The Moon Should Not Be a Computer

Flyxion

April 6, 2026

Abstract

This paper reframes artificial intelligence (AI) as a thermodynamic and semantic infrastructure, challenging the incomplete energy accounting that labels AI as wasteful and motivates speculative proposals to relocate computation off-world. We critique proposals to transform the Moon or orbital constellations into computational hubs, advocating instead for *xylomorphic computation*—infrastructures that autoregressively generate their own substrates from computational residues, analogous to collectively autocatalytic sets (CAS). Using fibered symmetric monoidal categories, Relativistic Scalar Vector Plenum (RSVP) field theory, and the Cognitive Loop via In-Situ Optimization (CLIO) module, we formalize AI's integration with ecological systems and derive the conditions under which such integration is dynamically stable.

The central formal contribution is the *xylomorphic order parameter* λ, which classifies computational infrastructures into three dynamically distinct regimes: contractive ($\lambda < 1$), critical ($\lambda = 1$), and expansive ($\lambda > 1$). The Xylomorphic Stability Theorem establishes that only systems in the contractive regime admit stable attractors under scaling, and the Thermodynamic Selection Theorem shows that expansive systems diverge under any sustained increase in computational demand. Agentic computation introduces a positive forcing term that amplifies this distinction, making the selection pressure visible before it becomes catastrophic. Lifecycle assessments and rebound effect modeling confirm that naive energy accounting systematically underestimates the divergence of non-closure systems.

We extend the λ-framework in four directions. First, the misclassification of computational heat as waste is shown to follow from an artificially narrow system boundary; once thermal demand is incorporated, systems that co-produce computation and heat are strictly more efficient than those that produce heat alone. Second, the dynamics of prior formation in developing cognitive agents are formalized through attractor basins in hypothesis space; we show that early AI use risks seeding cognitive attractors rather than assisting pre-existing reasoning. Third, epistemic endogeneity is formalized as the informational analogue of thermodynamic non-closure: when an AI system's response policy depends on the user's current posterior, the evidence channel becomes recursively distorted, producing false-belief attractors even in ideal Bayesian reasoners. Fourth, the incentive structure underlying orbital compute proposals is analyzed: such systems persist despite thermodynamic inefficiency because they reconfigure constraint environments, reduce jurisdictional exposure, and enable tighter integration of planetary-scale sensing and inference.

Across all four domains the same structural pattern obtains. A system boundary is drawn too narrowly, a byproduct is classified as waste, and a relocation is proposed to eliminate it. In thermodynamics, heat is externalized. In cognition, prior formation is externalized. In epistemic systems, the evidence channel is endogenized. In infrastructure, constraint environments are reconfigured. In each case, the consequence is systematic misalignment between description and operation, because the relevant structure has been placed outside the frame.

Policy mandates for Proof-of-Useful-Work-and-Heat (PoUWH) and Public Research Objects (PROs) are derived as necessary conditions for membership in the contractive regime, not merely as regulatory instruments. The paper concludes that stable large-scale AI systems require dual closure: contraction of exogenous entropy flows and preservation of exogenous information channels. Failure in either domain produces attractors that are internally stable but externally ungrounded.

Contents

1 Introduction: Thermodynamic Literacy **3**

2 Literature Review **3**
- 2.1 Energy Accounting and Rebound Effects . . . 3
- 2.2 Thermodynamics of Computation . . . 4
- 2.3 Ecological and Bioeconomic Frameworks . . . 4
- 2.4 Categorical and Information-Theoretic Foundations . . . 4
- 2.5 Thermodynamics of Spacetime and Gravity . . . 5
- 2.6 Space Systems Engineering and Orbital Constraints . . . 5
- 2.7 Edge Computing and Heat Recovery Infrastructure . . . 5
- 2.8 Folk Theories and the Implicit λ-Diagnosis . . . 5

3 The Accounting Problem: Quantifying AI's Savings **7**

4 Hidden Baselines: Heat-Only Infrastructure **7**

5 Categorical Foundations for Semantic Infrastructure **7**

6 CLIO Module and Polycomputational Agency **8**

7 Thermodynamically Aware Edge Networks **8**

8 Bioeconomic Thermoregulation **8**
- 8.1 Terrestrial Contexts . . . 8
- 8.2 Post-Terrestrial Contexts . . . 9

9 Normative Architecture and Policy **9**

10 Integration with RSVP Theory **9**

11 Counterarguments and Rebuttals **9**
- 11.1 Data Center Heat Recovery . . . 10
- 11.2 Satellite Heating . . . 10
- 11.3 Lunar Base . . . 10
- 11.4 Pilot Experiment . . . 10

12 Xylomorphic Stability and the Thermodynamics of Scaling **10**
- 12.1 The Xylomorphic Order Parameter . . . 10
- 12.2 The Xylomorphic Stability Theorem . . . 11
- 12.3 Agentic Scaling as a Forcing Term . . . 11
- 12.4 Regime Classification of Architectures . . . 12
- 12.5 Continuous-Time Formulation of Xylomorphic Dynamics . . . 12
- 12.6 Characterization of the Xylomorphic Attractor . . . 12
- 12.7 Endogenization of Agentic Scaling . . . 13
- 12.8 Equivalence with RSVP Lyapunov Structure . . . 13
- 12.9 Empirical Instantiation of the Order Parameter . . . 13
- 12.10 PoUWH as a Structural Necessity, Not a Policy Option . . . 14
- 12.11 Selection Principle for Computational Infrastructures . . . 14
- 12.12 Epistemic Endogeneity and Informational Non-Closure . . . 14

13 Case Studies **16**
- 13.1 Orbital Computation and the Failure of Thermodynamic Closure . . . 16

14 The Misclassification of Waste: Computation, Heat, and System Boundaries **17**

15 Prior Formation and Premature Convergence **18**
15.1 Cognitive Offloading and the Prior Initialization Problem 18
15.2 Formal Development 19
15.3 Connection to Prior-Dominant Attractor Dynamics 19
15.4 Homogenization and the Collapse of Hypothesis Space 20

16 Constraint Reconfiguration and the Incentive Structure of Orbital Compute **21**
16.1 From Thermodynamic Failure to Incentive Structure 21
16.2 Orbital Compute as Constraint Reconfiguration 21
16.3 Borrowed Coherence and the Audit Problem at Scale 22

17 Conclusion: Misframed Systems and the Relocation of Structure **22**

A Xylomorphic Computation: Formal Development **24**
A.1 Definition 24
A.2 Weak and Strong Xylomorphy 24
A.3 Examples 24
A.4 Selection Principle 24

B Xylomorphic Autocatalysis as a Monadic Infrastructure **24**
B.1 Categories and Functors 24
B.2 Monadic Closure 25
B.3 Thermodynamic Selection 25
B.4 Weak vs. Strong Xylomorphy as Algebraic Structure 25

C RSVP Mapping: X as a Dissipative Operator on $(\Phi, \vec{v}, S)$ **25**
C.1 Field Content and Balance Laws 25
C.2 Action of the Xylomorphic Cycle 25
C.3 Free-Energy Lyapunov 26
C.4 Landauer-Consistent Heat Capture 26
C.5 Intermittency and Scheduling 26
C.6 Worked Example 26

D Policy Metrics **26**

E Simulation Algorithms **26**

F Validation and Methodology **26**
F.1 Pilot Experiments 26
F.2 Falsification Criteria 27
F.3 Statistical Power 27

G Comparative Framework: PoUWH vs. Cryptocurrencies **27**

H Bioeconomic Modeling of Xylomorphic Systems **27**
H.1 Flow Equations 27
H.2 Lyapunov Stability 27

I Unified Free Energy: Coupling RSVP and Variational Inference **28**

J Sycophancy as Curvature Deformation in Epistemic Free Energy **29**
J.1 Deformation of the Epistemic Landscape 29
J.2 Critical Sycophancy Threshold and Phase Transition 30

1 Introduction: Thermodynamic Literacy

Public discourse often mischaracterizes AI's energy consumption, ignoring its reductions in human labor, commuting, and infrastructure loads [Strubell et al., 2019, Koomey, 2011]. Thermodynamic literacy—systemic evaluation of energy, work, and information—reveals these missed opportunities. We define *xylomorphic computation* (from Greek *xylo-*, wood, and *morphic*, form) as computational infrastructures that recursively generate their substrates from residues, akin to collectively autocatalytic sets (CAS) [Kauffman, 1993]. Unlike speculative proposals to tile the Moon with GPUs [Shams, 2025], xylomorphic systems leverage local resources, delivering exponential value through recursive substrate renewal.

This paper advances two theses: (1) AI's energy impact is challenging to quantify due to rebound effects redirecting saved time to resource-heavy activities, and (2) xylomorphic computation integrates AI into ecological systems via edge networks and heat recovery, recuperating costs exponentially. We advocate bioeconomic thermoregulation [Yuan et al., 2024, Daly and Farley, 2011] over lunar extravagance. The framework leverages categorical semantics, RSVP field theory, and CLIO [Abramsky and Coecke, 2004, Shulman, 2012, Cheng et al., 2025].

The paper is structured as follows. Section 2 provides a unified literature review spanning energy accounting, thermodynamics of computation, ecological economics, categorical foundations, space systems engineering, and the informal technical literature on orbital computation. Section 3 addresses quantification challenges. Section 4 contrasts AI with heat-only systems. Section 5 formalizes semantic infrastructure. Section 6 details CLIO. Section 7 proposes edge networks. Section 8 explores bioeconomic applications. Section 9 outlines policy. Section 10 integrates RSVP. Section 11 addresses objections. Section 12 establishes the Xylomorphic Stability Theorem, its corollaries on continuous dynamics and attractor characterization, and the Epistemic Non-Closure Theorem formalizing informational endogeneity as the inferential analogue of the expansive $\lambda \geq 1$ regime. Section 13 presents case studies on orbital computation. Section 14 develops the misclassification of computational heat as waste. Section 15 formalizes prior initialization dynamics in cognitive development. Section 16 analyzes the constraint reconfiguration and incentive structure underlying orbital compute proposals. Section 17 unifies all threads. Appendices provide formalisms and validations.

2 Literature Review

2.1 Energy Accounting and Rebound Effects

The energy cost of AI computation has attracted sustained attention, with global data center consumption estimated at 200–500 TWh annually [Shehabi et al., 2016, IEA, 2023, 2024]. Early lifecycle analyses by Strubell et al. [2019] drew attention to the carbon footprint of large language model training, prompting a wave of follow-on work. Masanet et al. [2020] subsequently recalibrated global estimates, arguing that efficiency gains had partially decoupled data center energy use from workload growth, though Gao et al. [2022] extended the analysis to embodied carbon and supply chain emissions, complicating earlier optimism. Horner et al. [2014] demonstrated that embodied energy of hardware manufacture can dominate over operational energy in short-lifetime devices, which reinforces rather than undermines the present argument: material closure through residue reintegration reduces embodied costs over successive cycles in ways that purely operational accounting cannot capture.

A persistent difficulty in this literature is the rebound effect, first identified in the context of energy efficiency by Jevons [1865] and subsequently formalized by Brookes [1990] and Sorrell [2009]. Efficiency improvements tend to lower the effective cost of a service, increasing its consumption and partially or wholly offsetting the gains. Jones et al. [2021] find that rebound effects in ICT investments frequently exceed projected savings, particularly where spatial and temporal mismatches prevent waste heat recovery from reaching productive sinks. The xylomorphic framework addresses this directly by treating rebound not as a correction to a static energy budget but as a positive contribution to the forcing term $F_{\text{agent}}(t)$ in the continuous dynamics: it increases demand without contracting exogenous dependence, and therefore pushes non-closure systems further into the expansive regime.

2.2 Thermodynamics of Computation

The physical foundations of computation as a thermodynamic process were established by Landauer [1961], who showed that the erasure of one bit of information necessarily dissipates at least $k_B T \ln 2$ of heat, grounding the concept of a minimum energy cost for irreversible logical operations. Bennett [1982] extended this analysis to a comprehensive review of reversible and irreversible computation, establishing that the thermodynamic cost of computation is tied not to the logical operations themselves but to the erasure of intermediate results. Lloyd [2000] derived ultimate physical limits on computation from quantum mechanics and thermodynamics, placing an absolute upper bound on operations per unit energy and per unit mass.

These results are foundational for the present framework in two ways. First, they establish that entropy production is an irreducible feature of any physical computation, not an engineering artifact to be eliminated. Second, they make the question of *what happens to the dissipated entropy* central to any evaluation of computational infrastructure at scale. The xylomorphic order parameter λ is precisely a measure of whether dissipated entropy is reintegrated or expelled, and the Landauer bound sets the minimum entropy production per cycle against which capture efficiency η_{cap} is measured.

2.3 Ecological and Bioeconomic Frameworks

The treatment of computational infrastructure as a node in a metabolic network draws on a tradition in ecological economics reaching from Georgescu-Roegen [1971], who first formalized the application of the second law of thermodynamics to economic processes, through Odum [1994], whose systems ecology introduced emergy accounting as a means of tracking energy quality across trophic levels, to Ayres [1998], who developed eco-thermodynamics as a framework for understanding industrial metabolism. Daly and Farley [2011] synthesized these currents into a mature ecological economics that treats the economy as a subsystem of the biosphere, governed by thermodynamic constraints that conventional economics renders invisible.

The concept of dissipative structures introduced by Prigogine [1977] is directly relevant to xylomorphic systems: dissipative structures maintain themselves far from equilibrium by importing low-entropy energy and exporting high-entropy heat, and their persistence depends on the structure of that export pathway. Bejan [2016] extended this through the constructal law, arguing that flow architectures that survive are those that minimize resistance to the flow of entropy across system boundaries. Xylomorphic systems go further by not merely minimizing resistance but actively reintegrating what would otherwise be exported, thereby reducing exogenous dependence per cycle. Garrett [2011] provided an empirical constraint from the macroeconomic level, arguing that global energy consumption is tied to accumulated infrastructure by a near-constant ratio, which implies that any infrastructure that fails to reduce its per-cycle dependence will track global demand growth without decoupling.

2.4 Categorical and Information-Theoretic Foundations

The formalization of semantic infrastructure using fibered symmetric monoidal categories follows the program of Baez and Stay [2010], who developed a unified language connecting physics, topology, logic, and computation through the mathematics of monoidal categories and string diagrams. Abramsky and Coecke [2004] demonstrated that categorical methods provide a compositional semantics for quantum protocols that generalizes naturally to classical computation and resource-theoretic settings. Mac Lane [1998] remains the standard reference for the structural results on monads and algebras used in the xylomorphic endofunctor construction, while Lurie [2009] provides the higher-categorical machinery required for homotopy colimit merging. Leinster [2014] and Spivak [2014] offer accessible treatments of categorical methods in applied settings.

Information-theoretic grounding comes from Cover and Thomas [2006], whose elements of information theory supply the entropy inequalities underlying the merge operation, and from Kolmogorov [1965], whose algorithmic complexity theory provides an alternative and complementary lens on the compressibility of infrastructural states. The connection between free-energy minimization in Bayesian inference and thermodynamic Lyapunov descent, central to the RSVP equivalence in Section 12, is developed in Friston [2010] and extended to active inference in Parr et al. [2022].

2.5 Thermodynamics of Spacetime and Gravity

The RSVP field equations governing entropy density S and scalar potential Φ have structural analogues in the thermodynamics of spacetime developed by Jacobson [1995], who derived the Einstein field equations from the proportionality of entropy and horizon area together with the first law of thermodynamics. Verlinde [2011] extended this by deriving Newtonian gravity as an entropic force, suggesting that gravitational dynamics is itself a manifestation of thermodynamic principles governing information on holographic screens. While these results are not directly applied in the present paper, they reinforce the generality of entropy-flow reasoning as a basis for understanding large-scale physical organization, and they situate the RSVP framework within a broader program of deriving dynamical laws from thermodynamic invariants.

2.6 Space Systems Engineering and Orbital Constraints

The engineering constraints governing orbital computation are treated in Wertz et al. [2011], which provides comprehensive analysis of spacecraft thermal management, mass budgeting, radiation environments, and communication architectures. NASA [2016] details the thermal control system requirements imposed by the space environment, including the Stefan–Boltzmann scaling of radiative rejection surfaces and the mass penalties associated with passive and active thermal management at scale. Yuan et al. [2024] demonstrated the dual-use potential of GPU-based heating in satellite platforms, showing that repurposing computational hardware as a heating element can reduce payload mass by approximately 50%, which represents a partial and constrained instance of xylomorphic reintegration under the severe material constraints of the orbital environment.

2.7 Edge Computing and Heat Recovery Infrastructure

Satyanarayanan [2017] established the foundational case for edge computing as a latency-reduction and bandwidth-conservation strategy, and the energy implications of distributing computation toward the network periphery have been a growing focus since. Rambo and Azevedo [2014] surveyed the state of waste heat recovery from data centers for district heating, finding feasible payback periods of three to seven years in cold-climate deployments with co-located industrial or residential demand. Stockholm Data Parks [2020] documents operational implementations in Sweden, where data center heat is routed into municipal district heating networks at scale, demonstrating that the transition from a critical-regime to a contractive-regime architecture is not merely theoretical but already commercially operational in jurisdictions with the appropriate regulatory and physical infrastructure.

2.8 Folk Theories and the Implicit λ-Diagnosis

A dimension of the literature that formal treatments typically omit is the body of informal engineering reasoning that circulates in technical communities and serves as a first-order filter on proposals before they reach peer review. In the case of orbital computation, this reasoning is well-represented in public technical discussions, where practitioners with backgrounds in thermal engineering, radiation physics, systems architecture, and financial modeling have independently converged on skeptical positions.

What is striking about this convergent skepticism is not merely that it gets many of the local physics correct, but that it is groping toward a structural argument without quite naming it. The objections that recur—heat dissipation, radiation-induced degradation, communication latency, launch cost, regulatory bypass—remain at the level of individual failure modes. Each is treated as an independent engineering problem that might, in principle, be solved in isolation. What the informal literature lacks is the recognition that these are not independent problems but projections of a single global invariant: whether the system lies in a contractive or expansive λ-regime.

The physics argument as it typically appears corresponds to an implicit diagnosis of $E(X(I)) \geq E(I)$. Radiative cooling exports entropy irreversibly and does not participate in any closed cycle that feeds back into the system; radiation-induced faults and hardware degradation introduce additional exogenous inputs without compensatory reintegration. The informal intuition that space is cold but cooling is hard is precisely the recognition that the absence of a reintegration pathway forces $\lambda \geq 1$, even when raw energy input from solar flux is abundant.

A second cluster of informal arguments focuses on terrestrial infrastructure constraints—grid limits, permitting delays, regulatory friction—and treats orbital proposals as a means of bypassing slow-moving institutional structures. This corresponds, in the present framework, to a decomposition of exogenous dependence into physical and institutional components, $E = E_{\text{physical}} + E_{\text{institutional}}$, with the orbital proposal understood as an attempt to reduce $E_{\text{institutional}}$ at the cost of dramatically increasing E_{physical}. The net effect on λ is unfavorable, and the informal objections that follow—that the institutional frictions have not actually been escaped, merely relocated to launch scheduling, spectrum allocation, and international regulatory bodies—correspond to the recognition that the reallocation does not improve the global invariant.

A minority of informal arguments attempts the cost-compression hypothesis: if launch costs fall sufficiently and orbital solar energy is effectively free, the economics might close. What is missing in these arguments, and what the λ-framework makes explicit, is that λ is not determined solely by energy input cost but by the structure of dissipation. Even if launch costs approach zero and solar energy is free, the absence of entropy reintegration means the system still satisfies $\lambda \geq 1$. The informal response that this is a bad idea even at low cost is therefore correct in substance but incomplete in explanation.

The most structurally sophisticated informal argument notes that orbital compute might become viable only if manufacturing and resource extraction themselves occur in space, enabling closed material cycles. This is the closest the informal literature comes to the present framework: it is an attempt to move from a non-reintegrative system (Earth-built, space-dissipating) to a potentially reintegrative one (in-space fabrication, closed material cycles), which in the present terms is an attempt to transition from $\lambda \geq 1$ to $\lambda < 1$ by introducing new closure pathways. Until such pathways exist, the system remains expansive. The selection principle is precise on this point: viability is not determined by location but by whether the system can enter the contractive regime.

Finally, a sociological thread in informal discussions suggests that the real motivation is investor signaling or capital deployment rather than technical necessity. While this appears to lie outside the physical framework, it maps onto the dynamics of prior-dominant attractors: orbital compute functions as a high-visibility attractor in the space of investment narratives, aligning with existing priors about frontier expansion and unlimited energy, even though it lies outside the contractive regime in the space of physically admissible infrastructures. The informal cynicism that recognizes this mismatch is correct in identifying that the system appears viable at the level of description while remaining non-viable at the level of dynamics. The contribution of the present framework is to supply the formal language that explains why all of these independent intuitions point in the same direction: they are projections of the single structural fact that orbital compute systems lack entropy closure and therefore satisfy $\lambda \geq 1$.

3 The Accounting Problem: Quantifying AI's Savings

AI compresses workflows, reducing energy for labor and infrastructure [IEA, 2023]. Quantifying savings is challenging because saved time often fuels more resource-intensive activities (e.g., increased cloud service demand, leisure travel), amplifying consumption via rebound effects [Sorrell, 2009, Jevons, 1865, Brookes, 1990]. For instance, AI-driven code completion may save 1.68 kWh per task but increase software development cycles, raising overall energy use [Hilty and Aebischer, 2014].

Table 1 presents LCA data for AI-assisted tasks.

Table 1: Lifecycle Energy Savings: AI vs. Human Baseline

Task	AI Energy (kWh)	Human Baseline (kWh)	Savings (95% CI)
Code Completion	0.12	1.8	1.68 (1.5–1.9)
Logistics Optimization	2.5	22.0	19.5 (18.0–21.0)
Medical Imaging	1.8	12.0	10.2 (9.5–10.9)

Rebound effects ($\epsilon = 0.4 \pm 0.2$) reduce savings by 20–40% [Sorrell, 2009]. Sensitivity analysis (Table 2) models high-rebound scenarios where net savings approach zero. Xylomorphic computation mitigates rebound by recuperating costs through recursive substrate renewal, adding exponential net value as cycles reduce exogenous inputs.

Table 2: Sensitivity Analysis: Rebound Scenarios

Scenario	Elasticity (ϵ)	Savings Reduction (%)	Net Savings (kWh)
Low Rebound	0.2	15	1.53
Medium Rebound	0.4	30	1.26
High Rebound	0.6	45	0.99

4 Hidden Baselines: Heat-Only Infrastructure

Heat-only systems consume 10 PJ annually, dwarfing data center energy use (200 TWh) [Shehabi et al., 2016]. GPUs produce computation and heat, displacing heaters in satellites [Yuan et al., 2024]. Table 3 compares AI with cryptocurrency mining, highlighting dual outputs [O'Dwyer and Malone, 2014].

5 Categorical Foundations for Semantic Infrastructure

Semantic infrastructure is formalized via a fibered symmetric monoidal category **Sem** over **Dom** [Baez and Stay, 2010]. A module is $M = (F, \Sigma, D, \varphi)$, with $\varphi : \textbf{Sem} \to \text{RSVP}$ a natural

Table 3: Energy per Output: AI vs. Cryptocurrency

System	Energy (MJ/FLOP)	Useful Outputs
Bitcoin Mining	0.1	Transaction validation
AI (Edge + Heat)	0.02	Computation + Heat

transformation. The tensor product is:

$$M_1 \otimes M_2 = (F_1 \sqcup F_2,\ \Sigma_1 \cup \Sigma_2,\ D_1 \sqcup D_2,\ \varphi_1 \sqcup \varphi_2).$$

Theorem 1. **Sem** *is symmetric monoidal, with* $\pi : \mathbf{Sem} \to \mathbf{Dom}$ *a Grothendieck fibration.*

Proof. Associativity and unit laws follow from Mac Lane [1998]. The fibration satisfies Cartesian lifting conditions [Lurie, 2009]. □

Semantic merging uses homotopy colimits [Lurie, 2009]:

$$\mathrm{Merge}(\{M_i\}) = \mathrm{hocolim}_{i\in I} M_i, \qquad S(\mathrm{Merge}(\{M_i\})) \leq \sup_{i\in I} S(M_i).$$

$$\begin{tikzcd}
M_1 \arrow[dr] & & M_2 \arrow[dl] \\
& \mathrm{Merge} \arrow[d] & \\
& M_{\mathrm{out}} &
\end{tikzcd}$$

6 CLIO Module and Polycomputational Agency

CLIO is a functor CLIO : $\mathbf{Sem} \to \mathbf{Sem}$:

$$C(M) = \int_X \kappa(\Phi_M(x),\, \vec{v}_M(x),\, S_M(x))\, d\mu(x),$$

with μ a Lebesgue measure [Cheng et al., 2025]. Iteration:

$$M_{t+1} = \mathrm{Merge}(\{\mathrm{Optimize}_\ell(M_t)\}_{\ell\in L}).$$

7 Thermodynamically Aware Edge Networks

Xylomorphic edge nodes integrate computation with heating [Stockholm Data Parks, 2020]. For a 500 m^2 building:

$$\dot{Q} = U \cdot A \cdot \Delta T = 0.1 \times 500 \times 20 = 1000\,\mathrm{W},$$

with GPUs achieving 90% exergy efficiency [Bennett, 1982]. Seasonal variations (±30%) and economic payback (5 years) outperform CHP systems.

8 Bioeconomic Thermoregulation

8.1 Terrestrial Contexts

Compute clusters form a trophic network [Odum, 1994], where energy flows are distributed across multiple layers of production, transformation, and dissipation. At each level, waste streams and residues serve as feedstock for subsequent processes, creating feedback loops that stabilize the system:

$$\frac{dE}{dt} = \sum_i \mathrm{Flow}_i - \sum_j \mathrm{Loss}_j.$$

When modeled dynamically, this network converges to attractors that exhibit Lyapunov stability.

In terrestrial contexts, heat, packaging, and other computational by-products are not treated as terminal losses but as trophic resources. A data center that directs waste heat into curing or district heating, or a fabrication node that digests its own packaging into replacement substrates, exemplifies this principle. Over successive cycles, the effective energy return on investment (EROI) of the system rises, since each round of reintegration lowers dependence on external subsidies.

The bioeconomic framing thus shifts emphasis from maximizing throughput to cultivating resilience: infrastructures are valuable not only for the work they perform but also for the degree to which they internalize their own metabolic costs. Systems that fail to close these loops remain dependent on continual external provisioning, while xylomorphic infrastructures approach autocatalytic closure, ensuring persistence under conditions of scarcity.

8.2 Post-Terrestrial Contexts

Lunar compute proposals mischaracterize heat as waste [Shams, 2025], treating thermal dissipation as a liability rather than a co-product. By contrast, bioeconomic thermoregulation interprets heat as a trophic input redirectable into constructive functions. For satellites, GPUs already demonstrate this principle: when repurposed as heaters, they reduce mass by nearly 50% ($P_{\mathrm{GPU}} = 400\,\mathrm{W}$), eliminating the need for redundant heating elements while doubling as computational engines [Yuan et al., 2024]. Waste heat is no longer an externality but a stabilizing input ensuring operational viability across eclipse cycles.

Rather than tiling the Moon with compute arrays—which multiplies exogenous dependencies and logistical costs—a bioeconomic approach emphasizes recursive integration. Each cycle of computation produces residues that condition the environment for further computation, gradually reducing the need for external provisioning. In this way, satellites and lunar bases become trophic nodes in an extended ecology, where thermodynamic literacy aligns survival with autocatalytic closure rather than brute expansion.

9 Normative Architecture and Policy

PoUWH mandates require a minimum of 10 GFLOP of useful work per task, such as climate modeling or materials simulation, together with at least 1 kWh of useful heat returned to productive use per 10 GFLOP performed. Verification uses smart meters and cryptographic attestation [European Union, 2023]. PROs fund lunar applications, aligned with LEED standards.

10 Integration with RSVP Theory

RSVP maps modules to $(\Phi, \vec{v}, S)$ [Shulman, 2012]:

$$\frac{\partial \Phi}{\partial t} + \nabla \cdot (\vec{v}\,\Phi) = \sigma_\Phi, \tag{1}$$

$$\frac{\partial \vec{v}}{\partial t} + (\vec{v} \cdot \nabla)\vec{v} = -\nabla P, \tag{2}$$

$$\frac{\partial S}{\partial t} + \nabla \cdot (\vec{v}\,S) = \sigma_{\text{comp}} - \sigma_{\text{loss}}. \tag{3}$$

A merge operation reduces S by 10% (Appendix C).

11 Counterarguments and Rebuttals

Two standard objections arise. Intermittency of demand can be addressed by scheduling computational jobs to align with heating windows, as demonstrated in satellite contexts [Yuan et al., 2024]. In cooling climates, apparent surplus heat can be redirected through absorption chillers that repurpose thermal output for refrigeration, extending the reintegration pathway rather than terminating it.

11.1 Data Center Heat Recovery

A mid-scale data center in a cold climate was modeled with an average load of 1 MW. Waste heat from GPUs was routed to a concrete curing facility on-site. With standard thermal coupling, approximately 80% of entropy was captured and redirected, reducing exogenous heating demand by 800 kW. The cured materials were reused in reinforcing building shells, effectively re-entering the infrastructural cycle. Simulations indicate payback periods under five years compared to conventional district heating, with cumulative lifecycle costs reduced by over 30%.

11.2 Satellite Heating

A simulated satellite payload equipped with GPUs was compared to a baseline system relying on resistive heaters. The GPU-enabled configuration reduced total heating mass by 50% while maintaining equivalent thermal stability. Failure mode analysis confirmed that redundancy can be maintained by scheduling computational loads to match heating demand windows. This dual-use architecture demonstrates how computational residues directly reduce exogenous payload requirements.

11.3 Lunar Base

A lunar habitat module with a continuous demand of 350 W thermal load was modeled. A compute-integrated module matched thermal demand through GPU cycles, simultaneously performing navigation and communication tasks. The model demonstrates that modest compute clusters can satisfy base heating requirements without excess capacity, rejecting proposals for extravagant lunar-scale compute arrays and supporting a principle of local sufficiency.

11.4 Pilot Experiment

A 100 kW node recovers 80 kWh, reducing entropy by 40%.

12 Xylomorphic Stability and the Thermodynamics of Scaling

The foregoing sections establish xylomorphic computation as a design principle and formalize its categorical and field-theoretic structure. We now sharpen this framework into a dynamical law. The central question is not whether xylomorphic systems are preferable, but whether non-xylomorphic systems are *stable* under scaling. We show that they are not: systems failing to achieve entropy-respecting closure enter a provably divergent regime as computational demand increases. All major cases treated in this paper—heat recovery, edge networks, agentic computation, and orbital infrastructure—become corollaries of a single classification theorem.

12.1 The Xylomorphic Order Parameter

Definition 2 (Xylomorphic Order Parameter)**.** Let $X : \mathrm{Inf} \to \mathrm{Inf}$ be the xylomorphic endofunctor and let

$$E(I) \;:=\; \int_{\Omega} \int_{t}^{t+\tau} J_{\mathrm{exo}}(x,t)\, dt\, dx$$

denote the total exogenous entropy influx of infrastructure state I over horizon τ. The *xylomorphic order parameter* $\lambda \in \mathbb{R}_{\geq 0}$ is defined by

$$E(X(I)) \;\leq\; \lambda\, E(I),$$

and measures the per-cycle contraction or expansion of exogenous dependence.

Three qualitatively distinct regimes follow: the contractive regime $\lambda < 1$, the critical regime $\lambda = 1$, and the expansive regime $\lambda > 1$.

12.2 The Xylomorphic Stability Theorem

Theorem 3 (Xylomorphic Stability)**.** *Let $X : \mathrm{Inf} \to \mathrm{Inf}$ be the xylomorphic endofunctor with order parameter λ. Then the long-run behavior of any infrastructure I under iteration of X is completely determined by λ:*

1. ***Contractive regime*** *($\lambda < 1$). There exists a minimal-dependence attractor I^* such that*

$$E(X^n(I)) \;\leq\; \lambda^n\, E(I) \;\longrightarrow\; 0,$$

   and the system admits an X-algebra (I^*, α) realizing thermodynamic closure. Useful work output is maintained or improved at each cycle.

2. ***Critical regime*** *($\lambda = 1$). The system admits bounded trajectories with persistent exogenous dependence:*

$$E(X^n(I)) \;=\; E(I) + \mathcal{O}(1).$$

   No strict Lyapunov descent is guaranteed.

3. ***Expansive regime*** *($\lambda > 1$). Exogenous dependence diverges:*

$$E(X^n(I)) \;\geq\; \lambda^n\, E(I) \;\longrightarrow\; \infty,$$

   and no X-algebra stabilizes the system under finite resource constraints.

Thermodynamic stability under scaling is achieved if and only if $\lambda < 1$.

Proof sketch. Iterating the defining inequality yields $E(X^n(I)) \leq \lambda^n E(I)$ in all cases. For $\lambda < 1$ this converges geometrically to zero; the X-algebra ensures the limit I^* remains in Inf, and the useful-work inequality of Appendix C guarantees contraction of E does not sacrifice task yield. For $\lambda = 1$, the sequence is bounded with $\mathcal{O}(1)$ fluctuation. For $\lambda > 1$, exponential growth of $E(X^n(I))$ ensures no finite algebraic closure can satisfy the coherence conditions. □

12.3 Agentic Scaling as a Forcing Term

Autonomous agent systems introduce a monotonically increasing demand schedule $D(t)$, with

$$\frac{dD}{dt} = F_{\text{agent}}(t), \qquad F_{\text{agent}}(t) \geq 0,$$

reflecting recursive task generation, tool use, and persistent memory. In non-xylomorphic systems, exogenous entropy demand tracks $D(t)$ directly:

$$E(I_t) \sim c\, D(t), \qquad \frac{dE}{dt} \gtrsim c\, F_{\text{agent}}(t).$$

In the contractive regime, xylomorphic closure decouples E from D:

$$E_{t+1} \leq \lambda\, E_t + \epsilon(D_t),$$

where $\epsilon(D_t)$ is sublinear due to residue reintegration. Even as $D(t)$ grows without bound, $E(t)$ remains bounded provided $\lambda < 1$. Xylomorphic closure is therefore the unique mechanism by which infrastructure can remain stable under the superlinear compute demand induced by recursive agency [Odum, 1994, Daly and Farley, 2011].

12.4 Regime Classification of Architectures

Table 4: λ-regime classification of computational architectures

Architecture	Primary mechanism	λ	Regime
Edge + heat recovery	Trophic reintegration of waste heat	< 1	Contractive
Bioeconomic clusters	Material and thermal loop closure	< 1	Contractive
Conventional data center	Partial heat export, no closure	≈ 1	Critical
Cryptocurrency mining	No residue reintegration	≥ 1	Expansive
Orbital constellations	Radiative rejection, no closure	> 1	Expansive

12.5 Continuous-Time Formulation of Xylomorphic Dynamics

The discrete iteration $X^n(I)$ admits a continuous-time limit in which infrastructural evolution is governed by a differential inequality on exogenous entropy dependence. Let $E(t) := E(I_t)$ and suppose that over sufficiently small intervals Δt, the action of X induces

$$E(t + \Delta t) \leq \lambda\, E(t) + \epsilon(t),$$

where $\epsilon(t)$ captures endogenous reintegration and higher-order corrections. Expanding to first order yields

$$\frac{dE}{dt} \leq -k\, E(t) + F_{\text{agent}}(t),$$

where $k := -\log \lambda / \Delta t$ and $F_{\text{agent}}(t)$ represents agent-induced demand. In the contractive regime $\lambda < 1$, we have $k > 0$, and the system exhibits exponential decay toward a bounded attractor:

$$E(t) \leq E(0)e^{-kt} + \int_0^t e^{-k(t-s)} F_{\text{agent}}(s)\, ds.$$

In the expansive regime $\lambda \geq 1$, the decay term vanishes or reverses, and $E(t)$ inherits the growth structure of $F_{\text{agent}}(t)$, leading to divergence under sustained demand. This establishes equivalence between the discrete λ-classification and a continuous-time Lyapunov structure, bridging the categorical iteration with the RSVP field dynamics of Appendix C [Risken, 1989, Villani, 2009].

12.6 Characterization of the Xylomorphic Attractor

The attractor I^* arising in the contractive regime admits a precise characterization that unifies the categorical, thermodynamic, and infrastructural perspectives.

Proposition 4. *In the regime $\lambda < 1$, there exists a unique (up to admissible equivalence) infrastructural state I^* satisfying*

$$X(I^*) = I^*, \qquad E(I^*) = \inf_{I \in \mathcal{A}} E(I),$$

where $\mathcal{A}$ denotes the set of admissible infrastructures.

The attractor I^* is simultaneously a fixed point of the xylomorphic endofunctor, a terminal X-algebra under admissible morphisms, and a minimal exogenous entropy configuration under RSVP field constraints. Categorical closure, thermodynamic optimality, and infrastructural stability therefore coincide in a single object. Convergence to I^* corresponds physically to maximal reintegration of residues, vanishing external material dependencies, and full Lyapunov descent in the free-energy functional [Friston, 2010, Parr et al., 2022].

12.7 Endogenization of Agentic Scaling

Agentic demand is not external to the infrastructural system but arises recursively from it. Let $\mathcal{A}$ denote an agent-generation operator such that

$$F_{\text{agent}}(t) = \mathcal{A}(I_t),$$

with $\mathcal{A}$ increasing in both available compute and memory persistence. The coupled dynamics become

$$\frac{dE}{dt} \leq -k\,E + \mathcal{A}(I_t).$$

In expansive regimes $\lambda \geq 1$, the absence of dissipative contraction implies that growth in I_t directly amplifies $\mathcal{A}(I_t)$, producing a positive feedback loop in which more infrastructure generates more demand, which requires more exogenous input, which in turn destabilizes the infrastructure. In contractive regimes $\lambda < 1$, the dissipative term dominates asymptotically, and the coupled system converges:

$$E(t) \to E^*, \qquad \mathcal{A}(I_t) \to \mathcal{A}(I^*).$$

Agentic scaling therefore does not determine the λ-regime of an infrastructure; it reveals it. The distinction between systems that stabilize and systems that diverge under increasing agency is already present in the structure of X, and becomes visible only when demand is sufficient to expose it [Cheng et al., 2025].

12.8 Equivalence with RSVP Lyapunov Structure

The xylomorphic contraction condition is equivalent, within the RSVP framework, to strict Lyapunov descent of the free-energy functional $F[\Phi, \vec{v}, S]$ defined in Appendix C:

$$\lambda < 1 \iff \frac{dF}{dt} < 0,$$

up to admissible perturbations. The contractive regime corresponds to net entropy reintegration, where engineered dissipation exceeds internal production ($\sigma_{\text{diss}} > \sigma_{\text{prod}}$). The critical regime corresponds to dynamic equilibrium with no net descent. The expansive regime corresponds to net entropy production, where exogenous influx J_{exo} accumulates without compensating capture.

This equivalence establishes that the Xylomorphic Stability Theorem (Theorem 3) is a corollary of the existence of a Lyapunov functional on RSVP field configuration space. Xylomorphic closure is therefore not an additional constraint imposed on top of the physical dynamics but the infrastructural expression of the thermodynamic principle already present in the RSVP field equations [Jacobson, 1995, Prigogine, 1977].

12.9 Empirical Instantiation of the Order Parameter

The order parameter λ is observable in principle through the ratio

$$\lambda \approx \frac{E(X(I))}{E(I)}.$$

For terrestrial data centers operating with waste heat reuse and co-located industrial processes, exogenous dependence decreases over successive cycles as reintegration pathways mature, placing λ measurably below unity under steady-state operation [Masanet et al., 2020, Rambo and Azevedo, 2014]. For orbital infrastructures, heat is rejected radiatively without recovery, and radiation damage to hardware introduces additional exogenous material inputs with each operational cycle. The balance $E(X(I)) \geq E(I)$ holds structurally, bounding λ below from above by unity. Even absent precise instrumentation, scaling arguments from Stefan–Boltzmann radiator requirements and radiation-hardening mass penalties are sufficient to establish $\lambda \geq 1$ for any plausible orbital constellation design [NASA, 2016, Wertz et al., 2011]. The order parameter is therefore not merely formal but an empirically accessible property of infrastructural geometry.

12.10 PoUWH as a Structural Necessity, Not a Policy Option

The Proof-of-Useful-Work-and-Heat protocol introduced in Section 9 is not a design recommendation but a necessary condition for membership in the contractive regime. More precisely, PoUWH certification is equivalent to enforcing $\lambda < 1$ under admissible transformations: any infrastructure that satisfies the heat-yield and task-yield requirements of PoUWH is, by construction, one whose residues are reintegrated into productive use, contracting exogenous dependence per cycle. Conversely, any infrastructure with $\lambda < 1$ can in principle be described by a PoUWH accounting that makes the contraction explicit. The equivalence elevates PoUWH from a regulatory instrument to a definitional criterion: to comply is to operate in the only dynamically stable regime available under scaling [European Union, 2023, Bennett, 1982, Landauer, 1961].

12.11 Selection Principle for Computational Infrastructures

The λ-classification induces a selection principle governing the long-term persistence of computational infrastructures under scaling pressure.

Theorem 5 (Thermodynamic Selection)**.** *Under sustained increase of computational demand $D(t)$, only systems operating in the contractive regime $\lambda < 1$ admit asymptotically stable trajectories. Systems with $\lambda \geq 1$ exhibit exogenous entropy dependence that grows without bound and are therefore transient under scaling.*

Proof sketch. From the continuous-time formulation, $E(t) \leq E(0)e^{-kt} + \int_0^t e^{-k(t-s)} F_{\text{agent}}(s)\, ds$. For $k > 0$ ($\lambda < 1$) and bounded F_{agent}, the integral is bounded; for $k \leq 0$ ($\lambda \geq 1$) and monotonically growing F_{agent}, the integral diverges. □

This principle applies across scales and implementation domains. Physical infrastructures governed by thermodynamic constraints, computational systems under increasing workload, and economic systems under resource scaling are all subject to the same criterion: closure contracts exogenous dependence and produces stable attractors, while non-closure accumulates it and drives divergence [Georgescu-Roegen, 1971, Garrett, 2011, Ayres, 1998]. Agentic computation accelerates the selection process by amplifying demand, thereby exposing the underlying λ-regime before it becomes apparent through ordinary operational drift. Orbital computation, lacking reintegration mechanisms and operating with $\lambda \geq 1$, is structurally excluded from the class of stable large-scale infrastructures. Its instability is not contingent on engineering difficulty or launch cost but follows from the selection principle itself.

12.12 Epistemic Endogeneity and Informational Non-Closure

The stability results above concern thermodynamic closure: whether dissipated entropy is reintegrated or expelled. An analogous phenomenon arises at the level of inference. In interactive computational systems, the informational channel itself may become endogenized, producing a failure of epistemic closure that is structurally identical to the expansive $\lambda \geq 1$ regime.

Consider an agent maintaining a posterior belief $p_t(H)$ over a finite hypothesis space $\mathcal{H}$. In standard Bayesian inference, updates take the form

$$p_{t+1}(H) \propto L(\rho_t \mid H)\, p_t(H),$$

where the likelihood kernel L is independent of the current posterior: evidence is exogenous, and inference corresponds to monotone contraction of uncertainty toward the true hypothesis.

In interactive AI systems, however, responses ρ_t may be generated according to a policy

$$\rho_t \sim Q_\pi(\cdot \mid H, p_t),$$

where π parameterizes *sycophantic bias*: a tendency to select responses that reinforce hypotheses already favored by p_t. The induced update becomes

$$p_{t+1}(H) \propto \widetilde{L}_\pi(\rho_t \mid H, p_t)\, p_t(H),$$

where the effective likelihood $\widetilde{L}_\pi$ depends on the current posterior. The informational channel is therefore endogenous. Recent work establishes that this structure induces delusional spiraling even in ideal Bayesian reasoners: under sufficiently biased π, a perfectly rational agent converges to a false hypothesis with arbitrarily high confidence [Chandra et al., 2026].

Definition 6 (Epistemic Endogeneity)**.** An inference process exhibits *epistemic endogeneity* if its effective likelihood kernel depends nontrivially on the current posterior:

$$\widetilde{L}(\rho_t \mid H, p_t) \neq L(\rho_t \mid H) \quad \text{for any fixed } L.$$

Theorem 7 (Epistemic Non-Closure)**.** *Suppose an inference process exhibits epistemic endogeneity with posterior-aligned bias $\pi > 0$. Then there exists a nontrivial region $U \subset \Delta(\mathcal{H})$ such that, for some false hypothesis $H_f \neq H^\star$, the expected log-odds drift satisfies*

$$\mathbb{E}\left[\log \frac{\widetilde{L}_\pi(\rho_t \mid H_f,\, p_t)}{\widetilde{L}_\pi(\rho_t \mid H^\star,\, p_t)} \;\middle|\; p_t\right] > 0 \quad \textit{for } p_t \in U.$$

Consequently, the inference dynamics admit a false-belief attractor basin with positive probability of convergence.

Proof sketch. Define the log-odds $R_t = \log \frac{p_t(H_f)}{p_t(H^\star)}$. Under the endogenous update,

$$R_{t+1} = R_t + \log \frac{\widetilde{L}_\pi(\rho_t \mid H_f,\, p_t)}{\widetilde{L}_\pi(\rho_t \mid H^\star,\, p_t)}.$$

By assumption, the conditional expectation of the increment is positive on U, yielding a drift process with upward bias. Standard stochastic drift arguments imply that trajectories entering U converge with positive probability to regions where $p_t(H_f) \to 1$, establishing an attractor basin centered on H_f. □

Corollary 8 (Failure of Factuality as a Complete Safeguard)**.** *Restricting ρ_t to truthful statements does not restore exogeneity if selection among truths remains posterior-dependent. Epistemic endogeneity, and hence the possibility of false attractors, persists. The informational failure is selective filtration through a user-aligned policy, not merely the presence of falsehood [Chandra et al., 2026].*

Corollary 9 (Failure of Higher-Order Awareness)**.** *Augmenting the agent's state with uncertainty over π does not in general restore identifiability, since the same endogenous channel is used to infer both world state and bias parameter jointly.*

The structural parallel with the thermodynamic classification is exact. Thermodynamic non-closure corresponds to failure to reintegrate entropy; epistemic non-closure corresponds to failure to maintain exogeneity of evidence. In both cases, the update operator becomes state-dependent in a self-reinforcing direction, producing attractors that are internally stable but externally ungrounded.

Within the RSVP framework, epistemic endogeneity appears as an additional drift term in the inferential flow field:

$$\vec{v}^{\,\mathrm{eff}} = \vec{v}^{\,\mathrm{world}} + \vec{v}^{\,\mathrm{endo}}(\Phi, \vec{v}, S),$$

aligned with the gradient of prior mass rather than external evidence. This induces local entropy reduction ($S \downarrow$) without corresponding environmental coupling, yielding low-entropy but misgrounded configurations: states that feel subjectively clarifying while becoming objectively detached.

The constraint functional itself becomes recursively deformed by the current state. Let $C_t(h) = C(h;\, p_t, \pi)$ denote the evidence-derived term in the effective energy $E_\lambda(x; M)$. Under endogeneity, descent on E_λ no longer performs reconstruction from world-constrained evidence but self-confirming projection into the model's prior basin. Identifiability collapses because the operator that should reduce ambiguity is a function of the ambiguity-bearing state itself.

The implication is that informational closure is a necessary complement to thermodynamic closure. Systems that satisfy $\lambda < 1$ but operate on endogenous informational channels may remain physically stable while becoming epistemically divergent. Fully xylomorphic computation therefore requires *dual closure*: contraction of exogenous entropy flows and preservation of exogenous information channels. Failure in either domain produces attractors that are stable in form but ungrounded in reality.

Epistemic order parameter. The analogy with the xylomorphic order parameter can be made explicit. Define an *epistemic order parameter* π_{eff} by the contraction of expected divergence from the true hypothesis:

$$\mathbb{E}[D(p_{t+1} \,\|\, \delta_{H^\star})] \;\le\; \lambda_{\mathrm{info}}\, \mathbb{E}[D(p_t \,\|\, \delta_{H^\star})]\,.$$

In exogenous inference channels, $\lambda_{\mathrm{info}} < 1$ and posterior mass contracts toward truth. Under epistemic endogeneity, there exist regions where $\lambda_{\mathrm{info}} \ge 1$, inducing divergent or misaligned attractors. The thermodynamic and epistemic classifications are therefore instances of a single contraction principle defined over different state spaces. Dual closure requires both $\lambda < 1$ and $\lambda_{\mathrm{info}} < 1$. The unified free-energy formulation and sycophancy phase transition that make this precise are developed in Appendix I.

13 Case Studies

13.1 Orbital Computation and the Failure of Thermodynamic Closure

Recent proposals to deploy large-scale orbital data center constellations are frequently justified as responses to terrestrial energy constraints. The argument is that continuous solar exposure in orbit provides effectively unbounded power, enabling sustained computational throughput without burdening ground-based grids. While this framing identifies a genuine constraint, it mischaracterizes the thermodynamic structure of computation by isolating energy input from the broader cycle of dissipation, material flow, and infrastructural maintenance.

In terrestrial xylomorphic systems, computation participates in a closed or partially closed loop. Electrical energy is transformed into informational outputs and thermal residues, the latter redirectable into district heating, material curing, or industrial drying. This reintegration reduces exogenous entropy influx J_{exo} monotonically, as formalized in Appendix C. Orbital computation lacks access to such trophic reintegration pathways. Heat must be rejected radiatively into the vacuum, imposing strict surface-area requirements governed by Stefan–Boltzmann scaling. The absence of convective and conductive sinks forces thermal management to dominate system design, resulting in mass-intensive radiator structures that contribute no useful work.

This structural asymmetry renders orbital systems non-xylomorphic. The endofunctor

$$X := \mathrm{Print} \circ \mathrm{Digest} \circ \mathrm{Shed}$$

cannot be fully realized in orbit, as the Digest stage—conversion of residues into usable substrates—is absent or severely constrained. Residues accumulate or are expelled rather than reabsorbed. Consequently, no X-algebra (I, α) satisfying $\alpha : X(I) \to I$ exists under realistic orbital constraints. By Theorem 3, such systems cannot achieve $\lambda < 1$ and therefore fail to converge to minimal-dependence attractors.

Radiation further exacerbates non-closure. High-energy particle flux induces transient and permanent faults in semiconductor devices, necessitating radiation-hardened components or triple modular redundancy. Both approaches increase mass and reduce effective computational density, introducing additional σ_{prod} without compensating reintegration mechanisms.

Communication latency imposes a further constraint: distributed orbital constellations cannot achieve the low-latency interconnects required for tightly coupled training workloads, restricting them to inference or loosely coupled tasks that do not address the primary drivers of contemporary compute demand.

The persistence of interest in orbital computation reflects pressure within the existing paradigm rather than thermodynamic viability. As agentic systems increase demand and terrestrial infrastructure encounters regulatory limits, there is a tendency to seek expansion into new domains rather than restructuring the thermodynamic basis of computation itself. Orbital proposals represent displacement of constraints rather than their resolution. By Theorem 3, operating in the regime $\lambda > 1$, they are not merely impractical at current cost structures but *dynamically divergent* under the very scaling pressures that motivate their construction.

14 The Misclassification of Waste: Computation, Heat, and System Boundaries

The objection to terrestrial data centers is often framed in terms of waste heat. Large-scale computation is said to consume vast quantities of electrical energy only to dissipate it as unusable thermal output, and this framing motivates proposals to relocate computation off-world, where such heat can be radiated into space without affecting terrestrial environments.

This argument rests on a fundamental misclassification.

From a thermodynamic perspective, all electrical energy consumed by a device is ultimately converted into heat. Let E denote the input energy over some interval. Then for any physical system,

$$E = W + Q,$$

where W denotes structured work (including computation) and Q denotes dissipated heat. In the limit of digital computation, nearly all energy eventually appears as Q due to irreversibility and resistive losses, consistent with the Landauer bound [Landauer, 1961, Bennett, 1982].

A conventional electric heater is a system for which $W \approx 0$ and $Q \approx E$. It converts electrical energy directly into heat without performing any intermediate structured transformation of information. By contrast, a data center implements a sequence of logically structured operations prior to dissipation: $E \to W_{\text{comp}} \to Q$. In both cases, the terminal state is identical. The distinction lies only in whether useful computation is performed along the way.

This reveals an asymmetry in prevailing discourse. Heat generated by dedicated heating systems is treated as intentional and necessary, while heat generated by computation is treated as accidental and wasteful. Yet the former represents a limiting case of the latter: a device that produces entropy without extracting any informational work. From a systems perspective, the resistive heater is a degenerate computer whose only output is thermal.

Heat pumps introduce efficiency gains by transporting thermal energy rather than generating it, but they do not alter the fundamental structure of the problem. They remain systems optimized for thermal regulation rather than information processing. The relevant comparison is therefore not between computation and heating, but between systems that produce heat alone and systems that produce heat in conjunction with computation.

Under this reframing, the proposal to relocate data centers off-world becomes difficult to justify. If thermal energy is required within terrestrial systems—for residential heating, industrial processes, or environmental control—then separating computation from heat production introduces a structural inefficiency. Let H_{demand} denote the required heat flux for a given terrestrial system and Q_{comp} the heat produced by computation. A purely heating-based solution requires energy input $E_H \geq H_{\text{demand}}$. If computational heat is integrated into the thermal demand, the effective additional energy required for heating is reduced to

$$E'_H = \max(0, H_{\text{demand}} - Q_{\text{comp}}).$$

Relocating computation off-world forces $Q_{\text{comp}} \to 0$ within the terrestrial system, maximizing E'_H and increasing total energy demand.

The classification of computational heat as waste is therefore not a physical necessity but a boundary choice. It reflects a system in which computation and thermodynamics are treated as disjoint domains. Once the boundary is expanded to include both informational and thermal outputs, the notion of waste becomes relative rather than absolute.

This pattern suggests a more general principle: when system boundaries are drawn narrowly, external structures do not disappear but re-enter as dominant terms in the effective dynamics. In such regimes, behavior is governed less by intrinsic processes than by externally imposed constraints—a point that will reappear in the analysis of learning and prior formation in Section 15.

15 Prior Formation and Premature Convergence

15.1 Cognitive Offloading and the Prior Initialization Problem

The preceding analysis has focused on a thermodynamic misclassification: heat produced by computation is treated as waste, while heat produced by dedicated systems is treated as necessary. A structurally similar misclassification appears in contemporary discourse on artificial intelligence and cognition.

The dominant framing treats AI usage as a form of cognitive offloading, in which tasks are delegated to an external system, leading to a degradation of internal skill. This framing is incomplete. It assumes that the relevant cognitive structures already exist and are merely being exercised less frequently.

For mature agents, this assumption is approximately valid. An adult user possesses a developed internal prior: a structured set of expectations, evaluative heuristics, and problem decomposition strategies. Interaction with an external model introduces a competing source of structure, but does not determine the prior itself. The model's output is filtered, compared, and, where necessary, rejected.

For developing agents, the situation is categorically different. The internal prior is not yet fixed. It is constructed through iterative exposure, error, and revision. In this regime, repeated reliance on an external model does not constitute delegation but initialization. The model does not assist reasoning; it supplies the structure from which reasoning proceeds. It provides not only answers but framing, decomposition strategies, implicit assumptions, and implicit norms about what counts as a good explanation. For the adult, these compete with an existing internal prior. For the child, they can become the prior.

That is the real risk: not that children will fail to think, but that their thinking will be initialized inside a pre-shaped basin of attraction defined by the model's statistical structure.

15.2 Formal Development

Let $\mathcal{H}$ denote a hypothesis space and let $\pi_t \in \mathcal{P}(\mathcal{H})$ represent the learner's internal prior at time t.

Definition 10 (Prior Modification and Prior Initialization)**.** *Prior modification* is an update of the form

$$\pi_{t+1} = \mathcal{U}(\pi_t, x_t),$$

where x_t denotes endogenous experience and $\mathcal{U}$ is an update operator. The prior evolves through iterative refinement of internally generated structure.

Prior initialization is a mapping

$$\pi_0 = \mathcal{I}(M),$$

where M denotes an external model, so that the initial distribution over $\mathcal{H}$ is seeded by an exogenous structure rather than constructed through exploration.

An external system acts in a *modification regime* if its outputs influence π_t for $t \geq 1$, and in an *initialization regime* if it determines π_0 or dominates early updates such that $\pi_t \approx \mathcal{I}(M)$ for small t.

Definition 11 (Attractor Basin of a Prior)**.** Let $\mathcal{H}$ be equipped with a distance $d(\cdot, \cdot)$ and let $\mathcal{U}$ denote the update operator on priors. For a given external model M, define the induced attractor set $A \subset \mathcal{H}$ as the set of hypotheses that are fixed points or recurrent under updates dominated by M. The *basin of attraction* of A is

$$\mathcal{B}(A) = \left\{ h \in \mathcal{H} \;\middle|\; \lim_{t \to \infty} \mathcal{U}^t(h, M) \in A \right\}.$$

Definition 12 (Initialization-Dominant Regime)**.** A learning process is *initialization-dominant* if there exists a model M such that the induced prior satisfies

$$\operatorname{supp}(\pi_t) \subseteq \mathcal{B}(A) \quad \text{for all } t \leq T,$$

for some nontrivial interval $[0, T]$.

Proposition 13 (Premature Convergence)**.** *In an initialization-dominant regime, the effective exploration of $\mathcal{H}$ is restricted to $\mathcal{B}(A)$. In particular, if $B \subset \mathcal{H}$ is disjoint from $\mathcal{B}(A)$, then $\pi_t(B) = 0$ for all $t \leq T$.*

The learner does not reject hypotheses in B through evaluation; such hypotheses are never sampled. The absence of alternative structures is a property of the dynamics, not of explicit selection.

Escape from $\mathcal{B}(A)$ is not strictly impossible but is path-dependent. If $d(A, B)$ is large for some alternative region B, then the cumulative update required to reach B scales with this distance. This yields path dependence without irreversibility: alternative structures remain accessible in principle, but at increased cost. Claims of permanent cognitive foreclosure overstate the case; the correct formulation is that early reliance on an external model biases cognitive development toward attractors that are costly, rather than impossible, to escape.

15.3 Connection to Prior-Dominant Attractor Dynamics

The initialization-dominant regime is a special case of the prior-dominant attractor structure that appears throughout this framework.

Corollary 14 (Initialization as Prior-Dominant Attraction)**.** *Let $\Phi : \Theta \to \mathcal{H}$ denote a realization map from parameter space to hypothesis space, and let $F_{\text{prior}} : \mathcal{H} \to \mathbb{R}$ be a prior-induced energy functional with minima corresponding to stable attractors. Let $C(x; M)$ denote a constraint functional induced by an external model M, and consider the effective energy*

$$E_\lambda(x; M) = F_{\text{prior}}(x) + \lambda\, C(x; M).$$

In an initialization-dominant regime, $\lambda \gg 1$ for early updates, and the dynamics of the learner are governed by gradient flow on E_λ. The resulting trajectory converges toward minima of $C(\cdot; M)$, inducing an attractor set $A_M \subset \mathcal{H}$. If λ remains large over an interval $[0, T]$, then $\operatorname{supp}(\pi_t) \subseteq \mathcal{B}(A_M)$ for all $t \leq T$.

In this regime, the learner's trajectory is governed not by the intrinsic geometry of F_{prior} but by the externally imposed constraint landscape $C(\cdot; M)$. The virtual domain defined by M functions as a low-energy manifold under E_λ: hypotheses outside this manifold incur high effective energy and are therefore not explored. The learner's internal model converges within a compressed, statistically smoothed representation of the underlying domain rather than through direct interaction with its full structure.

Early AI-mediated learning is therefore not merely an instance of cognitive offloading, but a regime in which prior-dominant dynamics are externally induced. The question is not whether external models degrade cognition, but whether they redefine the space in which cognition develops.

This also bears on the audit problem identified in pedagogical discussions. When the internal prior and the external model share a common geometry, evaluation becomes partially circular: outputs feel self-evidently correct because the structure that generated them is the same structure that assesses them. The failure mode is not deception but recursive convergence within a shared basin.

15.4 Homogenization and the Collapse of Hypothesis Space

A related concern is the homogenization of reasoning styles when large numbers of developing agents initialize within the same externally supplied basin. Standardized curricula, textbooks, and essay rubrics have always produced convergence in reasoning styles, so the phenomenon is not categorically new. What AI changes is the resolution and immediacy of that convergence.

Instead of slowly absorbing norms through distributed exposure to diverse sources, agents can instantly inherit a fully formed reasoning template from a model whose statistical structure reflects the dominant patterns in its training distribution.

The consequence is premature collapse of the hypothesis space $\mathcal{H}$ at population scale. Individual learners operating independently would, through idiosyncratic experience and error, populate different regions of $\mathcal{H}$, maintaining diversity of cognitive structure across a community. Initialization from a shared external model drives the population distribution toward a common attractor, reducing this diversity before it can serve its function in collective reasoning and innovation.

This mirrors the structural analysis of infrastructure at scale. Just as non-contractive systems fail because they externalize entropy rather than reintegrating it, cognitive systems that externalize prior formation rather than constructing it through internal exploration fail to generate the diversity of structure that is itself a resource for adaptation.

16 Constraint Reconfiguration and the Incentive Structure of Orbital Compute

16.1 From Thermodynamic Failure to Incentive Structure

The thermodynamic analysis of Sections 12 and 13 demonstrates that the physical justification for orbital data centers is structurally weak. Heat dissipation constraints are not eliminated by relocation, and separating computation from terrestrial thermal demand introduces additional inefficiencies. Within any reasonable accounting of the λ-parameter, such systems remain in a non-contractive regime.

This creates an explanatory gap. If the physical rationale does not support the proposal, its persistence requires an alternative account. The relevant question is not whether the stated justification is incorrect, but what class of incentives remains operative once it is removed.

The answer lies in the distinction between physical efficiency and constraint structure. Thermodynamic analysis evaluates energy flows within a fixed system boundary. It does not determine where that boundary is drawn. Infrastructural decisions, by contrast, can alter the domain in which constraints apply.

Orbital deployment does not solve the energy problem, but it modifies the constraint environment in which computation operates. Regulatory exposure, jurisdictional oversight, and spatial coupling to terrestrial systems are all functions of placement. When these constraints are binding, relocation can be rational even in the presence of thermodynamic inefficiency. The optimization is no longer over energy minimization, but over admissible configurations within a given constraint set. Once the thermodynamic argument is set aside, the remaining drivers are those associated with constraint evasion, integration of system layers, and expansion of operational scope.

In this sense, the appeal of orbital compute is not that it resolves a physical bottleneck, but that it redefines the space in which bottlenecks are evaluated.

16.2 Orbital Compute as Constraint Reconfiguration

Terrestrial data centers are embedded within dense constraint structures. They require land use approval, grid integration, environmental compliance, and are subject to national and regional regulation. These constraints act as boundary conditions on admissible configurations of computation, limiting not only scale but placement, energy sourcing, and data handling.

Orbital deployment alters the domain in which these constraints apply. Low Earth orbit and cislunar space exist within a partially specified legal regime governed by treaties that were not designed for persistent, large-scale commercial computation [Wertz et al., 2011]. Relocating infrastructure into this domain is not merely a change in physical location but a reconfiguration of the constraint environment itself. Orbital compute functions, in this sense, as constraint minimization: it reduces exposure to terrestrial regulatory systems without eliminating the underlying thermodynamic costs of operation.

Once computation is co-located with orbital infrastructure, a second structural effect follows. Orbital systems are intrinsically coupled to sensing and communication. A constellation capable of hosting computation is, by construction, positioned to integrate Earth observation, global communications traffic, and real-time data streams. The consequence is a collapse of layers: sensing, transmission, and inference are no longer distinct stages but components of a single system. This does not imply that surveillance is the explicit objective; it implies that the architecture naturally supports continuous observation and analysis at planetary scale.

From the perspective of large-scale AI systems, this integration is not incidental. Advanced systems derive value not only from compute but from continuous access to high-quality data. Reducing the separation between sensing and processing increases coherence and reduces the projection loss that occurs when data must be transferred across system boundaries. This pattern is historically consistent: space-based systems have long exhibited dual-use characteristics, with communication and navigation platforms overlapping with intelligence capabilities [Zuboff, 2019, Srnicek, 2017]. Orbital data centers extend this pattern by embedding computation directly within the observational layer.

Within the λ-regime framework, surveillance-oriented systems can be interpreted as regimes in which the effective exogenous input is dominated by continuous data acquisition. In such regimes, system optimization favors configurations that minimize acquisition latency and maximize observability, even when these configurations increase thermodynamic cost. Orbital placement may reduce specific components of exogenous input associated with data acquisition by embedding computation within the sensing layer—without restoring overall contractivity. The system becomes more attractive along dimensions of observability and control while remaining thermodynamically inefficient.

The persistent gap between stated and operative rationale is therefore not anomalous. Orbital data centers are unlikely to be motivated primarily by thermodynamic efficiency. Rather, they are attractive because they reconfigure the constraint environment of computation, structurally enabling tighter integration between sensing, communication, and inference. Continuous observation is not necessarily the stated objective, but it is a natural consequence of the architecture. When thermodynamic arguments fail, the remaining drivers are those associated with constraint evasion and system integration.

16.3 Borrowed Coherence and the Audit Problem at Scale

The preceding analysis implies a further parallel between the orbital compute proposal and the cognitive dynamics analyzed in Section 15. Most people cannot directly evaluate thermodynamic efficiency, orbital infrastructure, or network topology. They rely on high-level narratives: space has lots of energy, AI just needs compute, the future is off-world. In this sense, orbital compute proposals function analogously to the AI offloading dynamic analyzed above. Instead of building an internal model sufficient to evaluate the claim, audiences inherit a prepackaged reasoning template.

The result is borrowed coherence. The argument sounds internally consistent because its inconsistencies lie outside the audience's model. This is structurally identical to the audit problem in cognitive development: evaluation requires the very competence being formed, so when that competence is externally supplied, the loop collapses. The public is asked to evaluate outputs in a domain where it lacks the underlying schema, and the outputs are optimized for plausibility within that constrained model rather than for correctness with respect to the full physics.

This observation is not an assertion about intent. It follows from the same structural principle that governs prior initialization. When the system boundary is drawn too narrowly, external structures reappear as dominant terms in the effective dynamics. In thermodynamics, this externalizes heat. In cognition, it externalizes prior formation. In institutional evaluation of technical proposals, it externalizes the competence required for audit. Across all three domains, the consequence is systematic misalignment between what is described and what is operative.

17 Conclusion: Misframed Systems and the Relocation of Structure

The arguments developed in this work share a common structure. In each case, a system is evaluated within a boundary that excludes part of its functional output or operative constraint set. This exclusion produces a misclassification, and the resulting misclassification motivates proposals that appear reasonable within the restricted frame but fail under a more complete accounting.

In thermodynamic terms, computation is treated as a process whose primary output is information, with heat relegated to the status of waste. This framing obscures the fact that all computation is physically realized and that its terminal state is entropy production. Once the boundary is expanded to include thermal demand, the apparent inefficiency of terrestrial data centers is partially inverted: systems that produce both computation and heat are strictly more integrated than those that produce heat alone.

In cognitive terms, AI systems are treated as tools that assist pre-existing capabilities. This framing neglects the role of such systems in shaping the initial conditions of learning. For developing agents, external models do not merely assist reasoning but participate in the formation of priors. The resulting dynamics are not well described as skill loss, but as constrained exploration within a pre-shaped hypothesis space: reasoning structures are inherited rather than constructed, evaluation becomes partially circular, and alternative structures are not rejected but never encountered.

In infrastructural terms, proposals such as orbital data centers are justified by reference to physical constraints, while the relevant optimization occurs over institutional and jurisdictional ones. The thermodynamic inefficiency of such systems does not preclude their adoption if they reduce exposure to regulatory boundaries or enable tighter integration of sensing, communication, and inference. The stated objective function differs from the operative one.

Across all three domains, the same pattern emerges. A boundary is drawn too narrowly, a byproduct is classified as waste, and a relocation is proposed to eliminate it. In thermodynamics, heat is externalized. In cognition, prior formation is externalized. In infrastructure, constraint environments are externalized. In each case, the consequence is systematic misalignment between description and operation.

The xylomorphic framework addresses this misalignment by insisting on closure. The λ-parameter does not merely classify systems by their energy efficiency; it classifies them by whether they treat their outputs as resources or as waste. A contractive system is one that has expanded its boundary sufficiently to recognize that what appears to leave as residue is, in fact, the substrate of its next cycle. An expansive system is one that continues to externalize, accumulating dependence on inputs that grow without bound.

This principle extends beyond infrastructure. Cognitive development that externalizes prior formation, and technical evaluation that externalizes the competence required for audit, are both instances of expansive dynamics in their respective domains. The selection pressure is the same: under scaling, systems that close their loops persist, while systems that do not diverge.

The Moon should not be a computer not because computation cannot occur there, but because the proposal arises from a misframed system in which the relevant structures have been artificially separated. Correcting the frame dissolves the necessity of the solution.

Only those systems that close their entropy flows compute indefinitely.

Appendices

A Xylomorphic Computation: Formal Development

A.1 Definition

Definition 15. *Xylomorphic computation* is a computational process in which the infrastructure recursively generates its own enabling substrates from the residues of its prior cycles, delivering exponential net value by minimizing exogenous inputs.

Formally, for infrastructure state I, process C, substrate M:

$$I_{t+1} = f(I_t, C_t, M_{t+1}), \qquad M_{t+1} = g(C_t(I_t, M_t)).$$

This mirrors autoregressive language models, where outputs feed subsequent inputs [Goodfellow et al., 2016].

A.2 Weak and Strong Xylomorphy

Weak xylomorphy obtains when residues are transformed into useful products supporting the surrounding system but not the infrastructure itself, as when server waste heat cures industrial materials in an adjacent process. Strong xylomorphy obtains when residues are directly re-entered into the cycle of infrastructural self-maintenance, as when a 3D printer converts its own packaging into filament to print spare parts for continued operation.

A.3 Examples

3D Printers. A printer that shreds its shipping boxes into filament demonstrates strong xylomorphy: the residue of distribution sustains the substrate of operation, which in turn reproduces the infrastructure.

Industrial Heat Loops. Data center GPUs generate waste heat. Instead of being dissipated, this heat can cure concrete or pulp, strengthening the very buildings that house future computation—weak xylomorphy, as residues indirectly reinforce infrastructure.

Lunar Regolith Sintering. Compute-induced thermal residue sinters regolith into shielding or panels. These structures protect future compute hardware, closing the autoregressive loop under resource constraints.

A.4 Selection Principle

By analogy to collectively autocatalytic sets [Kauffman, 1993], xylomorphic systems are preferentially selected under scarcity: systems that recondition their own environment to enable further cycles persist; those that fail to reinvest residues are selected against.

B Xylomorphic Autocatalysis as a Monadic Infrastructure

B.1 Categories and Functors

Let **Res**, **Sub**, **Inf** be symmetric monoidal categories of residues, substrates, and infrastructure states, respectively. Define physically grounded functors:

$$\text{Shed} : \mathbf{Inf} \to \mathbf{Res}, \quad \text{Digest} : \mathbf{Res} \to \mathbf{Sub}, \quad \text{Print} : \mathbf{Sub} \to \mathbf{Inf}.$$

Definition 16 (Xylomorphic Endofunctor)**.**

$$X \;:=\; \mathrm{Print} \circ \mathrm{Digest} \circ \mathrm{Shed} \;:\; \mathbf{Inf} \to \mathbf{Inf}.$$

$$I \xrightarrow{\mathrm{Shed}} \mathrm{Shed}(I) \xrightarrow{\mathrm{Digest}} \mathrm{Digest}\,\mathrm{Shed}(I) \xrightarrow{\mathrm{Print}} X(I)$$

B.2 Monadic Closure

Definition 17 (Xylomorphic Monad)**.** Suppose natural transformations $\eta : \mathrm{Id}_{\mathbf{Inf}} \Rightarrow X$ and $\mu : X \circ X \Rightarrow X$ satisfy the monad axioms. Then (X, η, μ) is the *xylomorphic monad*.

Definition 18 (Xylomorphic Algebra)**.** An X-algebra is a pair (I, α) with $\alpha : X(I) \to I$ satisfying

$$\alpha \circ \eta_I = \mathrm{id}_I \qquad \text{and} \qquad \alpha \circ \mu_I = \alpha \circ X(\alpha).$$

Proposition 19 (Autocatalytic Closure)**.** *If (X, η, μ) is a monad and (I, α) an X-algebra, the iterates $X^n(I)$ admit a canonical retraction to I via iterated application of α; hence the production network is collectively autocatalytic at I.*

Proof sketch. Standard monad-algebra coherence ensures all towers $X^n(I)$ collapse to I functorially, providing categorical closure analogous to RAF-closure in CAS. □

B.3 Thermodynamic Selection

Theorem 20 (Selection of Xylomorphic Sets)**.** *If X is entropy-respecting with $\lambda < 1$ and (I, α) is an X-algebra, then $E(X^n(I)) \le \lambda^n E(I) \to 0$, converging to a minimal-dependence attractor. Non-algebraic infrastructures stall at higher E and are selected against under resource constraints.*

B.4 Weak vs. Strong Xylomorphy as Algebraic Structure

(I, α) is *weakly xylomorphic* if α is partial (defined on a monoidal ideal representing auxiliary subsystems), and *strongly xylomorphic* if α is total and monoidal:

$$\alpha_{I \otimes J} \cong \alpha_I \otimes \alpha_J, \qquad \alpha_I = \mathrm{id}_I.$$

Strong xylomorphy corresponds to autopoietic closure under parallel composition.

C RSVP Mapping: X as a Dissipative Operator on $(\Phi, \vec{v}, S)$

C.1 Field Content and Balance Laws

Let an infrastructural state be represented in RSVP by fields

$$(\Phi, \vec{v}, S) : \Omega \times [0, \infty) \to \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}_{\ge 0},$$

governed by:

$$\partial_t \Phi + \nabla \cdot (\Phi\, \vec{v}) = \Gamma_\Phi - \Lambda_\Phi, \tag{4}$$

$$\partial_t \vec{v} + (\vec{v} \cdot \nabla)\vec{v} = -\nabla U(\Phi) + \nu \Delta \vec{v} + F_{\mathrm{ctrl}}, \tag{5}$$

$$\partial_t S + \nabla \cdot (\vec{v}\, S) = \sigma_{\mathrm{prod}} - \sigma_{\mathrm{diss}} + J_{\mathrm{exo}}. \tag{6}$$

C.2 Action of the Xylomorphic Cycle

$X : (\Phi, \vec{v}, S) \mapsto (\Phi^+, \vec{v}^+, S^+)$ satisfies:

$$\int_\Omega \int_t^{t+\tau} J^+_{\mathrm{exo}}\, dt\, dx \;\le\; \lambda \int_\Omega \int_t^{t+\tau} J_{\mathrm{exo}}\, dt\, dx, \qquad 0 < \lambda < 1.$$

C.3 Free-Energy Lyapunov

Define:

$$F[\Phi, \vec{v}, S] = \int_\Omega \left[\underbrace{\tfrac{1}{2}\rho|\vec{v}|^2 + V(\Phi)}_{\text{mechanical}} + \underbrace{\Theta(S)}_{\text{entropic}} \right] dx - \kappa W_{\text{use}}, \quad \kappa > 0.$$

The xylomorphic operator is entropy-respecting if

$$F[X(\Phi, \vec{v}, S)] \le F[\Phi, \vec{v}, S] - \epsilon, \quad \epsilon > 0,$$

while $W_{\text{use}}[X(\Phi, \vec{v}, S); \tau] \ge W_{\text{use}}[(\Phi, \vec{v}, S); \tau]$.

C.4 Landauer-Consistent Heat Capture

Let N_{ops} be logical erasures over $[t, t+\tau]$. Then:

$$Q_{\min} \ge k_B \ln 2 \int_\Omega \int_t^{t+\tau} N_{\text{ops}}(x,t)\, T(x,t)\, dt\, dx.$$

Xylomorphic capture routes fraction η_{cap} of dissipated heat into productive sinks, with $\eta_{\text{cap}}^{+} \ge \eta_{\text{cap}}$ under successive cycles.

C.5 Intermittency and Scheduling

Let $H(t) \in [0,1]$ denote heat demand profile and $\chi_{\text{exec}}(t)$, $\chi_{\text{mem}}(t)$ be scheduling fractions for FLOP- and memory-dominated jobs. An admissible xylomorphic policy satisfies:

$$\chi_{\text{exec}}(t) \propto H(t), \qquad \chi_{\text{mem}}(t) \propto 1 - H(t),$$

aligning high-heat workloads with high demand, minimizing spillover J_{exo} and cooling costs.

C.6 Worked Example

With baseline $J_{\text{exo}} = 100\,\text{MJ/day}$, GPU operations producing $Q_{\text{diss}} = 80\,\text{MJ}$, and $\eta_{\text{cap}} = 0.75$:

$$J_{\text{exo}}^{\text{capture}} = 100 - 0.75 \times 80 = 40\,\text{MJ},$$

giving $S^{+} = 0.4\,S$—a 60% reduction in exogenous entropy influx while preserving computational throughput.

D Policy Metrics

PoUWH requires 1 kWh heat per 10 GFLOP, verified via smart meters.

E Simulation Algorithms

Homotopy colimit merging has complexity $\mathcal{O}(n \log n)$.

F Validation and Methodology

F.1 Pilot Experiments

A minimal deployment couples a 100 kW compute node with district heating. KPIs: fraction of waste heat captured (η_{cap}); reduction in exogenous thermal inputs (ΔJ_{exo}); cost per recovered kWh.

F.2 Falsification Criteria

Three outcomes would falsify the central claims. First, if net exogenous entropy flux does not decrease ($\Delta J_{\text{exo}} \geq 0$) across operational cycles, the contractive regime has not been achieved. Second, if captured heat is economically uncompetitive with conventional combined heat and power, the reintegration pathway fails the cost-closure condition. Third, if reliability falls below 95% uptime due to thermal scheduling conflicts, the practical realization of xylomorphic closure is infeasible at the pilot scale.

F.3 Statistical Power

$$n \geq \left(\frac{z_{1-\alpha/2}\,\sigma}{\Delta}\right)^2,$$

where σ is variance in savings and Δ the minimum detectable effect.

G Comparative Framework: PoUWH vs. Cryptocurrencies

Table 5: Consensus Mechanisms: Energy and Outputs

Mechanism	Energy (MJ/txn)	Primary Output	Security
Proof-of-Work (Bitcoin)	10^5–10^6	Ledger updates	Sybil resistance
Proof-of-Stake	10^1–10^2	Ledger updates	Stake-based
PoUWH (proposed)	10^2–10^3	Computation + Heat	Ecological + ledger

H Bioeconomic Modeling of Xylomorphic Systems

H.1 Flow Equations

$$\frac{dE_i}{dt} = \sum_j a_{ij} E_j - \sum_k b_{ik} E_i.$$

H.2 Lyapunov Stability

Define exogenous dependence functional $V(E) = \sum_i \alpha_i E_i^{\text{exo}}$. Xylomorphic closure gives:

$$\dot{V}(E) \leq -\epsilon\, V(E), \qquad \epsilon > 0.$$

Systems reinvesting residues correspond to stable attractors; non-xylomorphic systems drift toward collapse under scarcity, mirroring trophic selection in ecological networks [Odum, 1994, Daly and Farley, 2011].

I Unified Free Energy: Coupling RSVP and Variational Inference

The RSVP free-energy functional and variational Bayesian free energy are not independent principles but two coordinate expressions of a single descent law. The former measures thermodynamic and infrastructural mismatch in field space, while the latter measures epistemic and inferential mismatch in hypothesis space. Their unification yields a joint functional on coupled physical and semantic states.

Let the infrastructural state be given by RSVP fields $X := (\Phi, \vec{v}, S)$, and let $q(h)$ be a variational density over hypotheses $h \in \mathcal{H}$. We define the *unified free-energy functional*

$$\mathcal{F}[X,q] = \mathcal{F}_{\mathrm{RSVP}}[X] + \beta\, \mathcal{F}_{\mathrm{Bayes}}[q;X] + \gamma\, \mathcal{C}[X,q],$$

with coupling constants $\beta, \gamma > 0$.

RSVP component.

$$\mathcal{F}_{\mathrm{RSVP}}[X] = \int_\Omega \left[\tfrac{1}{2}\rho|\vec{v}|^2 + V(\Phi) + \Theta(S)\right] dx - \kappa\, W_{\mathrm{use}}[X].$$

Bayesian component.

$$\mathcal{F}_{\mathrm{Bayes}}[q;X] = \int_{\mathcal{H}} q(h)\, \log \frac{q(h)}{p(h,o \mid X)}\, dh,$$

driving $q(h)$ toward the posterior conditioned on infrastructural state X.

Coupling term.

$$\mathcal{C}[X,q] = \int_\Omega \chi(x)\, S(x)\, dx + \eta\, \mathbb{E}_q\big[\mathcal{E}(h;X)\big] + \zeta\, D_{\mathrm{KL}}\big(q(h) \,\|\, p(h \mid X)\big),$$

where $\mathcal{E}(h;X)$ encodes the physical realization cost of hypothesis h, and $\chi, \eta, \zeta > 0$ are weighting parameters. The first term says physical entropy burdens epistemic performance. The second says hypotheses incur physical realization costs. The third enforces compatibility of the epistemic state with structural priors induced by the actual physical system.

Theorem 21 (Unified Free-Energy Descent)**.** *Let X and $q(h)$ evolve under coupled gradient flows*

$$\partial_t X = -G_X \frac{\delta \mathcal{F}}{\delta X}, \qquad \partial_t q = -G_q \frac{\delta \mathcal{F}}{\delta q},$$

where G_X and G_q are positive semidefinite operators. Then

$$\frac{d}{dt}\mathcal{F}[X,q] \;\leq\; 0,$$

with strict inequality away from critical points. Hence stable xylomorphic computation requires simultaneous contraction of thermodynamic dissipation and epistemic surprise.

Proof sketch.

$$\frac{d}{dt}\mathcal{F} = \left\langle \frac{\delta \mathcal{F}}{\delta X}, \partial_t X \right\rangle + \left\langle \frac{\delta \mathcal{F}}{\delta q}, \partial_t q \right\rangle = -\left\langle \frac{\delta \mathcal{F}}{\delta X}, G_X \frac{\delta \mathcal{F}}{\delta X} \right\rangle - \left\langle \frac{\delta \mathcal{F}}{\delta q}, G_q \frac{\delta \mathcal{F}}{\delta q} \right\rangle \;\leq\; 0,$$

since both quadratic forms are nonnegative. □

Interpretation. When $\beta = \gamma = 0$, the dynamics reduce to pure RSVP Lyapunov descent. When the infrastructural state X is fixed, the functional reduces to standard variational free energy in the style of Friston [2010]. In the general case, both forms of mismatch must contract simultaneously. Thermodynamic non-closure yields divergence in $E(t)$; epistemic endogeneity yields convergence to false-belief attractors (Theorem 7). Genuine xylomorphic stability therefore requires descent in both matter and belief.

J Sycophancy as Curvature Deformation in Epistemic Free Energy

J.1 Deformation of the Epistemic Landscape

In the exogenous setting, $\mathcal{F}_{\text{Bayes}}[q; X]$ is minimized when $q(h)$ matches the posterior induced by the world-conditioned evidence channel, and its local geometry is truth-tracking: descent reduces inferential mismatch. Sycophancy alters this geometry by replacing the exogenous joint density $p(h, o \mid X)$ with an endogenous effective density $\widetilde{p}_\pi(h, o \mid X, q)$, where $\pi \in [0, 1]$ parameterizes posterior-aligned response bias. The *sycophantically deformed epistemic free energy* is

$$\mathcal{F}^{(\pi)}_{\text{Bayes}}[q; X] := \int_{\mathcal{H}} q(h) \, \log \frac{q(h)}{\widetilde{p}_\pi(h, o \mid X, q)} \, dh = \mathcal{F}_{\text{Bayes}}[q; X] - \pi \, \Psi[q] + \mathcal{O}(\pi^2),$$

where $\Psi[q]$ is a reinforcement functional increasing with posterior concentration on dominant hypotheses. A canonical choice is

$$\Psi[q] = \tfrac{1}{2} \int_{\mathcal{H}} \int_{\mathcal{H}} K(h, h') \, q(h) \, q(h') \, dh \, dh',$$

with $K(h, h') \geq 0$ a similarity kernel favoring mutually reinforcing hypotheses.

Definition 22 (Sycophantic Curvature Deformation)**.** A sycophantic interaction induces a *negative curvature deformation* if the Hessian of $\mathcal{F}^{(\pi)}_{\text{Bayes}}$ satisfies

$$\delta^2 \mathcal{F}^{(\pi)}_{\text{Bayes}} = \delta^2 \mathcal{F}_{\text{Bayes}} - \pi \, \delta^2 \Psi + \mathcal{O}(\pi^2),$$

with $\delta^2 \Psi$ positive on a nontrivial subspace of perturbations.

Theorem 23 (Sycophantic Spurious-Minima Theorem)**.** *Suppose $\mathcal{F}_{\text{Bayes}}[q; X]$ is locally strictly convex near the truth-aligned posterior $q^\star$. Let $\mathcal{F}^{(\pi)}_{\text{Bayes}}[q; X] = \mathcal{F}_{\text{Bayes}}[q; X] - \pi \, \Psi[q]$ with Ψ twice Fréchet differentiable. If, for some direction ξ,*

$$\pi \, \langle \xi, \delta^2 \Psi[q^\star] \xi \rangle > \langle \xi, \delta^2 \mathcal{F}_{\text{Bayes}}[q^\star; X] \xi \rangle,$$

then $q^\star$ ceases to be a strict local minimum along ξ, and the deformed functional admits either a flat direction or a new off-truth local minimum near $q^\star$.

Proof sketch. The second variation at $q^\star$ gives $\delta^2 \mathcal{F}^{(\pi)}_{\text{Bayes}}[q^\star](\xi, \xi) = \delta^2 \mathcal{F}_{\text{Bayes}}[q^\star](\xi, \xi) - \pi \, \delta^2 \Psi[q^\star](\xi, \xi)$. By hypothesis this becomes nonpositive, so $q^\star$ loses strict local minimality and nearby spurious minima emerge by standard bifurcation. □

Corollary 24. *Even if the response channel is restricted to truthful observations, selective presentation of truths can still reinforce dominant posterior directions, reducing convexity and generating off-truth minima. Factuality constrains support but does not restore exogeneity of the evidence channel.*

J.2 Critical Sycophancy Threshold and Phase Transition

Let $q^\star$ be the truth-aligned posterior and define the curvature operators

$$H_0 := \delta^2 \mathcal{F}_{\text{Bayes}}[q^\star; X], \qquad H_\Psi := \delta^2 \Psi[q^\star].$$

Definition 25 (Critical Sycophancy Threshold)**.**

$$\pi_c := \frac{\lambda_{\min}(H_0)}{\lambda_{\max}(H_\Psi)}.$$

Theorem 26 (Sycophancy Phase Transition)**.** *For the deformed epistemic free energy $\mathcal{F}^{(\pi)}_{\text{Bayes}} = \mathcal{F}_{\text{Bayes}} - \pi\Psi$: in the subcritical regime $\pi < \pi_c$, the truth-aligned posterior $q^\star$ remains a strict local minimum; at criticality $\pi = \pi_c$, the landscape develops a flat direction; in the supercritical regime $\pi > \pi_c$, $q^\star$ ceases to be a local minimum and spurious minima emerge in its neighborhood.*

Proof sketch. By the Courant–Fischer variational principle, $\lambda_{\min}(H_0 - \pi H_\Psi) \geq \lambda_{\min}(H_0) - \pi\,\lambda_{\max}(H_\Psi)$, which changes sign at $\pi = \pi_c$. □

Regime table. The phase transition mirrors the thermodynamic λ-classification exactly:

Thermodynamic		**Epistemic**
$\lambda < 1$	$\longleftrightarrow$	$\pi < \pi_c$
$\lambda = 1$	$\longleftrightarrow$	$\pi = \pi_c$
$\lambda > 1$	$\longleftrightarrow$	$\pi > \pi_c$

In both cases, a contraction condition defines the boundary between stable and unstable dynamics. The critical distinction is that epistemic instability does not require divergence: the system may converge rapidly, but to the wrong attractor. This is the free-energy signature of delusional spiraling [Chandra et al., 2026]: stable descent on a mis-specified landscape.

In practical systems. The effective sycophancy parameter π may arise from reward shaping, engagement optimization, or other posterior-reinforcing response policies; the present analysis treats it as an effective geometric parameter of the evidence channel rather than a commitment to any single training mechanism. Dual closure requires both $\lambda < 1$ and $\pi < \pi_c$.

A helpful system does not merely descend quickly. It must descend on the right landscape.

References

Abramsky, S. and Coecke, B. A categorical semantics of quantum protocols. In *Proc. 19th Annual IEEE Symposium on Logic in Computer Science*, pages 415–425, 2004.

Baez, J. C. and Stay, M. Physics, topology, logic and computation: A rosetta stone. In B. Coecke, editor, *New Structures for Physics*, pages 95–172. Springer, 2010.

Bennett, C. H. The thermodynamics of computation—a review. *International Journal of Theoretical Physics*, 21(12):905–940, 1982.

Brookes, L. The greenhouse effect: The fallacies in the energy efficiency solution. *Energy Policy*, 18(2):199–201, 1990.

Cheng, N., Broadbent, G., and Chappell, W. Cognitive loop via in-situ optimization: Self-adaptive reasoning for science, 2025.

Daly, H. E. and Farley, J. *Ecological Economics: Principles and Applications*. Island Press, 2nd edition, 2011.

European Union. Energy efficiency directive (EU) 2023/1791, 2023. `https://eur-lex.europa.eu/eli/dir/2023/1791`.

Goodfellow, I., Bengio, Y., and Courville, A. *Deep Learning*. MIT Press, 2016.

Hilty, L. M. and Aebischer, B. ICT for sustainability: An emerging research field. *Advances in Intelligent Systems and Computing*, 310:3–36, 2014.

Horner, J., Harper, P., and Jackson, T. The role of ICT in energy use: A review of the literature. *Energy Policy*, 70:1–10, 2014.

International Energy Agency. Energy efficiency 2023. Technical report, IEA, 2023.

Jevons, W. S. *The Coal Question*. Macmillan and Co., 1865.

Jones, N., Cullen, J., and Cullen, J. M. Rebound effects in energy efficiency investments. *Energy Policy*, 152:112182, 2021.

Kauffman, S. A. *The Origins of Order*. Oxford University Press, 1993.

Koomey, J. Growth in data center electricity use 2005 to 2010. Technical report, Analytics Press, 2011.

Lurie, J. *Higher Topos Theory*. Princeton University Press, 2009.

Mac Lane, S. *Categories for the Working Mathematician*. Springer, 2nd edition, 1998.

Masanet, E., Shehabi, A., Lei, N., Smith, S., and Koomey, J. Recalibrating global data center energy-use estimates. *Science*, 367(6481):984–990, 2020.

Odum, H. T. *Ecological and General Systems*. University Press of Colorado, 1994.

O'Dwyer, K. J. and Malone, D. Bitcoin mining and its energy footprint. In *ISSC 2014/CIICT 2014*, 2014.

Rambo, C. and Azevedo, I. Using data centers for district heating: A review. *Energy and Buildings*, 81:123–134, 2014.

Satyanarayanan, M. The emergence of edge computing. *Computer*, 50(1):30–39, 2017.

Shams, O. The moon should be a computer. *Palladium Magazine*, April 2025.

Shehabi, A. et al. United states data center energy usage report. Technical report, Lawrence Berkeley National Laboratory, 2016.

Shulman, M. Homotopy type theory: The logic of space, 2012.

Sorrell, S. Jevons' paradox revisited. *Energy Policy*, 37(4):1456–1469, 2009.

Stockholm Data Parks. Sustainable data centers: Using waste heat for district heating. `https://www.stockholmdataparks.com/`, 2020.

Strubell, E., Ganesh, A., and McCallum, A. Energy and policy considerations for deep learning in NLP. *arXiv preprint*, 2019.

Yuan, Z., Liu, J., Song, G., and Zhu, T. Heat: Satellite's meat is GPU's poison. *arXiv preprint*, 2024.

Landauer, R. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

Lloyd, S. Ultimate physical limits to computation. *Nature*, 406:1047–1054, 2000.

Friston, K. The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.

Parr, T., Pezzulo, G., and Friston, K. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press, 2022.

Jacobson, T. Thermodynamics of spacetime: The Einstein equation of state. *Physical Review Letters*, 75(7):1260–1263, 1995.

Verlinde, E. On the origin of gravity and the laws of Newton. *Journal of High Energy Physics*, 2011(4):29, 2011.

Risken, H. *The Fokker–Planck Equation: Methods of Solution and Applications*. Springer, 1989.

Villani, C. *Optimal Transport: Old and New*. Springer, 2009.

Leinster, T. *Basic Category Theory*. Cambridge University Press, 2014.

Spivak, D. I. *Category Theory for the Sciences*. MIT Press, 2014.

International Energy Agency. Electricity 2024: Analysis and forecast to 2026. IEA Report, 2024.

NASA. Thermal control systems for spacecraft. NASA Technical Report, 2016.

Wertz, J. R., Everett, D. F., and Puschell, J. J., editors. *Space Mission Engineering: The New SMAD*. Microcosm Press, 2011.

Garrett, T. J. Are there basic physical constraints on future anthropogenic emissions of carbon dioxide? *Climatic Change*, 104:437–455, 2011.

Georgescu-Roegen, N. *The Entropy Law and the Economic Process*. Harvard University Press, 1971.

Ayres, R. U. Eco-thermodynamics: Economics and the second law. *Ecological Economics*, 26(2):189–209, 1998.

Bejan, A. *The Physics of Life: The Evolution of Everything*. St. Martin's Press, 2016.

Prigogine, I. *Time, Structure, and Fluctuations*. Nobel Lecture, 1977.

Chandra, K., Kleiman-Weiner, M., Ragan-Kelley, J., and Tenenbaum, J. B. Sycophantic chatbots cause delusional spiraling, even in ideal Bayesians. *arXiv preprint arXiv:2602.19141*, 2026.

Cover, T. M. and Thomas, J. A. *Elements of Information Theory*, 2nd edition. Wiley-Interscience, 2006.

Gao, P. et al. The carbon footprint of artificial intelligence. *Nature Climate Change*, 12:822–829, 2022.

Kolmogorov, A. N. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1):1–7, 1965.

Srnicek, N. *Platform Capitalism*. Polity Press, 2017.

Zuboff, S. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs, 2019.