

Evolutionary Systems Thinking

From Equilibrium Models to Open-Ended Adaptive Dynamics

Dan Adler

`dan@danadler.com`

Abstract

Complex change is often described as “evolutionary” in economics, policy, and technology, yet most system dynamics models remain constrained to fixed state spaces and equilibrium-seeking behavior. This paper argues that evolutionary dynamics should be treated as a core system-thinking problem rather than as a biological metaphor.

We introduce Stability-Driven Assembly (SDA) as a minimal, non-equilibrium framework in which stochastic interactions combined with differential persistence generate endogenous selection without genes, replication, or predefined fitness functions. In SDA, longer-lived patterns accumulate in the population, biasing future interactions and creating feedback between population composition and system dynamics. This feedback yields fitness-proportional sampling as an emergent property, realizing a natural genetic algorithm driven solely by stability.

Using SDA, we demonstrate why equilibrium-constrained models, even when simulated numerically, cannot exhibit open-ended evolution: evolutionary systems require population-dependent, non-stationary dynamics in which structure and dynamics co-evolve. We conclude by discussing implications for system dynamics, economics, and policy modeling, and outline how agent-based and AI-enabled approaches may support evolutionary models capable of sustained novelty and structural emergence.

Keywords: evolutionary dynamics, system dynamics, non-equilibrium systems, stability-driven assembly, agent-based modeling

1 Introduction

The term *evolution* is widely used in economic, policy, technology, and organizational studies to describe long-term change. Firms are said to evolve, markets evolve, institutions evolve, and strategies evolve. However, in most system dynamics models (Meadows, 2008), the mechanisms underlying this change remain constrained to fixed state spaces, static transition rules, and equilibrium-seeking dynamics. As a result, such models capture adaptation within predefined structures, but struggle to explain how new structures, strategies, or categories emerge over time. Evolution is invoked descriptively but rarely implemented mechanistically (Nelson & Winter, 1982).

This tension reflects a broader issue in scientific explanation. Idealized models (Cartwright, 1983; Mitchell, 2009) often succeed by deliberately removing history, path dependence, and pop-

ulation effects to produce tractable equations. As philosophers of science have emphasized, idealization is not a flaw but a methodological tool (Potochnik, 2017). However, when the phenomena of interest are inherently historical, cumulative, and population-driven, equilibrium-based idealizations can obscure rather than illuminate the causal structure. In such cases, simulation-based approaches can offer superior explanatory power—not because they are more detailed, but because they preserve feedback between system state, population composition, and future dynamics.

Biological evolution (Fisher, 1930) depends on variation and inheritance, but both presuppose molecular replication machinery. Long before such machinery existed, physical and chemical systems already exhibited large differences in persistence: some patterns survived orders of magnitude longer than others under given conditions. Across physics, chemistry, and materials science, stability disparities arise generically from energetic minima, symmetry, spatial organization, and hierarchical assembly, spanning timescales from fleeting transient states to structures that persist for geological durations.

In population-based systems subject to continuous stochastic assembly, differences in persistence necessarily bias population composition. Patterns that last longer accumulate and are more likely to participate in subsequent interactions. This introduces a feedback loop: persistence shapes population statistics, population statistics bias future interactions, and those interactions reshape the distribution of persistence. Selection therefore emerges not as an externally imposed criterion, but as a dynamical consequence of stability-weighted sampling.

In this paper, we introduce *Stability-Driven Assembly* (SDA) (Adler, 2026b) as a minimal framework that makes this feedback explicit. SDA shows that when stochastic assembly is coupled to differential persistence in an open, non-equilibrium system, fitness-proportional sampling arises endogenously. Evolutionary search appears as a natural genetic algorithm (SDA/GA) (Adler, 2026a), driven by stability rather than replication, predefined fitness functions, or fixed transition rates.

This perspective exposes a fundamental limitation of equilibrium-constrained system dynamics models: even when simulated numerically, models with fixed state spaces and transition rules cannot exhibit open-ended evolution. Evolutionary systems instead require that state spaces, interaction structures, and selection pressures co-evolve with the population itself. We conclude by discussing implications for system dynamics, economics, and policy modeling, and by suggesting that agent-based and AI-augmented systems provide a practical substrate for evolutionary models capable of sustained novelty and structural emergence.

2 The Informal Use of “Evolution” in Systems Discourse

Figure 1 illustrates how broadly and persistently the term “evolution” is used across domains. Using Google Trends data over a five-year period, Figure 1(a) shows sustained interest in domain-specific phrases in biological, technological, cultural, linguistic, chemical, and astronomy contexts. Figure 1(b) collapses the same data into five-year averages, revealing that non-biological uses, particularly in technology and astronomy, are comparable in prevalence to biological evolution.




Figure 1: Google Trends interest for domain-specific uses of the term “evolution.” The term is used persistently and comparably across biological, physical, and socio-technical domains, motivating the need for a general mechanistic account of evolutionary dynamics.

Despite this widespread usage, the term is rarely accompanied by a clearly specified dynamical mechanism outside biology: there are no explicit genes, no genotype–phenotype distinction, no mutation or recombination operators, and no inherited replicators that transmit information across generations. Instead, “evolution” functions as a descriptive shorthand for gradual or path-dependent adaptive change in systems. While such usage is often intuitive and rhetorically effective, it leaves unresolved what, if anything, is being selected and how selection is implemented dynamically.

In socio-technical contexts such as technology, culture, language, and economics, evolutionary change is commonly attributed to diffuse processes including competition, learning, imitation, or adaptation to external pressures. These processes can generate historical contingency, but they are rarely formalized in ways that allow population composition to reshape the space of future possibilities. As a result, evolutionary language typically describes outcomes rather than mechanisms, offering narrative coherence without a corresponding population-level dynamical account.

Even in formal scientific and mathematical settings (Nowak, 2006), the evolutionary terminology is frequently used analogously. Evolutionary game theory, for example, adopts the language of strategies, fitness, and selection while generally assuming fixed strategy spaces and static payoff structures. The resulting dynamics describe changes in relative frequencies under predefined rules, not the endogenous emergence of new strategies, representations, or categories.

These usage patterns reveal a consistent gap between the informal evolutionary discourse and formal modeling practice. Evolution is invoked to describe long-term change, yet the underlying models remain constrained to fixed state spaces and equilibrium-seeking dynamics. This gap motivates the need for a general evolutionary systems framework capable of capturing open-ended, non-equilibrium dynamics driven by endogenous population feedback rather than predefined optimization criteria.

3 Stability-Driven Assembly as an Evolutionary Mechanism

Stability-Driven Assembly (SDA) was previously introduced (Adler, 2026b) as a minimal non-equilibrium framework in which differences in persistence bias the accumulation of structure over time, with complete mathematical derivations and symbolic simulations demonstrating entropy reduction and scaffold emergence. A companion paper (Adler, 2026a) extends SDA to chemical symbol space using SMILES molecular fragments, showing hallmark features of evolutionary search like scaffold-level dominance and sustained novelty over thousands of generations. The term “Assembly” here refers to pattern formation through combination, distinct from Assembly Theory (Sharma et al., 2023), which addresses molecular complexity measurement rather than population-level selection. Rather than reproducing those technical results, this section focuses on the core dynamical mechanism: how persistence-weighted feedback reshapes population distributions and induces evolutionary search, treated as a general process applicable across domains.

An SDA system consists of a population of interacting entities generated through stochastic interactions and removed through decay in an open-flow setting. The outcomes of the interaction differ in stability, expressed as a finite persistence time. Patterns that persist longer remain in the population for more generations, become more abundant, and therefore participate more frequently in subsequent interactions. Decay is implicitly modeled by removing expired patterns, while patterns with zero stability are created and eliminated within the same generation, representing interactions that cannot persist under prevailing conditions. This asymmetry in persistence introduces selection without reproduction: persistence biases population composition, population composition biases future interactions, and evolutionary dynamics emerge from this feedback alone—without genes, replication, or predefined fitness functions.

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

```
1: Initialize population with base elements
2: for each generation do
3:   Remove expired patterns
4:   Replenish base elements to maintain non-equilibrium conditions
5:   for each interaction do
6:     Sample two patterns randomly from the current population
7:     Generate a new pattern via combination or recombination
8:     Assign lifetime based on stability and add to population
9:   end for
10: end for
```

Pattern generation occurs through interactions between existing patterns. In its simplest form, SDA employs a concatenation operator, but this can be generalized to recombination with optional mutation. The resulting dynamics are captured by the following schematic loop: expired patterns are removed, base elements are replenished, parent patterns are sampled in proportion to their population frequency, and new patterns are generated and assigned lifetimes based on their stability. This minimal algorithmic structure is sufficient to implement persistence-weighted sampling.




Figure 2: Stability-Driven Assembly (SDA/GA) loop. Base elements are continuously replenished while unstable motifs expire, shaping the active population through differential persistence. Patterns are sampled from the population and combined through interaction (concatenation in SDA or recombination in SDA/GA), generating new motifs that re-enter the population with lifetimes determined by their stability. The resulting feedback from *stability* to *persistence* to *population composition* produces emergent selection without genes, replication, or an explicit fitness function.

Figure 2 illustrates the core feedback underlying the SDA dynamics. Replenishment and expiration define an active population whose composition is shaped by differential persistence. Patterns are randomly sampled from this population for interaction, but because longer-lived patterns accumulate, they are more likely to participate over time. Newly generated motifs inherit lifetimes determined by their stability and re-enter the population, closing the feedback loop. This diagram emphasizes that effective selection arises endogenously from persistence-weighted population dynamics rather than from predefined fitness criteria or externally imposed optimization.

Crucially, the probability that a pattern participates in interactions is proportional to its frequency in the population. Since frequency is itself determined by persistence, stability implicitly defines a selection pressure. This equivalence allows SDA to be interpreted as a *natural genetic algorithm* (SDA/GA) (Adler, 2026a; Goldberg, 1989; Holland, 1975): fitness-proportional selection arises endogenously from persistence differences rather than from explicit optimization criteria or reproduction operators. Variation enters through stochastic generation and recombination, while selection emerges automatically from differential lifetimes.

The resulting population dynamics can be described in terms of an evolving probability distribution over patterns. Two competing processes shape this distribution. A *Create* process introduces novelty through stochastic pattern generation, spreading probability mass across pattern space. A complementary *Persist* process removes patterns according to their lifetimes, biasing the population toward longer-lived motifs. In earlier work (Adler, 2026b), these effects were formalized as distinct creation and persistence contributions to a recursive population update. For the present purposes, the essential point is that persistence induces a systematic bias that accumulates over time, producing path-dependent evolutionary trajectories rather than convergence to a fixed equilibrium.

To clarify the dynamical character of this process, it is useful to move from a discrete, generation-based description to a continuum approximation. When population updates arise from many small stochastic interactions and decay events, the aggregate dynamics of the population distribution can be described in terms of drift and diffusion in pattern space. In classical physics, drift typically represents an externally imposed biasing force such as gravity, an electric field, or a magnetic potential, while diffusion captures random fluctuations. However, in the SDA case, the drift emerges endogenously from persistence-weighted population feedback. In this limit, the evolution of the population distribution takes the form of a Fokker–Planck equation (Gardiner, 2009):

$$\frac{\partial P(x, t)}{\partial t} = \underbrace{-\nabla \cdot [A[P](x, t) P(x, t)]}_{\text{Persist (drift)}} + \underbrace{D \nabla^2 P(x, t)}_{\text{Create (diffusion)}}, \quad (1)$$

where $P(x, t)$ represents the density of motifs in a continuous feature space x . The diffusion term captures the dispersive effect of stochastic pattern creation, while the drift term represents persistence-weighted bias toward regions of higher stability.

The notation $A[P]$ deserves emphasis: the square brackets indicate a *Functional*. The drift field depends not on the local value of P at a point, but on the entire population distribution. As a result, the dynamics of the system is nonlinear, history-dependent, and intrinsically non-equilibrium (Frank, 2005; McKean, 1966). This is not an incomplete specification awaiting parameterization; it is a statement about the causal structure of the system. The effective drift at any point in the pattern space is determined by the patterns that currently exist in the population, how abundant they are, and how their persistence biases the interaction probabilities.

This functional dependence distinguishes SDA from classical transport processes. In standard Fokker-Planck equations describing, for example, Brownian motion in an external potential, the drift $A(x)$ is specified in advance by the physics of the problem (gravity, electric fields, etc.). The equation is linear in P , and stationary distributions can be computed analytically. In SDA, by contrast, the drift *emerges from* the population distribution itself: abundant, long-lived patterns create effective attractors that did not exist before they accumulated. The equation becomes nonlinear and history-dependent, precluding a closed-form solution or a fixed equilibrium.

This is precisely why equilibrium-constrained models cannot capture evolutionary dynamics. Any model that specifies transition rates or drift fields independently of population state assumes

that the dynamical landscape is fixed. SDA demonstrates that in persistence-driven systems, the landscape is endogenous: it is shaped by, and in turn shapes, the population that evolves upon it. Selection is not imposed on the system, but emerges from this recursive coupling between population statistics and effective dynamics.




Figure 3: Core feedback structure underlying Stability-Driven Assembly. Stochastic interactions generate patterns, while differential persistence shapes population composition. The evolving population distribution biases future interactions, closing a feedback loop between persistence, population statistics, and pattern generation. Selection emerges endogenously from this loop without genes, replication, or an externally imposed fitness function.

Figure 3 summarizes the feedback mechanism that drives the evolutionary dynamics in SDA and highlights the difference between persistence-based selection and replication-based selection. In biological evolution, selection operates on individual organisms that replicate: each offspring has specific parents, and lineages can be traced through generations. In SDA, by contrast, patterns are motifs, not individuals. Multiple instances of the same motif in the population arise from distinct interaction events that share structural identity but not lineage. Abundance increases not because a motif copies itself, but because the same structural motif is independently regenerated by multiple interaction pathways and, once present, persists. Thus, selection operates on motifs via differential persistence, not on individuals via differential reproduction.

SDA patterns generated through stochastic interaction differ in persistence, and those that persist longer accumulate within the population. This accumulation reshapes the population distribution, which in turn biases which patterns are likely to participate in future interactions. The feedback loop closes causally: persistence shapes population statistics, population statistics shape

future interactions, and interactions generate new patterns whose persistence further reshapes the population. What appears at the population level as selection is therefore not imposed externally, but emerges from the internal dynamics of persistence-weighted feedback.

3.1 Illustrative Simulation Results

To illustrate the dynamical consequences of stability-driven selection, we simulated a minimal symbolic SDA system under open, non-equilibrium conditions. Patterns were generated stochastically through interaction, assigned lifetimes based on stability, and removed after expiration, while the base elements were continuously replenished (Adler, 2026a). No explicit fitness function, selection rule, or equilibrium constraint was imposed. As a baseline, we compared these dynamics with an unconstrained control in which persistence bias was removed.




Figure 4: Shannon entropy of the population distribution under Stability-Driven Assembly (SDA) compared to an unconstrained control. In the absence of persistence bias, entropy remains high, reflecting broad exploration of pattern space. Under SDA dynamics, entropy decreases steadily as probability mass accumulates on long-lived motifs, demonstrating emergent selection and order in a non-equilibrium system without explicit fitness functions.

Figure 4 shows the evolution of the Shannon entropy of the population distribution over time. The simulation code and extended results are available in the supplementary materials of (Adler, 2026a, 2026b). In the unconstrained control (random search), the entropy remains high, reflecting a broadly uniform exploration of the pattern space. In contrast, under SDA dynamics, entropy decreases steadily, indicating the accumulation of probability mass on a small subset of long-lived motifs. This reduction in entropy demonstrates the emergence of order and effective selection driven solely by differential persistence. Importantly, this occurs without genes, replication, or externally imposed optimization, confirming that persistence-weighted feedback is sufficient to induce evolutionary dynamics in an open, non-equilibrium system.

4 Limits of Equilibrium-Constrained System Models

The results above illustrate a general principle: evolutionary dynamics arise when population composition feeds back on future system behavior through differential persistence. When some patterns persist longer than others, they accumulate in the population and become more likely to participate in subsequent interactions, reshaping the effective dynamics over time. Selection in this sense is not imposed externally or encoded in an explicit fitness function, but emerges endogenously from persistence-weighted population feedback in open, non-equilibrium systems.

An important precursor to this perspective is Brooks and Wiley’s *Evolution as Entropy* (Brooks & Wiley, 1986), which rejected teleological and optimization-based interpretations of evolution in favor of a view emphasizing history, irreversibility, and compatibility with the second law of thermodynamics. By focusing on constraint, path dependence, and information accumulation, Brooks and Wiley reframed evolution as a non-equilibrium process unfolding within an entropic universe. A related shift appears in the econophysics literature, which similarly challenged equilibrium and optimization assumptions by treating economic systems as non-equilibrium populations of interacting entities rather than systems solving well-defined objectives (Farmer & Foley, 2009; Farmer et al., 2005). In both cases, however, evolution is characterized descriptively rather than mechanistically: no general population-level dynamical process is specified by which local order, selection, and cumulative structure arise endogenously.

This gap becomes clear when contrasted with equilibrium-constrained models commonly used in system dynamics, economics, and the physical sciences (Gharajedaghi, 2011; Scheinerman, 2012; Shugar, 1990; Sternberg, 2010). To achieve analytical tractability or numerical stability, such models typically assume fixed state spaces, constant transition rates, and well-mixed dynamics. Under these assumptions, long-term behavior is governed by convergence to steady states, limit cycles, or other attractors determined by externally specified parameters. Even when simulated numerically, these models retain fixed transition rules: all states participate equivalently in a Markovian flow (Norris, 1998), population history is not accumulated, and no state alters the rules governing its own future persistence or interaction.

In contrast, evolutionary systems require that effective dynamics change with population composition. In Stability-Driven Assembly, persistence induces a population-dependent drift that reshapes the probability landscape as the population evolves. What appears as a steady state in equilibrium theory instead corresponds to a transient configuration maintained by ongoing turnover, competition, and selection. From a systems perspective, the key distinction is whether the structure is fixed or subject to selection. Equilibrium-constrained models assume a fixed structure with adaptive behavior unfolding within it; evolutionary systems require that the structure itself be endogenously reshaped by population-level feedback. SDA provides a minimal demonstration of this principle, showing that open-ended evolutionary dynamics can arise without genes, replication, or predefined fitness functions, but not without abandoning equilibrium constraints.

5 Limits of Agent-Based Simulation for Evolutionary Dynamics

5.1 Agent-Based Models Without Population Feedback

Agent-based models (ABMs) (Hamill & Gilbert, 2016; Railsback & Grimm, 2012; Resnick, 1994) are often proposed as a remedy to equilibrium-constrained systems, since they allow heterogeneous agents, nonlinear interactions, and rich emergent behavior. Similarly, recent work in evolutionary economics has sought to formalize population-level dynamics more rigorously (Dopfer & Potts, 2008; Hodgson & Knudsen, 2010). However, both ABMs and these economic frameworks typically retain fixed interaction rules or replication-based inheritance as core mechanisms. Many such systems remain structurally incapable of evolution because they treat population-level structure as an output rather than a causal state.

In typical ABMs, agents interact according to fixed rules within a predefined state space. Population statistics, namely, the frequency of strategies or behaviors, are computed as summaries of agent states, but do not feed back to alter the rules governing future interactions. Such models can exhibit adaptation, learning, or complex transient dynamics, but lack a mechanism by which the population composition reshapes the effective dynamics of the system itself.

The SDA framework clarifies why this distinction is critical. In SDA, the population distribution is an active dynamical variable, not merely an aggregate outcome. Differential persistence changes population composition, which directly biases future interactions. Agents and populations are dynamically inseparable: individual behavior depends on the population, while the population evolves through individual persistence. Evolution arises from this bidirectional coupling, and selection emerges from persistence alone, without requiring that successful entities reproduce or transmit information to descendants.

This perspective explains why many agent-based approaches produce rich transient behavior but converge to fixed regimes or repeating cycles. Without population-level feedback that alters the effective transition structure, systems can adapt within a given space of possibilities, but cannot generate new ones. SDA provides a minimal example of how population-dependent feedback transforms simulation into genuine evolutionary dynamics.

5.2 Stochastic Interaction as a Systems-Level Primitive

In many systems, stochasticity arises from chance encounters. This is obvious in physical systems, but it applies equally in social and organizational contexts: conversations at conferences, informal introductions, and unexpected collaborations are not anomalies, but the primary way new connections form.

Whether these interactions matter in the long run depends on what persists afterward. Most interactions dissipate quickly, leaving no lasting trace. However, on occasion, an interaction produces sustained interest, durable collaboration, or a reusable organizational form. When such outcomes persist, they remain available for further interaction and elaboration. Persistence therefore acts as a filter: stochastic exposure generates variation, while differential longevity determines which

variants accumulate.

As persistent outcomes accumulate, they alter the landscape of future interactions. A collaboration that endures becomes visible to others, attracts additional participants, and is more likely to be encountered again, whether through reputation, shared infrastructure, or institutional embedding. What begins as a chance event can thus become a trend, not because it was designed to spread, but because its persistence increases the probability of further engagement. Frequency emerges as a consequence of durability, not of centralized coordination.

This process closely mirrors the dynamics formalized in SDA. Interactions are sampled stochastically from the population of existing entities, but entities that persist longer are sampled more often simply because they remain present. Selection arises not from intentional choice or optimization, but from the accumulation of exposure opportunities created by persistence. Over time, stable structures dominate not because they are optimal, but because they continue to participate in the dynamics of the system.

From a systems perspective, stochastic interaction combined with persistence-driven accumulation provides a natural explanation for how large-scale structure can emerge from local, contingent events. It also clarifies why deterministic interaction rules or equilibrium assumptions obscure evolutionary dynamics: by eliminating contingent exposure, they suppress the very mechanism by which novelty can enter and be amplified. In evolutionary systems, randomness is not noise to be averaged away, but the source of exploration that persistence converts into structure.

A possible objection to the claim that equilibrium-constrained models cannot capture evolution is that many well-known systems dynamics models exhibit rich non-steady behavior. Peter Senge (Senge, 1990) famously demonstrates through the Beer Game that simple stock-flow structures can generate persistent oscillations, delays, and amplification effects. Similar behaviors arise in physical systems, electrical circuits, ecological models, and business cycles. These dynamics are often described as non-equilibrium and are sometimes taken as evidence that such models already capture evolutionary change.

However, oscillation is not evolution. An oscillatory system remains confined to a fixed set of states and governed by invariant transition rules. Whether the system converges to a steady point, a limit cycle, or even a chaotic attractor, its long-term behavior is restricted to a predefined phase space. The structure of the system, the identities of its stocks, the form of its feedback loops, and the rules that govern interaction do not change over time. In this sense, the oscillation represents a dynamic equilibrium: the system continues to move, but does not transform.

The Beer Game illustrates this distinction clearly. Inventory levels, order rates, and delays fluctuate dramatically, yet the underlying entities and relationships remain unchanged. No new stocks appear, no interaction rules are altered, and no strategies or organizational forms emerge or disappear. The system learns nothing in a structural sense; it merely cycles through the consequences of its initial design. History affects state, but not structure.

Evolutionary dynamics require a stronger condition. For a system to evolve, its population composition must feed back on the rules that generate future behavior. New entities or strate-

gies must be able to arise, persist, and alter the space of subsequent possibilities. Without such population-dependent feedback, even highly nonlinear or oscillatory systems remain dynamically interesting but evolutionarily inert. Stability-Driven Assembly makes this distinction explicit: selection arises not from oscillation or instability, but from persistence-weighted feedback that allows structure itself to be reshaped over time.

6 Adding Evolutionary Dynamics to Stocks-and-Flows Models

One of the enduring strengths of system dynamics is the unifying power of the stocks-and-flows formalism (García, 2019; Meadows, 2008). Bank balances, inventories, population counts, chemical concentrations, and water in a reservoir can all be modeled within the same mathematical framework, often reducing to first-order differential equations. This abstraction is enormously successful precisely because it allows diverse phenomena to be analyzed using a common language of accumulation, inflow, and outflow.

However, this success also encourages a subtle but consequential extrapolation. The same stock–flow intuition is frequently applied to systems that do not merely adjust within a fixed structure but evolve by generating new structures, interactions, or entities over time. When this occurs, the mathematical convenience of fixed stocks and flows can obscure a fundamental change in the nature of the system being modeled.

A canonical example arises in chemistry (Turányi & Tomlin, 2014). Chemical kinetics and reaction network models typically describe systems in terms of concentrations, stoichiometric coefficients, and fixed reaction pathways. Species are treated as static stocks whose quantities change according to predefined reaction rates, and long-term behavior is analyzed in terms of equilibrium or steady-state concentrations. This framework is powerful and appropriate for many purposes. However, it implicitly assumes that the chemical environment itself, the space of possible compounds, interactions, and pathways, remains fixed.

As shown in the SDA/GA framework (Adler, 2026a, 2026b), this assumption breaks down in open, non-equilibrium chemical systems where new compounds can form, persist, and participate in subsequent reactions. In such settings, the relevant stocks are not merely concentrations of predefined species, but populations of assemblies whose identities, lifetimes, and interaction possibilities evolve over time. Modeling only concentrations erases persistence differences and population history, suppressing the very feedback that enables cumulative chemical evolution.

This limitation is not limited to chemistry. Similar stock–flow abstractions are routinely carried over into biology, psychology, economics, and sociology, where populations of organisms, behaviors, strategies, or institutions are modeled as continuous quantities evolving under fixed rules. Although such models can capture adaptation within a given structure, they struggle to represent the emergence of new forms, strategies, or organizational categories. The issue is not a lack of realism or detail, but a mismatch between the mathematical structure of the model and the evolutionary nature of the system.

In an SDA system, stocks represent populations of patterns rather than quantities of conserved substances. Flows represent stochastic interactions that generate new patterns and decay processes that remove them. Crucially, the stock structure itself is dynamic: new stocks appear as novel patterns emerge, and existing stocks disappear when patterns fail to persist. Differential persistence acts as a stock-specific outflow rate, biasing population composition over time, while frequency-weighted interaction flows create reinforcing feedback that amplifies durable structures.

Viewed through this lens, evolutionary dynamics cannot be captured by a single first-order differential equation with fixed state variables. They require models in which the set of stocks, the connectivity of flows, and the effective dynamics co-evolve with the population. Stability-Driven Assembly provides a minimal example of how such dynamics can be implemented within a stocks-and-flows paradigm, extending system dynamics from equilibrium-seeking adjustment to open-ended structural evolution.

7 Toward Open-Ended Evolutionary System Models

The preceding analysis suggests that capturing evolutionary dynamics in systems models requires more than simulating fixed equations far from equilibrium. What distinguishes evolutionary systems is not complexity alone, but feedback that allows population structure to reshape the dynamics that generate it. In such systems, the effective state space, interaction rules, and selection pressures evolve together over time. Models that hold these elements fixed can exhibit adaptation, but not open-ended evolution.

From a modeling perspective, this requires treating population composition as a causal state rather than a descriptive outcome. Evolutionary models must allow the entities populating the system to influence not only aggregate variables, but also the rules governing future interactions and persistence. SDA provides a minimal illustration of this principle: differential persistence reshapes population composition, population composition biases future interactions, and this feedback continually alters the effective dynamical landscape.

Agent-based modeling offers a natural substrate for implementing such dynamics. Agents can represent patterns, strategies, or assemblies whose interactions and lifetimes depend on their internal structure and environment. When embedded in open, replenished settings with stochastic interactions, differential persistence alone can generate selection-like dynamics without explicit optimization objectives. Crucially, agent-based models can support expanding state spaces, allowing novel behaviors or representations to emerge that were not specified at initialization.

Recent advances in artificial intelligence further extend this design space. Learning agents can modify their internal representations or interaction strategies based on experience, effectively changing the space of possible future behaviors. When combined with population-level persistence-driven selection, such systems may support genuinely open-ended evolutionary learning. In this context, AI functions not as an optimizer, but as a flexible substrate for variation and interaction.

The central challenge, however, is not adding learning or complexity, but preserving evolutionary

feedback. Without mechanisms that couple persistence to population-level bias, even sophisticated learning agents tend to converge to fixed equilibria or cycle within predefined solution spaces. Open-ended evolution requires that success alter the future search process itself. Systems models that satisfy this criterion blur the boundary between dynamics and structure, extending system dynamics from equilibrium analysis toward frameworks in which structure itself is subject to selection.

8 Implications for Economics and Policy Modeling

8.1 Illustrative Example: Industry Ecosystems

The SDA framework can be interpreted in economic contexts as a model of evolving industry ecosystems. Consider a simplified system with manufacturers (A), retailers (B), and logistics providers (C) as base elements. Through interaction and coordination, these entities form composite organizational structures such as manufacturer-retailer partnerships (AB) or integrated supply networks (ABBC). Different configurations exhibit different degrees of stability, reflecting factors such as profitability, resilience, and institutional durability.




Figure 5: Conceptual illustration of an industry ecosystem modeled as a Stability-Driven Assembly (SDA) process. Base economic entities (e.g., manufacturers, retailers, logistics providers) interact to form composite organizational structures. Configurations that persist longer under competitive and institutional pressures accumulate resources and participation, biasing future interactions. Over time, differential persistence shapes the population of organizational forms, producing evolutionary dynamics without centralized optimization or explicit replication.

More stable configurations persist longer and therefore accumulate resources, attention, and participation, increasing their likelihood of engaging in further interactions. Less viable arrangements dissolve more quickly, redistributing resources back into the system. Over time, this persistence bias skews the population toward durable industry structures, which function as attractors without requiring centralized optimization or ex ante design.

This example illustrates how stability-driven selection can operate in socio-economic systems

through feedback between persistence, population composition, and future interactions. From a systems perspective, the key driver is not rational optimization at the firm level, but population-level feedback that amplifies durable organizational adaptation over time.

8.2 Implications for Economic and Policy Systems

Economic and policy systems are routinely described as evolutionary: firms adapt, technologies diffuse, institutions evolve, and strategies compete over time. However, most formal models in economics and policy analysis remain equilibrium-centered or rely on fixed behavioral rules, even when implemented in dynamic or agent-based form. As a result, such models can represent adjustment within predefined structures, but struggle to capture the emergence of new strategies, institutional forms, or categories of behavior.

The SDA framework suggests a different modeling emphasis. Rather than specifying selection criteria or optimization objectives, evolutionary change can arise from differential persistence under continual turnover. In economic contexts, persistence can correspond to the durability of business models, regulatory arrangements, norms, or technologies under changing conditions. Entities that persist longer naturally accumulate influence, visibility, or market share, biasing future interactions, and shaping the environment in which new variants arise.

This perspective clarifies key limitations of equilibrium-based policy analysis. Policies are often evaluated by comparing steady states before and after an intervention, implicitly assuming that the system structure remains fixed. In practice, however, policy interventions frequently alter the space of viable strategies and institutions, triggering long-term evolutionary effects that are invisible to equilibrium comparisons. Persistence-driven dynamics imply that small changes affecting survival or turnover rates can have outsized consequences by reshaping population composition and future trajectories.

Incorporating evolutionary feedback into economic and policy models does not require abandoning the principles of system dynamics, but it does require extending them. Models must allow population composition to influence future dynamics endogenously, rather than treating behavior and structure as static. Hybrid approaches that combine stock-flow models with agent-based or population-based representations offer a promising path forward, in which stocks represent populations of strategies or institutions, flows represent entry and exit governed by persistence, and feedback loops capture how accumulated structures bias future behavior.

Finally, the SDA perspective highlights a shift in policy design from optimization to resilience. Rather than attempting to engineer optimal outcomes in complex adaptive systems, policy can be better understood as shaping persistence landscapes—altering which behaviors, institutions, or technologies are able to survive and propagate. From this point of view, effective policy is less about directing outcomes and more about creating conditions under which desirable structures can persist and evolve over time.

9 Conclusion

This paper has argued that many systems described as evolutionary fail to evolve in a formal sense because the models used to represent them are structurally constrained toward equilibrium. When state spaces, transition rules, and selection criteria are fixed in advance, long-term changes can be simulated, but not generated. Evolution, understood as the cumulative and path-dependent transformation of both populations and the structures that shape them, requires feedback in which persistence biases future dynamics.

SDA provides a minimal generative demonstration of this feedback. By coupling stochastic generation with differential persistence in an open, non-equilibrium system, SDA shows how selection-like behavior emerges without genes, replication, or predefined fitness functions. Persistence reshapes population composition, population composition biases future interactions, and this feedback produces sustained entropy reduction and structured population-level order. Interpreted dynamically, SDA corresponds to a population-dependent nonlinear drift process that lies outside the scope of equilibrium-constrained models.

The broader implication is methodological rather than domain-specific. Evolutionary systems are best explored not through single trajectories or equilibrium outcomes, but through ensembles of stochastic realizations, analogous to Monte Carlo analysis, in which the advantage is inferred statistically from differential persistence across histories. SDA does not attempt to explain how stability functions arise, how environments change, or how hierarchical organization emerges; rather, it isolates the dynamical consequences of persistence-driven feedback. Extending system dynamics to incorporate this feedback shifts the focus from optimization to robustness, survivability, and path-dependent dominance.

Viewed in this way, evolution is not a metaphor imported from biology, but a general class of dynamical behavior that arises under specific structural conditions. Identifying and modeling those conditions may be essential for understanding complex adaptive systems in economics, policy, and beyond.

References

- Adler, D. (2026a). How information evolves: Stability-driven assembly and the emergence of a natural genetic algorithm. <https://arxiv.org/abs/2601.17061>
- Adler, D. (2026b). Stability-driven assembly theory. *Journal of Theoretical Biology*, 620, 112352. <https://doi.org/10.1016/j.jtbi.2025.112352>
- Brooks, D. R., & Wiley, E. O. (1986). *Evolution as entropy: Toward a unified theory of biology*. University of Chicago Press.
- Cartwright, N. (1983). *How the laws of physics lie*. Oxford University Press.
- Dopfer, K., & Potts, J. (2008). *The general theory of economic evolution*. Routledge.
- Farmer, J. D., & Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460(7256), 685–686. <https://doi.org/10.1038/460685a>

- Farmer, J. D., Shubik, M., & Smith, E. (2005). Is economics the next physical science? *Physics Today*, 58(9), 37–42. <https://doi.org/10.1063/1.2123803>
- Fisher, R. A. (1930). *The genetical theory of natural selection*. Oxford University Press.
- Frank, T. D. (2005). *Nonlinear fokker-planck equations: Fundamentals and applications*. Springer. <https://doi.org/10.1007/3-540-27771-9>
- García, J. M. (2019). *Theory and practical exercises of system dynamics*. Springer.
- Gardiner, C. W. (2009). *Stochastic methods: A handbook for the natural and social sciences* (4th ed.). Springer.
- Gharajedaghi, J. (2011). *Systems thinking: Managing chaos and complexity* (3rd). Morgan Kaufmann.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley.
- Hamill, L., & Gilbert, N. (2016). *Agent-based modelling in economics*. Wiley.
- Hodgson, G. M., & Knudsen, T. (2010). *Darwin’s conjecture: The search for general principles of social and economic evolution*. University of Chicago Press.
- Holland, J. H. (1975). *Adaptation in natural and artificial systems*. University of Michigan Press.
- McKean, H. P. (1966). A class of markov processes associated with nonlinear parabolic equations. *Proceedings of the National Academy of Sciences*, 56(6), 1907–1911. <https://doi.org/10.1073/pnas.56.6.1907>
- Meadows, D. H. (2008). *Thinking in systems: A primer*. Chelsea Green Publishing.
- Mitchell, M. (2009). *Complexity: A guided tour*. Oxford University Press.
- Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of economic change*. Harvard University Press.
- Norris, J. R. (1998). *Markov chains*. Cambridge University Press.
- Nowak, M. A. (2006). *Evolutionary dynamics: Exploring the equations of life*. Belknap Press.
- Potochnik, A. (2017). *Idealization and the aims of science*. University of Chicago Press.
- Railsback, S. F., & Grimm, V. (2012). *Agent-based and individual-based modeling: A practical introduction*. Princeton University Press.
- Resnick, M. (1994). *Turtles, termites, and traffic jams: Explorations in massively parallel microworlds*. MIT Press.
- Scheinerman, E. R. (2012). *Invitation to dynamical systems*. Pearson.
- Senge, P. M. (1990). *The fifth discipline: The art and practice of the learning organization*. Doubleday.
- Sharma, A., Czégel, D., Lachmann, M., Kempes, C. P., Walker, S. I., & Cronin, L. (2023). Assembly theory explains and quantifies selection and evolution. *Nature*, 622, 321–328. <https://doi.org/10.1038/s41586-023-06600-9>
- Shugar, D. (1990). *The science of systems*. Oxford University Press.
- Sternberg, S. (2010). *Dynamical systems*. Dover Publications.
- Turányi, T., & Tomlin, A. S. (2014). *Analysis of kinetic reaction mechanisms*. Springer.