or doxycycline as the near-term option, JNJ0966 as the long-term option), iron chelation (deferiprone as the near-term option), and complement modulation (anti-C1q antibody as the developmental option), with the three agents deployed simultaneously rather than sequentially. The simultaneous deployment is required because the three arms cooperate in positive feedback once the convergence has begun, and the interdiction of any single arm in isolation produces only partial protection while the remaining two arms continue to operate.

The temporal targeting of the combinatorial intervention requires a biomarker for the Phase II Phase III transition that the framework does not yet possess at the level of specificity that clinical trial design requires. The candidate biomarkers include CSF or plasma levels of aggrecan and brevican degradation fragments (the *neoepitope* products of MMP-9 and ADAMTS-4/5 cleavage that are signatures of active perineuronal-net dissolution), CSF iron and ferritin levels (as indices of ferroptotic activity in the affected parenchyma), CSF C1q and C3a levels (as indices of complement activation), and imaging-based measures of perineuronal-net coverage using emerging Wisteria-Floribunda-agglutinin-targeted PET tracers. The development of these biomarkers is a prerequisite for the clinical trial of the combinatorial strategy, and the biomarker development should proceed in parallel with the trial design rather than as a sequential precondition.

The therapeutic logic of the Phase II III window also implies a specific architecture for the broader temporal-pharmacology intervention strategy. The Phase I intervention class (PARP inhibitors, NAD-sparing strategies, sirtuin activators) is best deployed continuously from the third or fourth decade of life onward, as a substrate-preservation strategy whose therapeutic effect is maximal over decades and whose maintenance through the Phase II III window provides the bioenergetic backdrop against which the combinatorial Phase II III intervention will operate. The Phase II ferroptosis-inhibitor class is best deployed from the late fifth or early sixth decade onward, as a disease-modifying strategy whose effect is to slow the rate of oligodendrocyte cell death and therefore to delay the moment at which the Phase II III transition becomes imminent. The

combinatorial Phase II III intervention is best deployed in the late sixth or early seventh decade, ideally guided by the biomarker indices specified above, with the goal of interdicting the convergent attack on the perineuronal net at the moment when its initiation can still be prevented. The Phase III intervention class, if applied after the bridgehead has consolidated, is necessarily palliative and preservative rather than disease-modifying, and its goal is to extend the window of residual cognitive function rather than to reverse the pathology that has already occurred.

The implication for clinical trial design is that the conventional architecture of Alzheimer's trials — single-agent intervention in a clinically diagnosed cohort, with cognitive endpoint over an eighteen-month window — is mismatched to the structure of the disease as the three-phase framework describes it. The appropriate trial architecture for the combinatorial Phase II III intervention is a multi-agent intervention in a biomarker-defined late-Phase-II cohort, with biomarker endpoints measuring perineuronal-net coverage and proteolytic-fragment levels over a window of three to five years and with cognitive endpoints serving as the long-term validation rather than the proximate trial endpoint. The regulatory and trial-design implications of this requirement are substantial and are taken up in the broader therapeutic-landscape paper that accompanies the trilogy.

9. What This Resolves, What This Does Not

The transition model developed in this paper resolves the Phase II Phase III boundary problem in the three-phase framework by specifying a tripartite mechanism that converges on a single cellular substrate, and it does so in a way that integrates the existing Homeostatic Microglial Collapse, Bioenergetic Collapse, and Convergent Synaptic Collapse theses without modifying the architecture of any of them. It identifies the parvalbumin-positive interneuron's perineuronal net as the cellular substrate on which the three arms converge, and it explains the clinical phenotype of Phase III as the predictable consequence of the simultaneous attack on the three protective functions of the perineuronal net by three mechanistically distinct effector pathways. The therapeutic implications for the Phase II III window follow directly from the transition mechanism and specify a combinatorial intervention strategy whose architecture is different in kind from the single-target interventions of the bracketing phases.

The model also clarifies several features of the three-phase framework that had previously been articulated only schematically. The substrate-margin framing developed in section 7 specifies why the Phase III decompensation is acute despite the gradualness of the upstream attacks, and it identifies the layered logic by which Phase I bioenergetic preservation operates synergistically with

Phase III matrix-preservation rather than as a separate sequential intervention. The bridgehead framing of the PV+ interneuron parallels the bridgehead framing of the hippocampus in the LC bridge paper, and it suggests that each phase of the framework may be productively understood as the consolidation of upstream pressures on a specific cellular environment that contains within itself the substrates of the next phase's vulnerability — a recursive structure whose generalization across the three boundaries is a candidate organizing principle for the framework as a whole. The complement arm framing in section 5 connects the three-phase framework to the Stevens-Shatz-Lemere complement-mediated-pruning program in a manner that the prior literature had not articulated, and it positions the developing C1q-targeting therapeutic class as a Phase II III intervention with mechanistic justification rather than as an opportunistic application of a developmental biology finding to a degenerative context.

Several questions remain unresolved and are flagged here as candidates for further work. The first is the empirical status of the iron-binding step in the second arm. The Snow and Bishop work on glycosaminoglycan iron binding is well-established at the biochemical level, and the Ayton-Bush work on labile iron in Alzheimer's tissue is well-established at the histological level, but the specific binding of liberated iron to perineuronal-net glycosaminoglycans in late-Phase-II tissue has not been directly demonstrated in the specificity that the transition mechanism requires. The demonstration is technically tractable through a combination of iron-imaging modalities (Perls' Prussian blue staining, synchrotron-based X-ray fluorescence microscopy) and PNN imaging (Wisteria-Floribunda-agglutinin lectin labeling), and the necessary tissue cohorts are available through the standard human Alzheimer's postmortem archives. The experiment is a near-term priority.

The second is the precise threshold at which the proteolytic switch in the first arm becomes irreversible. The Marschallinger LDAM characterization establishes that lipid-droplet accumulation reaches a phenotypic threshold at which microglial function is qualitatively altered, but the threshold itself is poorly characterized in absolute terms (number of droplets per cell, total lipid mass per cell, NLRP3 activation state), and the question of whether the threshold is a hard discontinuity or a graded function with a steep slope is unresolved. The clinical implication is substantial: if the threshold is graded, then partial interdiction of the lipid load through metabolic intervention may produce proportional benefit, whereas if it is discontinuous, then interventions are only useful if they prevent the threshold from being crossed at all. The resolution of this question requires single-cell metabolic and transcriptional profiling of LDAM populations across the late-Phase-II to Phase III transition in human tissue, and the necessary platforms are now available.

The third is the role of the dorsal raphe nucleus and the serotonergic system in the Phase II III transition. The dorsal raphe is identified alongside the locus coeruleus in the Phase I locus, but its contribution to the Phase II III transition has not been articulated. The 5-HT receptor expression on microglia, the serotonergic regulation of complement deposition in some contexts, and the

documented vulnerability of serotonergic neurons in late-stage Alzheimer's disease together suggest that a parallel serotonergic arm of the transition mechanism may operate alongside the noradrenergic-derived mechanisms developed in the LC bridge paper and in the present synthesis, but the parallel arm has not been specified and is a candidate for future work.

The fourth is the integration of the transition mechanism with the APOE-isoform and lipid-availability variables identified in the broader framework. The APOE4 isoform has been associated with elevated complement activation, with altered microglial lipid handling, and with elevated MMP-9 activity in multiple studies, and the integration of the APOE genotype with the three-arm transition mechanism is a natural extension that would specify the genetic gating of the transition risk and would inform the stratification of intervention candidates by genotype. The integration has not been undertaken in the present paper and is left as a candidate for the next synthesis.

These open questions notwithstanding, the central claim is straightforward and, we believe, mechanistically defensible. The Phase II Phase III transition is the convergent attack on the perineuronal-net sheath of the parvalbumin-positive interneuron by three mechanistically distinct effector pathways — the proteolytic output of post-homeostatic microglia after the LDAM-NLRP3 metabolic switch has licensed matrix-metalloproteinase transcription; the catalytic conversion of the iron-loaded glycosaminoglycan backbone into a self-degrading Fenton substrate after the liberation of redox-active iron from ferroptotic oligodendrocytes; and the complement priming of the synaptic interface by C1q deposition and downstream C4d opsonization that directs the proteolytic and phagocytic output specifically to the PV+ interneuron's perisomatic territory. The three arms converge on a single substrate because the perineuronal net is the only structure in the parenchyma that simultaneously performs three biologically distinct protective functions whose loss tips the PV+ interneuron's already-narrowed substrate margin below the threshold required to sustain its gamma-frequency drive. The PV+ interneuron is the Phase III bridgehead, the perineuronal net is the substrate of its protection, and the three-arm convergence is the mechanism by which the Phase II homeostatic collapse executes the Phase III matrix disintegration. The three-phase framework's transition problem is, on this account, not a separate mechanistic puzzle but the expected consequence of the cellular biology of the substrate that the framework has already identified as the Phase III locus, and the resolution of the transition problem completes the mechanistic specification of the framework that the temporal-pharmacology argument requires.

◆

Companion to The Locus Coeruleus Bridge, Convergent Synaptic Collapse, Homeostatic Microglial Collapse, and Bioenergetic Collapse. Third in the series of inter-phase transition syntheses prepared under the Organic Network Synthesis methodology.