From Episodes to Populations: Evolutionary Explanation Requires a Constructive Epistemology

Dan Adler

Corresponding author(s). E-mail(s): dan@danadler.com;

Abstract

The sciences divide into those that discover laws and those that reconstruct histories. We argue that this division does not reflect a difference in subject matter, but a difference in *epistemic regime*. Law-based sciences operate under *episodic closure*: systems are idealized so that the outcomes of prior interactions do not alter the rules governing future ones. This *regime-defining idealization* (distinguished from *pragmatic idealization*) underlies the predictive successes of physics, but creates a systematic blind spot for evolutionary dynamics. We formalize this distinction using Stability-Driven Assembly (SDA), a minimal non-equilibrium framework in which differential persistence couples episodes into population-level evolutionary dynamics without genes, replication, or predefined fitness functions. Representing compositional objects as $\boldsymbol{\lambda}$-calculus terms, we show that episodic science studies isolated $\boldsymbol{\lambda}$-reductions under fixed rules, while evolutionary science studies populations of $\boldsymbol{\lambda}$-instantiations whose outputs re-enter the space of operators. The resulting dynamics are self-modifying and irreducibly sequential: each step rewrites the conditions for the next. A four-quadrant taxonomy locates episodic science, evolutionary science, and two commonly conflated intermediate cases: formal possibility and constructive potential, within a single framework. From this analysis we derive the "No Free Telos" constraint: in constructive systems where population feedback reshapes the effective dynamics at each step, the cost of predicting future states cannot be reduced below the cost of simulating the generative history. The resulting framework bridges episodic and historical sciences, not by reducing one to the other, but by identifying population-level memory as the structural condition that transforms law-governed episodes into open-ended evolutionary processes.

Keywords: episodic and constructive regimes, stability-driven assembly, differential persistence, evolutionary explanation, No Free Telos, prebiotic evolution, idealization, lambda calculus

1 Introduction: The Boundary of "Unreasonable Effectiveness"

The apparent divide between the physical and biological sciences is widely treated as a difference in subject matter: physics studies simple, law-governed systems, while biology studies complex, historically contingent ones. We argue that this framing is mistaken. The divide is not necessarily between physics and biology but between two *epistemic regimes*, and any science can occupy either one. What we call *episodic regimes* idealize systems so that experimental trials are independent, state spaces are fixed in advance, and dynamics are governed by transition rules invariant across instantiations. What we call *constructive regimes* are systems in which the outputs of one interaction become the operators for the next, creating a feedback loop in which history accumulates as structure. Physics typically studies the former, whereas biology necessarily studies the latter. But the distinction is structural, not disciplinary: a chemical system allowed to remember its history would become evolutionary, and a biological system forced to forget would become episodic.

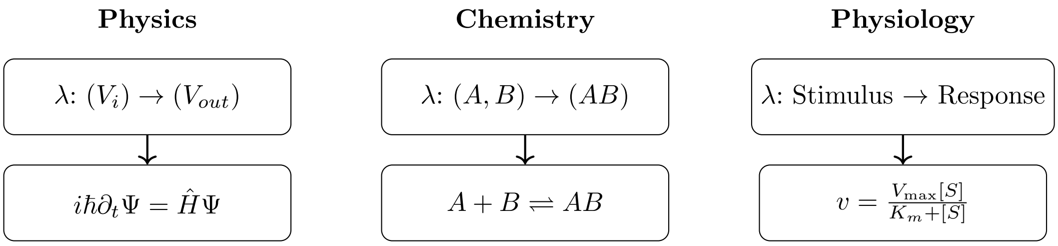

Fig. 1: Episodic modeling as λ-inference. In physics, chemistry, and physiology, predictive laws (bottom) are induced from discrete observational episodes (top). Each episode is treated as an application of an invariant transformation rule (λ-term), discovered by induction over episodes.

The "unreasonable effectiveness of mathematics" in the physical sciences [1] is thus a regime selection effect. Mathematics is effective because physical inquiry has historically restricted itself to episodic regimes: fixed phase spaces and deterministic flow in classical mechanics; Markovian approximations in statistical mechanics and chemical kinetics [2–4]; fixed Hilbert spaces and invariant operator algebras in quantum mechanics (Figure 1). Under these conditions, inductive inference succeeds: many realizations of an interaction yield a compact causal schema: a closed-form equation or a statistical law that applies universally across contexts [5, 6]. When mathematics falters in biology, economics, or geology [7–9], the issue is not a failure of formalism but a mismatch of regime: episodic tools applied to constructive regimes.

Constructive regimes, shown schematically in Figure 2, arise when episodes are allowed to remember each other. The outputs of each λ-instantiation do not disappear. They persist, aggregate, and form a population that feeds back to alter the probability distribution of future interactions. The "laws" of the system are no longer fixed invariants, but rather transient constructions of the population history. The framework developed here seeks to explain how such self-modifying dynamics, and the entropy reducing organization they produce, can arise in non-living, open, non-equilibrium systems through stochastic assembly and differential persistence.

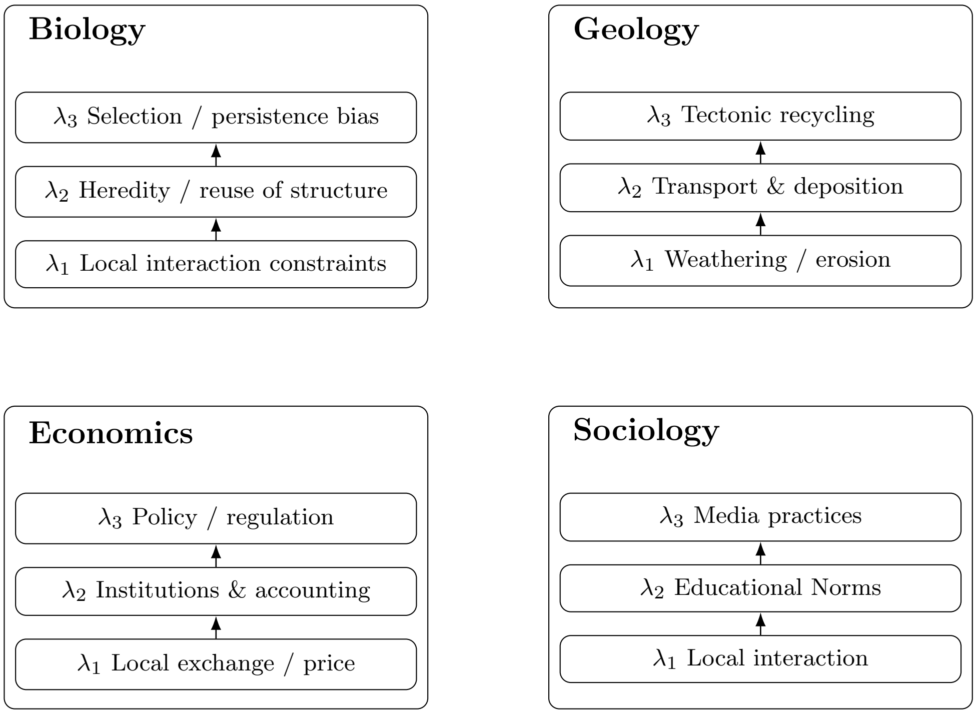

Fig. 2: Hierarchical λ-inference in historical sciences. In complex domains, the invariant "law" is not a single term but a stack of operators. Each level of constraint $(\lambda_1, \lambda_2, \dots)$ emerges from the population dynamics occurring under the constraints of the level below (e.g., stable patterns of local interactions λ_1 give rise to institutional rules λ_2). These higher-order, population-level constraints act as shaping contexts for future episodes, creating an evolving system. Labels are illustrative.

We demonstrate this distinction using Stability-Driven Assembly (SDA), a minimal framework in which differential persistence creates the population-level memory required for evolution without genes, replication, or predefined fitness functions. We show that such systems exhibit *generative irreducibility*: the operator governing the system is a functional of the population distribution, predicting the future requires simulating the constructive history. We term this the "No Free Telos" constraint: there is no compressed shortcut to the final state because the state space itself is under construction. The resulting framework bridges physics and biology by identifying the epistemic boundary where episodic closure gives way to constructive dynamics.

2 Episodic Regimes as Methodological Artifacts

If the distinction between episodic and evolutionary regimes is so fundamental, why has the physical sciences' reliance on the former not been fatal to its success? The answer lies in a specific form of selection bias: the "unreasonable effectiveness" of mathematics in physics is largely the result of restricting inquiry to systems that can be effectively isolated from their histories.

Standard experimental design in physics is an exercise in enforcing Methodological Markovianity. Consider a particle scattering experiment or a chemical reaction study. Great care is taken to ensure that the initial conditions are precisely reset between trials. The apparatus is cleaned; the vacuum is restored; the reactants are purified. By design, the output of Trial N is prevented from becoming the input to Trial $N+1$. This methodology artificially suppresses the population-level accumulation that characterizes evolutionary systems. It forces the system to behave as a sequence of independent λ-episodes, ensuring that the inferred laws are time-invariant and the state space remains fixed. In many physical systems, such effective memorylessness reflects genuine equilibration or thermalization dynamics; the claim is not that invariance is artificial, but that the explanatory strategy of physics depends on identifying and isolating regimes in which invariance holds.

This methodological stance forces us to distinguish two senses of "memorylessness," often conflated in scientific practice. *Phenomenological Markovianity* occurs when the system genuinely "forgets" its history due to thermalization or equilibration (e.g., a gas in a box). *Methodological Markovianity* occurs when the system is *prevented* from remembering its history by the intervention of the experimenter (e.g., resetting the apparatus). The former is a property of the target system. The latter is a property of the epistemic framework. When we model evolutionary phenomena using the tools of physics, such as differential equations on fixed manifolds, we are implicitly imposing Methodological Markovianity on systems whose defining feature is the violation of it.

2.1 Regime-Defining Idealization

Philosophers of science have long recognized the role of idealization in managing complexity. Potochnik [10] argues that idealizations are not merely simplifications but essential tools for isolating causal patterns. We propose a distinction between two fundamentally different types: *Pragmatic idealization* suppresses causal factors that are present but negligible, like ignoring air resistance in a falling body. It acts as a filter, removing noise to reveal a signal, and the suppressed factors can, in principle, be restored without altering the model's structure. *Regime-defining idealization*, in contrast, structurally eliminates the capacity for state-space expansion, as when a biosphere is modeled as a fixed system of differential equations. This acts not as a filter but as a cage: what is excluded is not a negligible perturbation but the very mechanism by which the system generates novelty.

This distinction refines Cartwright's (1983, 1999) influential argument that the laws of physics "lie", holding exactly only within idealized models. For Cartwright, the issue is that real-world complexity prevents the antecedent conditions of laws from being satisfied. In our account, the deeper issue is structural: regime-defining idealization does not merely limit a law's domain of application but eliminates the capacity for the system to generate its own constraints. The distinction is between a world where fixed laws have limited scope and a world where the effective laws are themselves under construction.

Standard dynamical models, particularly those intended to capture long-run equilibria, fall into the second category. When we write a system of differential equations $\dot{x} = F(x)$, we pre-specify the dimensionality of the vector x and the functional form of F. This formalism structurally forbids the emergence of new dimensions (new observables) or the modification of F by the system itself. It does not just simplify the evolutionary process; it deliberately eliminates it.

2.2 The "Cage" of Fixed Phase Space

This leads to what we term the *Cage Effect*. By modeling a constructive system using episodic tools, we trap the dynamics in a pre-defined phase space. In the episodic regime, a "novelty" is the visitation of a previously unvisited coordinate in a fixed space (a state discovery). In a constructive regime, novelty is the generation of a new coordinate axis altogether (a dimension discovery). Other accounts have shown that idealizations do constitutive work beyond simplification [12, 13]; our claim is more specific: regime-defining idealization eliminates the capacity for state-space expansion, separating filtering from caging.

Physics excels at state discovery because its methodology is designed to keep the axes fixed. Biology, geology, and economics struggle under this methodology because their central phenomena involve the proliferation of new axes: new species, new minerals, and new markets. By treating these historical sciences as "messy physics," we force constructive territories into episodic maps, not merely approximating but structurally mischaracterizing the target phenomenon. We are searching for laws of motion governing the studs, while the phenomenon of interest is the construction of the castle. To model construction, we require a formalism in which composition is not a separate capability, but the engine of dynamics itself. We find this in the lambda calculus.

3 The Stability-Driven Assembly (SDA) Framework

3.1 Stability-Driven Assembly: The Constructive Mechanism

Stability-Driven Assembly (SDA) serves as our operational model for the transition from episodic to constructive dynamics. The framework was introduced by Adler [14] as a minimal non-equilibrium system in which differences in persistence bias the accumulation of structure, with full mathematical derivations demonstrating entropy reduction and scaffold emergence. Further work Adler [15] extended SDA to chemical symbol space, showing hallmark features of evolutionary search, such as the emergence of dominant scaffolds and sustained novelty, over thousands of generations. Here, *stability* and *persistence* denote resistance to dissipation over time, measured as a characteristic lifetime, analogous to half–life in physics. For the present argument, we focus on the core dynamical mechanism: how persistence-weighted feedback reshapes the population distribution and induces evolutionary search without explicit replication.

An SDA system consists of a population of interacting entities in an open-flow setting: energy and base elements flow through a bounded interaction volume, maintaining the system far from equilibrium while ensuring that the population is finite and that a well-defined probability distribution over patterns exists at each generation. The dynamics are governed by a simple asymmetry: patterns are generated stochastically but persist differentially. New patterns are created through the interaction of existing ones, e.g., concatenation, bonding, or recombination, depending on the domain. Each pattern is assigned a finite lifetime based on its stability, and patterns that persist longer accumulate in the population, increasing their frequency, and thus their probability of participating in future interactions. This closes the feedback loop: persistence shapes the composition, composition biases the interaction, and the interaction generates new patterns whose persistence further reshapes the population.

A simple example illustrates the mechanism. Consider a system with base elements A and B, where the compound AB has a lifetime of 10 generations, $ABAB$ has a lifetime of 50 generations, and all other compounds degrade instantly. After the first generation, only A, B, and AB persist; unstable combinations such as AA and BB are eliminated immediately. As AB accumulates due to its longer life, its frequency in the population increases, raising the probability that two copies of AB will be sampled together and recombined to form $ABAB$. Once present, $ABAB$ persists five times longer than AB, accumulating further and biasing subsequent interactions toward still higher-order configurations. No fitness function was specified; no selection operator was imposed. The population-level bias toward $ABAB$ emerged entirely from the differential between persistence times: stability created selection.

Algorithm 1 Stability-Driven Assembly (SDA): Minimal Dynamics

Initialize population with base elements
for each generation **do**
  Decay: Remove expired patterns from the population
  Flow: Replenish base elements to maintain non-equilibrium
  for each interaction step **do**
    Sample two parents A, B proportional to their abundance $P(x)$
    Generate offspring $C = \text{Interact}(A, B)$
    Assign lifetime τ_C based on stability function $S(C)$
    Add C to population
  end for
end for

Algorithm 1 formalizes the core loop. At each generation, patterns whose lifetimes have elapsed are removed, and base elements are replenished, maintaining the system far from equilibrium. Pairs of patterns are then sampled in proportion to their current abundance (the roulette wheel) combined to produce an offspring, and the offspring is assigned a lifetime determined by its stability. Stable offspring persist into future generations, raising their frequency and thus their probability of being sampled as parents. Unstable offspring decay quickly and contribute nothing more. No fitness function is specified; the population-level bias toward increasingly stable structures emerges from the asymmetry between creation and decay alone.

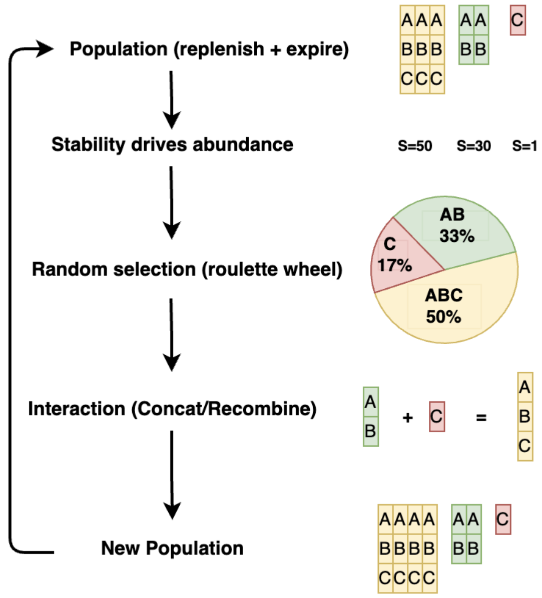

Fig. 3: **The Constructive Loop.** Base elements are continuously replenished while unstable motifs expire. Patterns are sampled from the population, combined, and re-enter with lifetimes determined by their stability. The resulting feedback from stability to persistence to population composition produces emergent selection without externally specified fitness.

Stability asymmetry introduces selection without the need for reproduction. In standard biological models, selection acts on the *rate of reproduction* (the Malthusian parameter). In SDA, selection acts on the *rate of decay*. As shown in Figure 3, this creates a closed causal loop: persistence biases population composition, population composition biases the trajectory of future interactions, and evolutionary dynamics emerge from this feedback alone.

To understand why this system breaks out of the *episodic cage effect* (Section 2), it is useful to approximate population dynamics as a continuous transport process. In the limit of large populations, the evolution of the pattern density $P(x, t)$ over a feature space x can be described by a Fokker–Planck equation Gardiner [3]:

$$\frac{\partial P(x,t)}{\partial t} = \underbrace{-\nabla \cdot \big[A[P](x,t)\, P(x,t)\big]}_{\text{Selection (Drift)}} + \underbrace{D\nabla^2 P(x,t)}_{\text{Variation (Diffusion)}} \tag{1}$$

Here, the diffusion term $D\nabla^2 P$ captures the dispersive effect of stochastic pattern creation (the "Create" process). The drift term $A[P]$ captures the biasing effect of differential persistence.

The crucial feature of Equation 1 is the notation $A[P]$. In standard physical systems (e.g., a particle in a magnetic field), the drift field $A(x)$ is an *external potential* fixed by the boundary conditions. It is an independent operator. However, in SDA, the drift is a *functional of the population distribution*. The "force" pulling the system toward specific regions of pattern space is not pre-existing; it is generated by the current population of stable scaffolds acting as attractors for future binding events.

This nonlinearity (A depends on P, which depends on A) is the mathematical signature of a constructive regime. The system is historically closed: the effective laws of motion at time t are determined by the accumulated structures from time $t - 1$. This precludes the existence of a static equilibrium distribution that can be solved analytically (as in the Boltzmann distribution). To know the future distribution, one must integrate the path-dependent history of the functional $A[P]$.

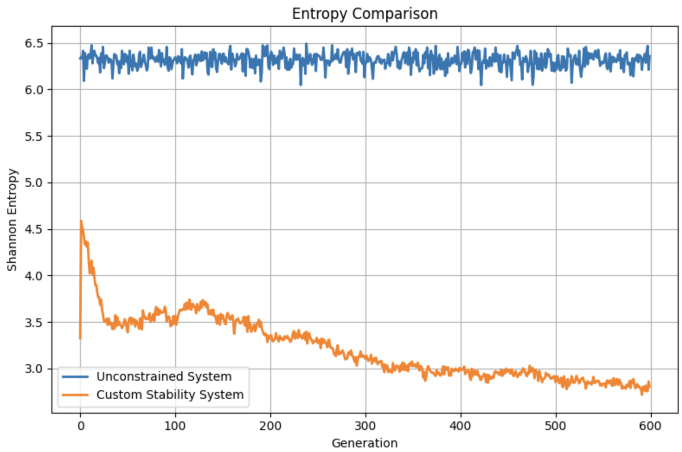

Fig. 4: **Entropy Reduction via Persistence.** Evolution of the Shannon entropy of the population distribution in a symbolic SDA simulation Adler [15]. In the unconstrained control (blue), entropy remains high as the system explores pattern space randomly. Under SDA dynamics (orange), entropy decreases as probability mass concentrates on families of stable, long-lived motifs.

The consequence of this functional feedback is the spontaneous reduction of entropy. Figure 4 compares the Shannon entropy of a population under SDA dynamics against a neutral control. In the control case (random assembly without persistence bias), the system diffuses broadly through the combinatorial library (a regime we will later formalize as quadrant Q3), resulting in high entropy. In the SDA case, the population effectively discovers stable sub-structures. These structures persist, creating local

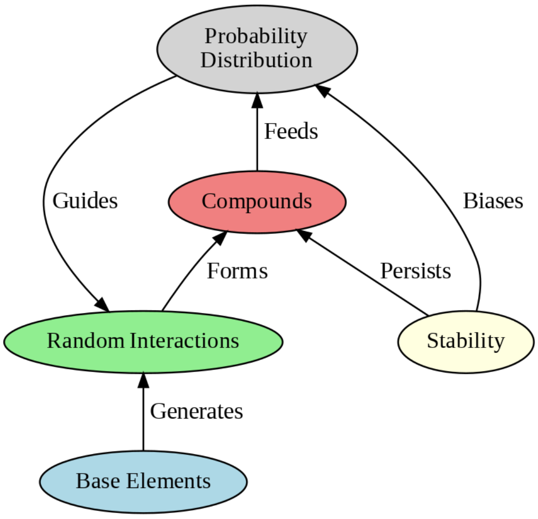

Fig. 5: Core feedback structure underlying SDA. Differential persistence biases population composition, which biases future interactions, closing a causal loop in which selection emerges without genes, replication, or an imposed fitness function.

density spikes that bias future sampling, leading to a cascading collapse of entropy as the system locks onto a set of dominant, stable scaffolds.

This result demonstrates that evolutionary ordering is not a unique property of biological replication. It is a general physical property of open systems where persistence depends on structure. By simply allowing the output of an interaction to influence the probability of future interactions, SDA transitions the system from a fixed episodic probability distribution to an evolving constructive history.

3.2 From Persistence to Selection

Because sampling probability is proportional to population frequency, and frequency is determined by persistence, stability implicitly defines a fitness landscape. This feedback implements fitness-proportional sampling [16] without an external fitness function: the system realizes a *natural genetic algorithm* (SDA/GA) Adler [15], in which variation enters through stochastic assembly and selection emerges from differential lifetimes alone. Figure 5 summarizes this feedback loop: persistence shapes population composition, composition biases future interactions, and interactions generate new patterns whose persistence further reshapes the population.

It is important to be precise about what "selection" means here. Since the Modern Synthesis, evolution has been understood through Lewontin's triad: variation, differential fitness, and heredity Lewontin [17], Godfrey-Smith [18]. In practice, heredity is almost always operationalized as template copying: a mechanism by which individuals transmit information to offspring. This identification has become sufficiently entrenched that evolutionary dynamics, a population-level phenomenon, are often treated as inseparable from reproduction, an individual-level mechanism.

SDA satisfies all three conditions through different means. *Variation* arises from stochastic assembly. *Selection* arises from differential decay. *Heredity* arises not from template copying but from what we call *structural recurrence*: a persistent pattern remains available as a reactant, increasing the frequency of interactions that regenerate it or structures containing it. The “offspring” of a pattern are not copies manufactured by the pattern itself but new instances independently generated by the system's interaction dynamics, sharing structural identity without sharing lineage, as shown in the example of Figure-3.

This generalization engages directly with Griesemer's (2000, 2005) influential critique of the Modern Synthesis's conflation of replication with reproduction. Griesemer argues that heredity requires *material overlap* between generations: progenitor and offspring must share physical stuff for inheritance to occur. Structural recurrence is weaker than Griesemer's reproduction (no material continuity required) and weaker than replication (no copying mechanism required), yet sufficient to couple episodes into evolutionary trajectories. Whether this further generalization illuminates or obscures the special character of biological inheritance is an open question, but it suggests that heredity admits of degrees, with template-based replication representing one particularly efficient implementation of a more general population-memory function.

This generalization is not ad hoc. As Maynard Smith and Szathmáry [21] argued, the earliest evolutionary systems likely operated under conditions of limited heredity without high-fidelity replication. What replication provides is a particularly efficient implementation of structural recurrence, but it is not the only one. The shift from replication to persistence as the primitive selective quantity reveals evolution as a statistical property of population histories rather than a biological property of individuals, extending evolutionary explanation to domains where replication is absent or undefined: prebiotic chemistry, institutional change, and the constructive regimes that are the subject of this paper. In the formalism introduced below, these interaction pathways correspond to instantiated λ-terms: each episode is a λ-application, and structural recurrence preserves λ-identity across independently generated instances.

Mitchell [8] argues that biological phenomena require *integrative pluralism*: explanations that combine multiple levels and causal factors without expecting a reduction to a single framework. The SDA framework identifies a structural reason why: populations generate level-specific constraints, with the “rules” at each level of the explanatory hierarchy (Figure 2) emerging from persistence-weighted dynamics at the level below. Mitchell's pluralism thus finds a dynamical grounding in the recursive feedback structure of constructive systems.

This perspective also clarifies the debate over the Extended Evolutionary Synthesis (EES). Proponents argue that developmental plasticity, niche construction, and extra-genetic inheritance require revisions to neo-Darwinism [22, 23]; critics respond that these phenomena are already accommodated [24]. Our analysis suggests that the debate conflates two questions: the *mechanism* of evolution (which EES seeks to extend) and the *conditions* under which evolutionary dynamics arise (which SDA characterizes). Both camps may be correct about the mechanism while sharing an unexamined assumption about what makes a system evolutionary in the first place, an assumption the episodic/constructive distinction makes explicit.

3.3 Turing vs Church: Library vs Lego

The equivalence of Turing machines and the λ-calculus is a cornerstone of computability theory: both capture the same class of effectively computable functions and confront the same undecidable boundaries Turing [25], Church [26]. Yet they achieve this equivalence through radically different architectures, and these differences matter when we move from the theory of *what can be computed* to the question of how computational structures *evolve*.

A Turing machine solves exactly one problem. Its state-transition table is fixed at design time; if you need a new function, you build a new machine. Each machine is like a *book in a library*: complete, static, and self-contained. Einstein captured this intuition: "We are in the position of a little child entering a huge library filled with books in many languages... The child dimly suspects a mysterious order in the arrangement of the books but doesn't know what it is" Isaacson [27]. In the episodic regime, the library is the correct metaphor: the book of nature is already written, and discovery is retrieval. But to expand the collection one must write a new volume from scratch. There is no reuse of internal structure across books except by deliberate offline copying.

The λ-calculus, in contrast, is compositional by nature. A λ-term is not a book but a *Lego block*: it can be applied to other terms, partially evaluated, passed as an argument, or returned as a result. New functions are assembled by applying existing terms to each other. The Church numeral $(\lambda f.\lambda x.f(fx))$ does not need to know in advance what f will be. It is a *scaffold* that awaits completion.

One might object that Von Neumann's universal constructor shows a Turing machine *can* replicate itself von Neumann [28]. But the replication is total and literal: the descendant is an identical copy, and the constructor must be built into the original design. There is no mechanism by which a machine, solving its own problem, spontaneously gives rise to a new machine solving a different problem. Innovation requires an a-priori architect. In the λ-calculus, composition is not a separate capability installed alongside a computation. It *is* itself a computation. A term that transforms A into B can be combined with a term that transforms B into C, producing a new term without rewriting either predecessor. The population remembers solutions not as discrete artifacts, but as recombinable capabilities. The library catalogs the past; the Lego set builds the future.

This is why the operators in Figure 2 are written as λ-terms, not Turing tables: λ_2 arises from the population dynamics of λ_1 interactions because composition is already what λ_1 *does*. The machinery of emergence need not be installed from without; it is already operative within.

3.4 Formalizing SDA in the λ-Calculus

This structural equivalence was first recognized by Fontana and Buss [29], who mapped chemical objects to λ-calculus terms precisely because the λ-calculus makes no distinction between program and data: every term is simultaneously a potential operator and a potential operand. Fontana and Buss showed that this reflects the ontological structure of constructive chemistry, where a molecule is both a thing acted upon (a reactant) and a thing that acts (a catalyst or scaffold). Their brilliant insight generalizes beyond chemistry: a firm is both an entity in a market and a set of routines that transform inputs; a string is both a pattern to be recombined and a component that biases future recombination. The λ-calculus captures this shared constructive structure not because these domains are "really" computational, but because they are all systems in which the products of interaction re-enter the space of operators.

What SDA adds to Fontana's ontological insight is a selection mechanism. Where Fontana required self-maintenance to filter the space of possible organizations, SDA shows that differential persistence alone is sufficient. Each recombination event in Algorithm 1 corresponds to a λ-application: two terms are drawn from the population, one acts upon the other, and the result either persists into P_{t+1} or dissipates. The population is thus not a vector in a fixed state space but an evolving multiset of λ-terms whose composition at each generation determines the operators available for the next.

3.5 Self-Reference and the Failure of Spectral Decomposition

In episodic regimes, the transition rules governing a system are time-invariant. A fixed linear operator can be spectrally decomposed into eigenvalues and eigenvectors, yielding closed-form expressions for long-run behavior. This is why physics admits symmetry arguments, dimensional reduction, and predictive shortcuts.

In SDA, this factorization is not available and the reason is structural. The interaction rules at generation t depend on the population distribution P_t: which patterns are present, how abundant they are, and therefore which λ-applications are likely to occur. Since P_t is itself the cumulative result of all prior interactions, the effective dynamics at each step encode the population history that produced them. To factor an operator from its operands, one must be able to specify the operator independently of the states upon which it acts. In SDA, the two are intertwined: the interaction rules cannot be written down without knowing P_t, and P_t cannot be computed without evaluating the interaction rules in every preceding step.

Kleene's recursion theorem provides the formal basis for this observation: any computational system powerful enough to represent its own operations can construct self-referential programs containing their own description as a sub-expression Kleene [30], Rogers [31]. Once compounds-as-operators feed back into the population: a compound's persistence alters the sampling distribution, which alters which compounds are produced, which alters persistence. The system instantiates exactly this kind of self-reference. There is no fixed matrix to diagonalize and no time-independent Hamiltonian to solve. The dynamics are irreducibly sequential: each step rewrites the conditions for the next, and no reordering or parallelization of the history can preserve the outcome. This *generative irreducibility* is the formal basis of the constraint we develop in Section 6.

This does not mean that prediction is never possible. When the population distribution changes slowly relative to the observation timescale, when a small number of high-stability scaffolds dominate, and the probability flow through the interaction network reaches a temporary steady state, the effective dynamics is approximately fixed, and standard analytical tools apply locally. The system behaves, for a time, *as if* it were episodic. But this equilibrium is always provisional: the same persistence-weighted feedback that stabilizes the current configuration also concentrates probability mass on the interactions most likely to produce novel high-stability compounds. When such a compound emerges, it can reorganize the entire population distribution, invalidating the local approximation. Predictability in constructive systems is thus intermittent: present during quasi-stable periods, absent across the transitions between them.

4 The Four Quadrants

The preceding analysis suggests that the distinction between episodic and evolutionary explanation is not a matter of domain but of dynamical structure. We organize scientific explanation along two independent dimensions: whether the objects of study are abstract formal schemata or physically instantiated populations, and whether the governing dynamics are fixed or endogenously modified by the process itself. The resulting map (Figure 6) locates episodic science, evolutionary science, and two commonly conflated intermediate cases within a single framework.

Q1: The Formal Regime (abstract, fixed). Pure mathematics, formal logic, and the λ-calculus *as a formal system* occupy this quadrant. Objects are purely abstract, not instantiated in any physical reality, but rather defined by axioms and inference rules that do not change. A theorem proved today was provable yesterday and will be provable tomorrow. Q1 provides the schemata, the λ-terms, that other quadrants instantiate, but it contains no populations, no time, and no history. Whether these formal structures exist independently of human cognition (as mathematical Platonism holds) or are epistemic constructions is not at issue here. What matters for our framework is that they constrain the dynamics of the quadrants in which they are instantiated.

Q2: Episodic Science (instantiated, fixed). This is the home of classical and quantum mechanics, chemical kinetics, physiology, and equilibrium statistical mechanics, systems where physically realized populations exist but episodes are independent. The apparatus is reset between trials, the state space is fixed in advance, and the transition rules are invariant. Prediction succeeds here because the operator F in $x_{t+1} = F(x_t)$ can be specified independently of the states upon which it acts, allowing spectral decomposition, and all the predictive machinery of Section 3.5. Classical game theory and mean-field economics also reside here: agents interact within a fixed strategy space, and selection operates over states rather than structures. Q2 is where Wigner's "unreasonable effectiveness" holds, precisely because Methodological Markovianity is enforced.

Q3: The Library (abstract, constructive morphospace). The space of all possible proteins, the adjacent possible [32] as a static set, these are possibility spaces whose structure is combinatorially rich and, in principle, open-ended, but which contain no physically instantiated populations. Iterative and recursive systems such as cellular automata [33], Conway's Game of Life, and fractal generators also reside here: they exhibit complex, seemingly evolutionary behavior, but their dynamics are governed

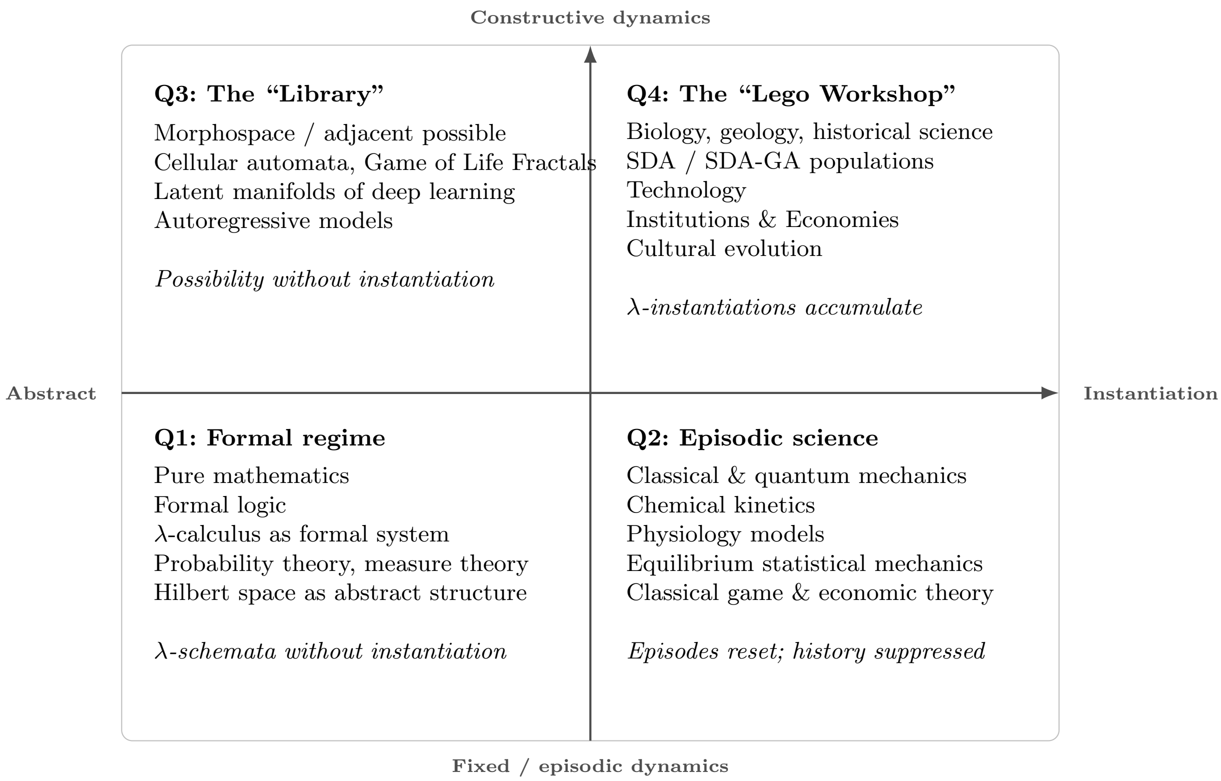

Fig. 6: **The Four Quadrants**: The horizontal axis distinguishes abstract formal schemata from physically instantiated populations. The vertical axis distinguishes fixed dynamics (invariant operators on fixed state spaces) from constructive dynamics (operators and state spaces co-evolve).

by fixed rules applied to fixed grids, with no population-level feedback and no instantiation in the physical world. Because they pay no thermodynamic cost, they erase no information and construct nothing. They merely enumerate possibilities within the cage of their initial conditions. Autoregressive models and single-agent reinforcement learning similarly inhabit Q3: they carry memory of their own outputs but lack population-level selection operating on physically instantiated variants. Q3 answers the question *what could exist* but not *what does exist*. It catalogs the Library, but it does not write new books.

Q4: The Lego Workshop (instantiated, constructive). Biology, geology, real (non-equilibrium) economics, technology, and cultural evolution occupy this quadrant. Here, populations of physically realized structures persist differentially, they feed back on interaction probabilities, and generate novel structures that re-enter the space of operators. The effective dynamics at each step are shaped by the population that previous steps produced. Q4 is where the λ-schemata of Q1, constrained by the laws discovered in Q2, are instantiated as evolving populations whose trajectories cannot be predicted without simulation.

The quadrants are not hermetic. Movement between them defines the central transitions in the history of science. The passage from Q1 to Q2 is *instantiation*: a formal schema acquires physical referents and becomes an empirical law. Newton's inverse-square law exists as a mathematical relation in Q1; applied to planetary orbits, it becomes episodic science in Q2. The passage from Q2 to Q4 is the main claim of the paper: episodic science becomes evolutionary science when population-level memory couples episodes over time. A chemical reaction studied in isolation (Q2) becomes a constructive system when its products persist, accumulate, and bias future reactions (Q4). The passage from Q3 to Q4 is *actualization*. The Library contains every possible protein; the Workshop contains only those that a specific history of persistence-weighted assembly has constructed. No amount of contemplation of Q3 produces a single entry in Q4. For that, one needs energy, time, and a population.

The distinction between Q3 and Q4 echoes Kant's observation [34] that one cannot reason one's way from a concept to an object: the concept of one hundred dollars is identical to the concept of one hundred *existing* dollars, yet the financial consequences are quite different. The *possibility* of a complex organism is a mathematical fact given by the laws of chemistry and combinatorics, but its *actuality* requires a constructive history. A specific sequence of persistence-biased interactions that assembled it from available precursors.

The distinction between Q2 and Q4 reframes the long-standing divide between sciences that seek general laws and those that reconstruct particular histories [35, 36]. What determines whether a science is historical is not its subject matter (chemistry studies repeatable reactions, yet geochemistry is historical), but whether its explanatory framework structurally permits or suppresses population-level memory.

Kauffman's concept of the adjacent possible [32] can be understood as the dynamic boundary between Q3 and Q4. For any instantiated population in the Workshop, the adjacent possible constitutes the immediate subset of the Library that is exactly one constructive step away from realization. As the population actualizes novel structures through persistence-biased interaction, the boundary expands into previously inaccessible regions of Q3. Because the adjacent possible of tomorrow depends entirely on what is constructed today, the trajectory through the Library cannot be mapped in advance; it must be built.

Relational approaches such as [37] correctly identify Q4 as the target regime by emphasizing organizational closure, but by treating closure as a definitional primitive rather than an emergent outcome of physical dynamics, they leave the constructive mechanism unspecified: how biological closure is achieved from non-living precursors. The "No Free Telos" constraint, developed in Section 6, is the computational cost of traversing the Kantian gap.

5 Thermodynamic Erasure and the Cost of Construction

The Kantian gap between Q3 and Q4 is not only epistemic but also thermodynamic. In the episodic regime, dynamics are fundamentally reversible: unitary quantum evolution and Hamiltonian classical mechanics preserve information. Given the final state, one can, in principle, recover the initial state. Evolutionary construction is irreversible because selection is information erasure: to construct a specific structure from a vast combinatorial space requires discarding the configurations that were *not* selected.

This connection is formalized by Landauer's Principle Landauer [38]: erasing a single bit of information requires dissipating at least $k_B T \ln 2$ of heat into the environment. In SDA, every decay event is a Landauer erasure: the system forgets a failed configuration. When the population converges on a stable scaffold, it reduces the entropy of the probability distribution (Figure 4), and this reduction must be paid for by exporting entropy to the environment. The population acts as a collective Maxwell's Demon [39]. Interaction tests for stability. Decay erases what fails. Energy flow pays the bill, flushing the entropy of rejected configurations into the environment so that the survivors can accumulate and compound. The path through Q4 is a tunnel dug through entropy, and the process of digging takes time.

6 The "No Free Telos" Constraint

6.1 Formal Statement and Scope

The preceding analysis yields a constraint on prediction in constructive systems. The generative irreducibility identified in Section 3.5 (the interdependence of operator and population) means that no general algorithm can compress the constructive history into a predictive shortcut.

The undecidability results of [25] establish that no *general* algorithm can predict the outcome of an arbitrary self-referential computational process. SDA systems, in which the interaction operator is a functional of the population history and composition is unconstrained, exhibit the self-referential structure to which these results apply; whether they are Turing-complete in the strict sense remains an open formal question, but the dynamical mechanism where operators rewrite the conditions for their own future application leads to a halting problem. This does not entail that every constructive system resists prediction at all times. Just as specific programs can be proven to halt even though the halting problem is undecidable in general, specific evolutionary trajectories may admit local prediction even though the class of constructive systems does not admit a universal predictive shortcut. What the formal results *do* establish is that any such predictability is *fragile* and vulnerable to disruption by emergent structure.

Figure 7 illustrates the mechanism. In Figure 7a, patterns ABA (S=50, N=33) and ABC (S=50, N=35) have accumulated substantial population mass, creating a quasi-stable configuration. In Figure 7b, the emergence of ABCABA (S=100, N=70) reorganizes the entire system: this "disruptor node" redirects probability mass flow and alters the effective dynamics of the entire system. The emergence of ABCABA

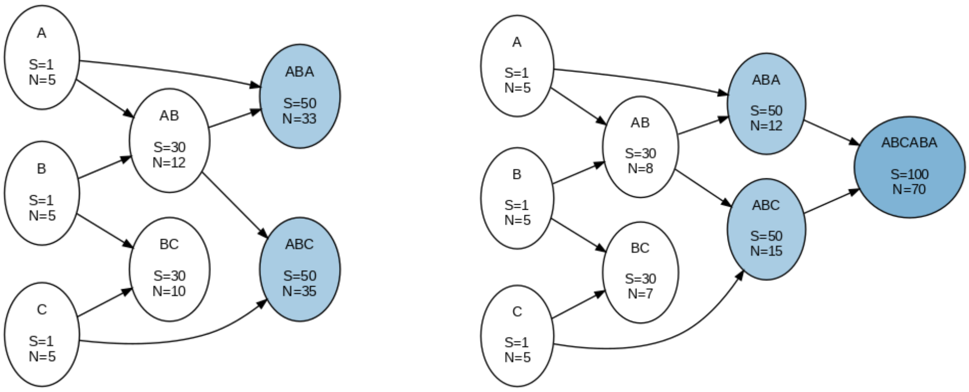

(a) Early Generation (b) Late Generation

Fig. 7: Punctuated dynamics in SDA: **(a)** Early generation with patterns ABA and ABC dominating, **(b)** Later generation after emergence of ABCABA (S=100), which reorganizes probability mass flow throughout the network.

required a specific interaction sequence (ABC + ABA → ABCABA) whose probability depended on the prior population state, but whose *consequences* (the reorganization of the fitness landscape) could not be computed without simulating the assembly. This punctuated pattern, long periods of stasis interrupted by rapid reorganization, has been independently documented in constructive domains: punctuated equilibrium in paleontology Eldredge and Gould [40], creative destruction in economics Schumpeter [41], and paradigm shifts in the history of science Kuhn [42]. What has been lacking is a dynamical explanation for why such punctuation is generic. The SDA framework provides one solution: stable scaffolds concentrate the probability mass, increasing the likelihood of interactions that produce novel combinations with still higher stability. Stasis creates the conditions for its own punctuation.

Constraint 1 (No Free Telos - NFT)**.** *For a constructive system in which the effective interaction rules at generation t depend functionally on the population history $P_0 \dots P_t$:*

1. ***Local prediction** is possible during quasi-stable periods when the population distribution changes slowly relative to the observation timescale.*
2. ***Global prediction**: determining the dominant structures at time $T \gg t$ requires running the generative history. The emergence of disruptor structures cannot be anticipated without computing the interaction sequence that produces them.*
3. ***Transitions** between quasi-stable regimes are generically unpredictable: perturbations of all sizes, including system-reorganizing cascades, occur with nonzero probability.*

The claim is not that prediction is impossible, but that predictability is *intermittent*. The telos is not free because the cost of crossing the Kantian gap, from possibility (the space of potential disruptors) to actuality (which disruptor emerges and when), is paid in computational steps that cannot be compressed.

6.2 The Assembly Index as Empirical Indicator

If constructive history is physically encoded in structure, we should expect a measurable signature. Assembly Theory Sharma et al. [43] provides one. The Assembly Index (M_A) of an object is the minimum number of recursive join operations required to construct it from basic building blocks. In our framework, M_A serves as a lower-bound indicator for the depth of the constructive path. Objects with low M_A can arise from thermodynamic fluctuation. Objects with high M_A, such as hemoglobin or a ribosome, inhabit a region of combinatorial space so vast that spontaneous formation is effectively impossible. Their existence is evidence of a history biased by persistence.

A crucial distinction must be drawn between M_A and the actual trajectory. The Assembly Index calculates the *shortest* construction path; it is a compression metric. But evolution selects for persistence, not parsimony. As Koza [44] demonstrated, evolved solutions rarely converge to minimal implementations. They exhibit "bloat," redundant structures carried along by the process, unless parsimony is explicitly included in the fitness function, which is rarely the case in natural evolution. The population finds *a* path, not the shortest path. Thus, M_A represents the theoretical minimum cost of construction; the actual cost is typically much higher.

7 Conclusion

The success of the physical sciences has long encouraged the hope that if we could find the fundamental laws and the initial conditions, the structure of the complex world would follow as a necessary consequence. This is the dream of the Library: the book of the universe is already written, and science is learning to read it.

We have argued that this dream rests on a regime-defining idealization. The methods of episodic science: spectral decomposition, fixed state spaces, and invariant operators are effective precisely because they are restricted to systems where population-level memory is structurally suppressed. By imposing Methodological Markovianity, we filter out the mechanism that makes evolution possible: the feedback by which persistent structures reshape the interactions that produce them.

The shift to a constructive epistemology does not abandon episodic science, but extends it. The four-quadrant map developed here resolves a persistent tension in the philosophy of science: the apparent divide between law-based and historical explanation reflects neither a hierarchy of rigor nor an accident of subject matter, but a structural difference in epistemic regime. Physics and biology are not in conflict; they occupy different quadrants of the same map. The domain-specific λ-schemata discoverable by episodic science define the constraints, the geometry of the studs, but the population trajectory through the space of possible constructions is determined by persistence-weighted feedback that rewrites the effective dynamics at every step to build a specific castle. The Library (Q3) catalogues what is possible; the Workshop

(Q4) constructs what is actual. The "No Free Telos" constraint, grounded in the self-referential structure of constructive dynamics and the undecidability results of Turing and Rice, establishes that no shortcut bridges the gap between them.

A recurring impulse in the history of thought, from Platonic forms to Pythagorean harmonics, to contemporary proposals that identify physical reality with the space of all possible computations [33], has been to locate causal priority in the space of possibility. In this view, the actual world is governed by, or derived from, an abstract order that precedes it. In the language of our framework, these accounts treat Q1-Q3 as causally prior to Q4. The SDA framework inverts this priority.

This inversion does not diminish the role of episodic science. The laws discoverable in Q1–Q3 are the λ-schemata that constrain every constructive step in Q4: conical symmetries give rise to inverse-square force laws, molecular orbital geometry determines which bonds are stable, and thermodynamic inequalities set the cost of each Landauer erasure. These constraints are real and inviolable. What they do not do is determine the trajectory. The geometry of the studs is fixed by physics; which castle gets built is determined by the population history of persistence-weighted assembly. SDA does not replace episodic science but identifies what episodic science cannot, in principle, provide: an account of how particular constructions emerge from the space of possibilities that physical law permits.

The constructive order of the Workshop is not derived from the Library; it is assembled within it, through stochastic interactions, differential persistence, and population-level feedback. No pre-existing blueprint, no external fitness function, and no privileged formal structure is required; Only an open thermodynamic system with asymmetries in stability, conditions that are generic to any universe whose initial symmetry is broken. SDA provides a minimal demonstration of how persistence-driven feedback generates selection, not a complete evolutionary theory. The framework does not address how stability functions themselves arise, how environmental boundary conditions change, or how the hierarchical organization depicted in Figure 2 bootstraps across levels. These questions will be subjects of future work.

The No Free Telos constraint establishes that global prediction in constructive systems is irreducible, but it does not preclude *local anticipation*. During quasi-stable periods, the current population distribution determines which λ-applications are probable: the roulette wheel of fitness-proportional sampling is observable. This means that the adjacent possible: the set of structures exactly one constructive step from realization, is not uniformly accessible but is biased by the current composition of the Workshop. Future work could exploit this asymmetry: by characterizing which high-stability configurations lie in the high-probability region of the adjacent possible, one could identify the candidate disruptors most likely to punctuate the current equilibrium. This would not predict *when* disruption occurs, but it would narrow *what* the disruption is likely to look like, transforming the NFT constraint from a prohibition on prediction into a guide for anticipatory modeling. Such an approach could ground empirical research programs in prebiotic chemistry, where the relevant stability landscapes are increasingly measurable. We cannot deduce the organism from the equation because the organism is not a deduction. It is a creative construction, paid for in joules.

References

[1] Wigner, E.P.: The unreasonable effectiveness of mathematics in the natural sciences. Communications on Pure and Applied Mathematics **13**(1), 1–14 (1960)

[2] Kampen, N.G.: Stochastic Processes in Physics and Chemistry, 3rd edn. North-Holland, Amsterdam (2007). https://doi.org/10.1016/B978-044452965-7/50006-4

[3] Gardiner, C.W.: Stochastic Methods: A Handbook for the Natural and Social Sciences, 4th edn. Springer, Berlin (2009). https://doi.org/10.1007/978-3-662-05389-8

[4] Norris, J.R.: Markov Chains. Cambridge University Press, Cambridge, UK (1998)

[5] Pearl, J.: Causality: Models, Reasoning, and Inference, 2nd edn. Cambridge University Press, Cambridge (2009)

[6] Woodward, J.: Making Things Happen: A Theory of Causal Explanation. Oxford University Press, Oxford (2003)

[7] Cartwright, N.: How the Laws of Physics Lie. Oxford University Press, Oxford (1983)

[8] Mitchell, S.D.: Biological Complexity and Integrative Pluralism. Cambridge University Press, Cambridge (2003)

[9] Garte, S., Marshall, P., Kauffman, S.: The reasonable ineffectiveness of mathematics in the biological sciences. Entropy **27**(3) (2025)

[10] Potochnik, A.: Idealization and the Aims of Science. University of Chicago Press, Chicago (2017)

[11] Cartwright, N.: The Dappled World: A Study of the Boundaries of Science. Cambridge University Press, Cambridge (1999)

[12] Batterman, R.W.: The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence. Oxford University Press, Oxford (2002)

[13] Morrison, M.: Unifying Scientific Theories: Physical Concepts and Mathematical Structures. Cambridge University Press, Cambridge (1999)

[14] Adler, D.: Stability-driven assembly theory. Journal of Theoretical Biology **620** (2025) https://doi.org/10.1016/j.jtbi.2025.112352

[15] Adler, D.: How Information Evolves: Stability-Driven Assembly and the Emergence of a Natural Genetic Algorithm. arXiv:2601.17061 (2026)

[16] Goldberg, D.E.: Genetic Algorithms in Search, Optimization, and Machine Learning. Addison-Wesley, Reading, MA (1989)

[17] Lewontin, R.C.: The units of selection. Annual Review of Ecology and Systematics **1**(1), 1–18 (1970)

[18] Godfrey-Smith, P.: Darwinian Populations and Natural Selection. Oxford University Press, Oxford (2009)

[19] Griesemer, J.R.: Reproduction and the reduction of genetics. The Concept of the Gene in Development and Evolution, 240–285 (2000)

[20] Griesemer, J.R.: The informational gene and the substantial body: On the generalization of evolutionary theory by abstraction. Idealization XII: Correcting the Model—Idealization and Abstraction in the Sciences, 59–115 (2005)

[21] Maynard Smith, J., Szathmáry, E.: The Origins of Life: From the Birth of Life to the Origin of Language. Oxford University Press, Oxford, UK (1999)

[22] Laland, K.N., Uller, T., Feldman, M.W., Sterelny, K., Müller, G.B., Moczek, A., Jablonka, E., Odling-Smee, J.: The extended evolutionary synthesis: Its structure, assumptions and predictions. Proceedings of the Royal Society B **282**(1813), 20151019 (2015) https://doi.org/10.1098/rspb.2015.1019

[23] Pigliucci, M., Müller, G.B.: Evolution: The Extended Synthesis. MIT Press, Cambridge, MA (2010)

[24] Wray, G.A., Hoekstra, H.E., Futuyma, D.J., Lenski, R.E., Mackay, T.F.C., Schluter, D., Strassmann, J.E.: Does evolutionary theory need a rethink? Nature **514**, 161–164 (2014) https://doi.org/10.1038/514161a

[25] Turing, A.M.: On computable numbers, with an application to the entscheidungsproblem. Proceedings of the London Mathematical Society **42**(1), 230–265 (1936)

[26] Church, A.: An unsolvable problem of elementary number theory. American Journal of Mathematics **58**(2), 345–363 (1936)

[27] Isaacson, W.: Einstein: His Life and Universe. Simon & Schuster, New York, NY (2007)

[28] Neumann, J.: Theory of Self-Reproducing Automata. University of Illinois Press, Urbana and London (1966)

[29] Fontana, W., Buss, L.W.: "the arrival of the fittest": Toward a theory of biological organization. Bulletin of Mathematical Biology **56**(1), 1–64 (1994) https://doi.org/10.1007/BF02458289

[30] Kleene, S.C.: On notation for ordinal numbers. The Journal of Symbolic Logic **3**(4), 150–155 (1938)

[31] Rogers, H.: Theory of Recursive Functions and Effective Computability. MIT Press, Cambridge, MA (1987). Reprint of the 1967 edition

[32] Kauffman, S.A.: A World Beyond Physics: The Emergence and Evolution of Life. Oxford University Press, New York (2019)

[33] Wolfram, S.: A Project to Find the Fundamental Theory of Physics. Wolfram Media, Champaign, IL (2021)

[34] Kant, I.: Immanuel Kant's Critique of Pure Reason. Macmillan, London (1781)

[35] Windelband, W.: Geschichte und naturwissenschaft. Präludien: Aufsätze und Reden zur Philosophie und ihrer Geschichte (1894). English translation: "History and Natural Science," Theory & Psychology, 8(1), 1998, pp. 5–22

[36] Beatty, J.: The evolutionary contingency thesis. Concepts, Theories, and Rationality in the Biological Sciences, 45–81 (1995)

[37] Rosen, R.: Life Itself: A Comprehensive Inquiry Into the Nature, Origin, and Fabrication of Life. Columbia University Press, New York (1991)

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

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

[40] Eldredge, N., Gould, S.J.: Punctuated equilibria: An alternative to phyletic gradualism, 82–115 (1972)

[41] Schumpeter, J.A.: Capitalism, Socialism and Democracy. Harper & Brothers, New York (1942)

[42] Kuhn, T.S.: The Structure of Scientific Revolutions. University of Chicago Press, Chicago (1962)

[43] Sharma, A., Czégel, D., Lachmann, M., Kempes, C.P., Walker, S.I., Cronin, L.: Assembly theory explains and quantifies selection and evolution. Nature **622**, 321–328 (2023) https://doi.org/10.1038/s41586-023-06600-9

[44] Koza, J.R.: Genetic Programming: On the Programming of Computers by Means of Natural Selection. MIT Press, Cambridge, MA (1992)