1 Beyond Current Frontiers: Three Theoretical Advances

1.1 Paper 1: Adaptive Conductance Economics → Predictive Topological Liquidity Theory

1.1.1 What the current reframe gets right

The Physarum-style adaptive conductance model is solid but it is still *reactive* — conductance updates follow flow, flow follows gradient. Every existing network-flow model in econophysics (including the mycorrhizal-inspired ones) shares this limitation: the network responds to signals after they arrive.

1.1.2 The move beyond

Claim: Mycorrhizal networks are not just adaptive transport systems. They are predictive allocation systems with topological memory.

The biological literature already contains the seed. Mycorrhizal common networks show pre-allocation behavior: resource routing shifts *before* a connected host shows measurable stress, correlated with seasonal cues, neighbor identity, and historical demand patterns. This is not gradient-following. This is learned anticipation encoded in network structure.

The novel theoretical object is **conductance memory** — the idea that the topology and weight distribution of the network at time t encodes a compressed model of the demand-supply history, and that this structural memory generates *predictive flows* that precede price signals.

1.1.3 Formal extension

Introduce a **memory kernel** into the conductance dynamics:

$$\frac{dK_{ij}}{dt} = \alpha |J_{ij}|^m - \beta K_{ij} + \underbrace{\int_0^t \mathcal{M}(t-\tau, \mathcal{G}_\tau)\, d\tau}_{\text{topological memory term}}$$

where $\mathcal{M}$ is a memory kernel that depends not just on past flows but on the *graph Laplacian spectrum* $\mathcal{G}_\tau$ at earlier times. This means the network remembers not just "how much flowed" but "what shape the network was in when crisis happened."

This creates three phenomena that do not exist in any current model:

1. Anticipatory liquidity routing. The network pre-positions capital at nodes whose spectral neighborhood resembles past crisis configurations — before any local signal arrives. This is analogous to how the hippocampus pre-activates place cells for upcoming locations during navigation, but applied to economic topology. The prediction is falsifiable: in a simulated network with memory, liquidity should begin moving toward a node *before* that node's shadow price rises, if the surrounding topology matches a previously learned stress pattern.

2. Topological anti-fragility. Standard adaptive networks strengthen used edges (Hebbian). The memory kernel additionally allows the network to *pre-strengthen edges that were critical during past recovery*, even if they carry no current flow. This produces a structural reserve — dormant high-conductance

paths that activate only when the graph Laplacian approaches a previously experienced crisis eigenspectrum. The analog in mycorrhizal biology is the maintenance of hyphal connections to temporarily non-productive hosts — connections that appear wasteful until environmental shift makes them essential.

3. Information-theoretic risk as an emergent observable. Define systemic risk not as a function of node-level variables but as the *spectral entropy* of the flow-weighted graph Laplacian:

$$\mathcal{R}(t) = -\sum_k \hat{\lambda}_k \ln \hat{\lambda}_k, \qquad \hat{\lambda}_k = \frac{\lambda_k}{\sum_j \lambda_j}$$

where λ_k are the eigenvalues of the weighted Laplacian. When spectral entropy collapses (eigenvalue concentration), the network has lost its distributed routing capacity and is approaching a brittle, centralized configuration — the mathematical signature of too-big-to-fail. When spectral entropy is high, liquidity has many independent routing paths. This gives you a single scalar that is both computable from the network state and predictive of cascade failure, something no current systemic risk measure achieves from topology alone.

1.1.4 What this changes

The existing adaptive conductance model is a better metaphor. This version is a *different kind of dynamical system* — one where the network's own history-dependent structure becomes a computational substrate for prediction. The closest mathematical relative is not Physarum optimization but reservoir computing on graphs: the network's transient dynamics in response to perturbation contain information about the future because the network has been shaped by the past.

Revised title: *Predictive Topological Liquidity: A Memory-Kernel Network Model for Anticipatory Capital Routing and Spectral Risk Detection*

1.2 Paper 2: Bioelectric Oncology → Topological Defect Theory of Carcinogenesis

1.2.1 What the current reframe gets right

The PDE framework for V, T, I coupling is the right mathematical skeleton. Bioelectric control variables are real. But the field currently treats membrane potential as a *scalar* property of individual cells that needs to be "corrected." This is the wrong level of description.

1.2.2 The move beyond

Claim: Cancer is not a cellular disease with bioelectric correlates. It is a topological defect in the tissue-scale bioelectric field.

The Levin lab and related work already show that what matters for morphogenesis is not the membrane potential of any single cell but the *spatial pattern* of V_{mem} across the tissue. Normal tissue maintains a bioelectric pattern with specific symmetries and boundary conditions. Individual cell depolarization does not cause cancer. *Disruption of the pattern's topological structure* does.

The novel theoretical object is the **bioelectric order parameter field** and its **topological defects**.

1.2.3 Formal extension

Treat the tissue-scale bioelectric state not as a scalar field $V(x,t)$ but as a vector order parameter $\mathbf{n}(x,t)$ encoding both the local membrane potential and its gradient orientation:

$$\mathbf{n}(x,t) = \frac{\nabla V}{|\nabla V|}$$

In healthy tissue, $\mathbf{n} \approx \mathbf{n}$ — points or regions where the field orientation is undefined because the surrounding gradients form a vortex, source, or sink pattern that cannot be smoothly continued through the center.

Define the topological charge at a point:

$$q = \frac{1}{2\pi} \oint_{\gamma} d\theta$$

where γ is a closed path around the suspected defect and θ is the angle of $\mathbf{n}$. A topological charge $q \neq 0$ is a defect. The key insight: **a $q = +1$ defect is a region where all bioelectric gradients point inward or outward — a sink or source in the field. This is precisely the bioelectric signature of a tumor: a depolarized island surrounded by normally polarized tissue, creating a radial gradient pattern with undefined orientation at the center.**

This gives you three consequences that go beyond anything in the current literature:

1. Topological protection of the malignant state. In condensed matter physics, topological defects cannot be removed by smooth, local deformations — they are *topologically protected*. This offers a novel explanation for why tumors are so resistant to local intervention: you cannot fix a topological defect by adjusting the field values near it. You can only annihilate it by bringing it into contact with an anti-defect ($q = -1$) or by globally restructuring the field. This reframes therapeutic resistance not as molecular evolution but as geometric inevitability.

2. Defect-mediated metastasis. In liquid crystals and superfluids, topological defects are mobile — they move through the medium and interact with each other. If tumor initiation sites are field defects, then metastasis is *defect migration*: the topological singularity propagates through the tissue's bioelectric field, and a secondary tumor forms where it stabilizes. The prediction: metastatic sites should show bioelectric gradient vortex signatures *before* detectable tumor cell colonization. The defect arrives before the cells.

3. Topological therapy: annihilation by anti-defect injection. Instead of trying to locally repolarize tumor cells (which topological protection guarantees will fail for the core defect), the therapeutic strategy becomes: *inject an anti-defect*. Create a controlled $q = -1$ bioelectric pattern (a region where gradients form the complementary vortex) near the tumor. When defect and anti-defect meet, they annihilate — the field smoothly reconnects, the topological singularity disappears, and the cells at the former defect core lose the field-level instruction that maintained the malignant state.

Formally, track defect positions $\mathbf{r}_\alpha(t)$ as emergent degrees of freedom:

$$\frac{d\mathbf{r}_\alpha}{dt} = -\mu_\alpha \nabla_{\mathbf{r}_\alpha} \mathcal{F}[\mathbf{n}] + \sum_{\beta \neq \alpha} \mathbf{f}_{\alpha\beta}(|\mathbf{r}_\alpha - \mathbf{r}_\beta|) + \boldsymbol{\xi}(t)$$

where $\mathcal{F}$ is the free energy of the bioelectric field configuration, $\mathbf{f}_{\alpha\beta}$ is the defect–defect interaction (attractive for opposite charges, repulsive for same), and

ξ captures stochastic microenvironmental forcing. Therapy enters as a controlled boundary condition that creates $q = -1 \approx \nabla V$ across tissue, compute the winding number around suspected lesions, and determine whether $q \neq 0$ at tumor sites and $q = 0$ in healthy tissue. If this holds, you have a topological biomarker for cancer that precedes histological diagnosis.

Revised title: *Topological Defect Theory of Carcinogenesis: Tumor Initiation, Metastasis, and Annihilation as Dynamics of Bioelectric Field Singularities*

1.3 Paper 3: Carbon Removal → Autotrophic Reactor Architecture

1.3.1 What the current reframe gets right

The bounded electrochemical reactor model is the only defensible version of the original idea. But it still inherits the standard DAC problem: every reactor is a cost center. You build it, you run it, you pay for energy, the catalyst degrades, you replace it. Scaling is linear in cost. This is why no DAC pathway currently projects below $\approx \frac{dN}{dt} = \eta_{\text{struct}} \cdot R_{\text{carbon}}(N) - \delta N$\$

where N is the number of operational reactor units, $R_{\text{carbon}}(N)$ is the total solid carbon production rate of the fleet, η_{struct} is the fraction of output carbon structurally suitable for reactor construction, and δ is the unit failure/degradation rate.

The critical threshold is:

$$\eta_{\text{struct}} \cdot r_{\text{carbon}} > \delta$$

where r_{carbon} is the per-unit carbon production rate. Above this threshold, the fleet grows *faster than it degrades* — it is autotrophic. Below it, you need external manufacturing input forever.

This reframes the entire carbon removal problem from "how do we pay for enough reactors" to "how do we engineer the output material to be reactor-grade."

Three consequences:

1. Self-constructing carbon infrastructure. Recent work has demonstrated room-temperature liquid-metal catalytic conversion of CO_2 to solid carbonaceous materials including amorphous carbon, graphitic flakes, and in some configurations, carbon nanotubes and nanofibers. These are not waste products — they are *structural materials*. Carbon nanotubes have tensile strength exceeding steel. Graphitic carbon is chemically stable and electrically conductive. The theoretical move: design the electrochemical conversion pathway so that its output is optimized not just for carbon mass but for structural utility. The reactor's waste becomes its offspring's skeleton.

Formally, replace the simple conversion rate with a product-quality tensor:

$$R_{\text{carbon}} = \int_{\Omega} k_{\text{red}}(\phi, T, \text{cat}) \cdot c(x)\, dx, \qquad \text{with product distribution } P(s|\phi, T, \text{cat})$$

where s indexes the structural quality of the solid carbon product (amorphous powder → graphitic sheets → nanotube/nanofiber) and the distribution P is tunable via electrode potential ϕ, temperature T, and catalyst composition. The engineering problem becomes: find the (ϕ, T, cat) operating point that maximizes η_{struct} while maintaining acceptable overall R_{carbon}.

2. Thermodynamic bootstrapping from dissolution free energy. Current DAC systems waste the free energy of CO_2 dissolution. When CO_2 dissolves in

alkaline or amine solutions, the reaction is exothermic — this enthalpy is currently discarded as low-grade heat. The novel thermodynamic architecture: couple the capture exotherm *directly* to the electrochemical conversion endotherm via a heat-integrated reactor where the capture and conversion zones are in thermal contact.

Define the net energy balance:

$$E_{\text{net}} = E_{\text{electro}} - Q_{\text{capture}} \cdot \eta_{\text{thermal-coupling}} - W_{\text{expansion}}$$

where Q_{capture} is the capture exotherm, $\eta_{\text{thermal-coupling}}$ is the efficiency of thermal integration between capture and conversion zones, and $W_{\text{expansion}}$ is recoverable work from gas expansion during absorption. Even modest thermal coupling ($\eta_{\text{thermal-coupling}} \sim 0.3$) reduces the external energy input by 15–25%, fundamentally changing the cost curve.

3. Exponential vs. linear scaling regime. The key quantitative prediction of the autotrophic framework:

In the standard (non-autotrophic) regime, fleet carbon removal scales as:

$$C_{\text{removed}}(t) = N_0 \cdot r_{\text{carbon}} \cdot t$$

Linear. Every unit must be externally manufactured and funded.

In the autotrophic regime (above the critical threshold), fleet growth is approximately exponential during the unconstrained phase:

$$N(t) \approx N_0 \exp\left[(\eta_{\text{struct}} \cdot r_{\text{carbon}} - \delta) \cdot t\right]$$

and cumulative removal goes as:

$$C_{\text{removed}}(t) \approx \frac{N_0 \cdot r_{\text{carbon}}}{\eta_{\text{struct}} \cdot r_{\text{carbon}} - \delta} \left[\exp((\eta_{\text{struct}} \cdot r_{\text{carbon}} - \delta)t) - 1\right]$$

This is the exponential scaling curve that all current DAC projections lack, achieved *without* releasing anything into the open atmosphere. The self-propagation is industrial, not molecular. The "reproduction" is factory-building-factory, not catalyst-catalyzing-catalyst.

The saturation constraints are real and must be modeled: CO_2 feedstock concentration declines as the fleet scales (success changes the boundary condition), catalyst raw materials besides carbon (metals, electrolyte) remain external inputs, and land/energy footprint grows with N. So the full model includes logistic saturation:

$$\frac{dN}{dt} = (\eta_{\text{struct}} \cdot r_{\text{carbon}} - \delta)N\left(1 - \frac{N}{N_{\text{max}}(t)}\right)$$

where $N_{\text{max}}(t)$ is the carrying capacity set by non-carbon resource constraints and declining atmospheric CO_2.

1.3.2 What this changes

Current carbon removal engineering treats every reactor as an independent cost center requiring external capital, external manufacturing, and external materials for its entire operational life. The autotrophic architecture treats the reactor fleet as a *self-replicating industrial organism* whose growth rate is determined by the structural quality of its own carbon output. This converts the scaling problem from "how much money" to "how good is the material" — an engineering optimization rather than an economic negotiation.

The falsifiable prediction: if η_{struct} for room-temperature liquid-metal CO_2 conversion can be pushed above ≈0.15 (15% of output carbon is structurally reactor-grade), and per-unit r_{carbon} exceeds degradation rate δ by a factor of 10, the fleet doubling time drops below 2 years. At that point, a 100-unit seed fleet reaches gigatonne-scale removal within 15 years without proportional capital scaling. That is the number to beat.

Revised title: *Autotrophic Carbon Removal: Self-Constructing Reactor Architectures via Structural Solid Carbon Output*

1.4 Cross-Paper Theoretical Unification

There is a hidden structural unity across all three papers that goes beyond the surface-level "bio-inspired systems" framing:

All three theories concern the same abstract problem: how a distributed system uses its own output to modify its own structure, and how this self-modification either stabilizes or destabilizes the system.

- In Paper 1, the network's flow history modifies its topology (conductance memory), and this self-modification either produces anticipatory resilience or brittle lock-in depending on spectral entropy.
- In Paper 2, the tissue's bioelectric field pattern either maintains its topological integrity (health) or nucleates defects that are self-stabilizing (cancer), and therapy is the controlled destruction of those self-reinforcing topological structures.
- In Paper 3, the reactor fleet's carbon output either builds or fails to build new reactor capacity, and the system either achieves autotrophic exponential growth or remains trapped in linear cost-scaling.

The unifying mathematical structure is a **self-modifying dynamical system** where the state x evolves on a structure $\mathcal{S}$ that is itself a function of x's history:

$$\dot{x} = F(x, \mathcal{S}), \qquad \dot{\mathcal{S}} = G(x, \mathcal{S}, \mathcal{H}[x])$$

where $\mathcal{H}[x]$ is a history functional. The stability, defect structure, and growth properties of such systems are not well-characterized in current mathematics. A fourth paper — purely mathematical — characterizing the bifurcation theory of self-modifying dynamical systems with history-dependent structure would be the capstone of this portfolio.

Capstone title: *Bifurcation Theory of Self-Modifying Dynamical Systems: Stability, Defects, and Autotrophic Growth in History-Dependent Networks*