The Problem of Persistence

Alyx Dubrow

Abstract

Highly organized, hierarchically structured systems are rare in a universe governed by physical processes that tend to disperse and degrade fine-grained organization. Yet such systems: living organisms, nervous systems, and technological societies, exist and endure. Their existence presents a puzzle that is often overlooked. Deep complexity cannot arise instantaneously; it must accumulate structure over extended periods. Over such timescales, however, environmental fluctuations are far more likely to disrupt organized configurations than to enhance them. The continued presence of multi-level organization therefore calls for explanation. This paper argues that the long-term persistence of complex systems in a changing universe depends on adaptive mechanisms that track environmental regularities. Two broad classes of mechanisms are identified. The first operates across populations and generations: replication with variation and differential survival, the Darwinian process of evolution. The second operates within individual systems: the construction and continual updating of internal models that allow possible futures to be simulated and actions to be selected accordingly, a capacity that provides a functional characterization of intelligence. Although typically studied in separate domains, these processes share a common adaptive logic of variation, evaluation under environmental constraint, and retention of effective structure. From this perspective, evolution and intelligence are not isolated biological or cognitive phenomena but scale-dependent implementations of mechanisms that make long-term persistence of complex organization possible. Life and mind thus appear not as anomalies superimposed on an indifferent physical world, but as natural outcomes of the conditions required for highly organized systems to endure over time in the face of continual change.

The Problem of Persistence

1. Introduction

The universe evolves through time. Physical conditions shift over time, energy gradients dissipate, and interactions among components continually disturb existing structures. Against this background of ongoing change, the existence of long-lived, highly organized systems is not something that can be taken for granted. Complex structures: organisms, ecosystems, brains, and civilizations form and then persist across changing conditions. Their persistence across changing conditions calls for explanation.

Much scientific and philosophical attention has been devoted to specific kinds of complex systems, especially living organisms and intelligent agents. Yet a more basic question lies beneath these familiar topics: how can any deeply structured system persist at all in a changing universe? Most physical processes tend to disperse, mix, or degrade organized states. The continued existence of multi-layered, functionally integrated structures is therefore not a default outcome of physical processes, but something that requires explanation.

This paper approaches life, evolution, and intelligence from this more general standpoint. Instead of beginning with biological or cognitive definitions, it begins with the problem of persistence itself. Any complex system we are able to observe has already survived a history of environmental variation. Its present existence is evidence that, in some way, it has managed to counteract the destabilizing effects of change long enough for complex organization to accumulate. This observation suggests a unifying question: what kinds of mechanisms make such persistence possible?

I argue that, in a time-evolving universe with incomplete information about the future, long-term persistence of complex systems requires adaptive strategies that effectively anticipate and compensate for environmental change. These strategies fall into two broad classes. The first operates at the level of populations across generations: replication with variation, coupled with differential survival and reproduction. This is the familiar Darwinian mechanism of evolution, here understood as a general process by which populations track environmental regularities over time. The second operates within individual systems: the construction and continual updating of internal models that allow possible futures to be simulated and actions to be selected accordingly. This model-based anticipatory capacity provides a functional characterization of intelligence.

Although these two strategies are often treated separately, one as the domain of evolutionary biology and the other as the domain of cognitive science, I suggest that they can be understood as different implementations of a common underlying logic. Both are ways in which systems build compressed representations of environmental structure and use variation and selection to improve their fit to a changing world. From this perspective, evolution and intelligence are not isolated phenomena but scale-dependent solutions to a single physical problem: persistence under change.

The goal of this paper is to articulate this general framework and explore its implications. I begin by clarifying the kind of complexity at issue and by distinguishing persistence in stable physical regimes from persistence across variable conditions. I then develop the two adaptive strategies in more detail and examine their structural parallels. Finally, I consider several objections and briefly explore the broader consequences of this perspective, including its relevance to anthropic reasoning and to philosophical questions about mind and agency.

2. From Physics to Persistence: A Generalized Anthropic Step

The argument developed in this paper does not begin from the assumption that complex systems must arise everywhere or at all times. On the contrary, much of the universe appears hostile to long-lived, highly organized structure. Stars burn out, planetary environments shift, and many physical processes rapidly erase fine-grained organization. It is therefore entirely consistent with current physical understanding that large regions of the universe, or many possible universes, contain little or no deep hierarchical complexity.

This observation motivates an important selection effect. The only complex systems available for analysis are those that have already persisted for some nontrivial span of time. We do not observe fragile, short-lived configurations that immediately dissolve under environmental fluctuations, except perhaps as transient phenomena. Instead, the systems that attract scientific and philosophical interest: living organisms, ecological networks, nervous systems, and technological civilizations are precisely those that have maintained organized structure across changing conditions.

This reasoning parallels anthropic arguments, but with a broader and more physical emphasis. Traditional formulations of the anthropic principle focus on the conditions necessary for observers. Here the relevant filter is more general: we can only encounter persistent complex systems in regions of reality where persistence of such systems is possible. The existence of any such system is therefore evidence that its environment, and the laws governing it, permit mechanisms capable of counteracting destabilizing change.

This does not imply that complexity, life, or intelligence are inevitable outcomes of physics in every context. Rather, it establishes a conditional framework. Given that a complex system exists and has persisted long enough for organized structure to accumulate, its continued presence is not a brute fact but the result of processes that have enabled it to withstand, compensate for, or otherwise remain viable under environmental variation.

This perspective shifts the explanatory burden. Rather than focusing on the historical circumstances that led to the existence of a particular complex system, we ask what general classes of mechanisms make the long-term persistence of complex organization possible in a changing universe. The following sections develop the claim that there are only a limited number of such mechanisms, and that the familiar phenomena of biological evolution and intelligent behavior can be understood as two major implementations of this more general solution to the problem of persistence.

3. What Kind of Complexity Is at Issue?

The term complexity is used in many different and often incompatible ways across disciplines. In information theory, a sequence may be called complex if it is incompressible or random. In statistical mechanics, high entropy states may be described as complex because they correspond to a large number of possible micro-configurations. While these notions are precise and useful in their respective domains, they do not capture the kind of complexity relevant to the problem of persistence addressed here.

The focus of this paper is on organized, hierarchical complexity: systems composed of many interacting parts arranged in multiple levels of structure, where higher-level patterns both emerge from and constrain lower-level dynamics. Such systems typically exhibit functional integration, historical layering, and emergent regularities that are not apparent at the level of their individual components.

Examples include living cells, multicellular organisms, nervous systems, ecosystems, and technological societies. In each case, the system's organization is not a simple repetition of a basic pattern, as in a crystal, nor a maximally disordered configuration, as in a gas at equilibrium. Instead, it consists of differentiated components coordinated into larger wholes, often with specialized subsystems and feedback loops operating at multiple scales.

A central feature of this form of complexity is its historical depth. Highly organized systems are rarely assembled in a single step. Rather, they accumulate structure over time, with relatively stable intermediate configurations serving as platforms for further elaboration. This layered construction makes such systems powerful and versatile, but also potentially fragile. Disruption at one level can propagate to others, and the maintenance of higher-level organization depends on the continued integrity of lower-level processes.

This conception of complexity also helps clarify why persistence is nontrivial. The more levels of organization a system contains, and the more tightly coordinated its components are, the more ways there are for environmental change to disrupt its structure. Organized complexity occupies a comparatively small region of the space of all possible configurations. Random perturbations are therefore more likely to degrade such organization than to enhance it.

For these reasons, the persistence problem addressed in this paper concerns not complexity in the sense of randomness or mere multiplicity of parts, but the long-term maintenance of hierarchically organized, functionally integrated structure. It is this kind of complexity that accumulates over time, supports adaptive behavior, and gives rise to the phenomena we associate with life and intelligence. The question is how such systems manage to endure at all in the face of continual environmental change.

4. Why Persistence Is Nontrivial in a Changing Universe

Once complexity is understood as hierarchically organized, functionally integrated structure, its persistence in a changing environment can no longer be treated as the default outcome of physical processes. On the contrary, such organization is typically fragile. Most environmental changes do not preserve fine-tuned arrangements of components; they disrupt them.

One way to see this is combinatorial. For any system composed of many interacting parts, the number of possible configurations grows astronomically with the number of degrees of freedom. Only a very small subset of these configurations correspond to coherent, functionally integrated structures. The overwhelming majority correspond to disordered, uncoordinated arrangements. Random perturbations are therefore far more likely to move a system away from an organized configuration than toward one.

Biological mutation provides a familiar example. Genetic variation is essential for evolution, yet most random mutations are neutral at best and harmful at worst. Only a small fraction improve an organism's fit to its environment. The fact that beneficial variations are rare does not undermine evolution; it is precisely what makes selection necessary. Without a filtering process, accumulated random changes would rapidly erode functional organization.

A similar asymmetry applies more generally. Environmental fluctuations include changes in temperature, resource availability, chemical composition, or the behavior of other systems, none of which are tailored to preserve any particular structure. They act as perturbations drawn from a vast space of possibilities, most of which do not align with the narrow requirements for maintaining a given organized state.

These observations can be understood as a consequence of the second law of thermodynamics in a broad sense. While local decreases in entropy are possible in open systems driven by energy flows, the maintenance of low-entropy, highly structured states requires continual work. Absent mechanisms that actively counteract degradation, organized structures tend to dissolve into more probable, less constrained configurations over time.

The persistence of deeply structured systems therefore demands explanation. It is not enough that such systems form; they must also withstand an ongoing stream of perturbations drawn from a space in which disorder vastly outnumbers order. This asymmetry between the space of possible disruptions and the narrow set of configurations compatible with continued organization is what makes the problem of persistence both general and profound.

In the next section, I argue that there are only a limited number of ways in which systems can meet this challenge. These involve not merely passive resistance to change, but active processes that track, anticipate, and adapt to environmental variation.

5. Two General Solutions to the Persistence Problem

If organized, hierarchical systems are fragile under arbitrary environmental change, then their continued existence requires more than passive stability. Persistence across variable conditions demands mechanisms that allow a system to track, anticipate, and adapt to the structure of its environment. I argue that there are two broad classes of such mechanisms. Although they are usually studied in different scientific domains, they can be understood as alternative implementations of a common adaptive logic.

5.1 Population-Level Adaptation: The Darwinian Strategy

The first strategy operates at the level of populations across generations. Systems reproduce with variation, and their variants differ in how well they maintain organization under prevailing conditions. The environment then acts as a filter: variants that better fit environmental constraints are more likely to persist and reproduce, while poorly matched variants are eliminated.

At this level, persistence does not mean the continued existence of a particular organism, but the continuation of an organized form across generations, as individuals survive long enough to reproduce and transmit that organization forward in time.

This process, replication with variation and differential survival, is the core of Darwinian evolution. Understood in the present framework, evolution functions as a distributed, long-timescale learning process. The genome can be seen as a compressed record of environmental regularities accumulated through past selection. Each generation constitutes a new round of hypothesis testing, in which genetic variations are evaluated against the environment. Successful structures are retained and recombined, while unsuccessful ones are pruned away.

This mechanism does not require foresight within any individual organism. Prediction is implicit and statistical, embodied in the changing distribution of traits within the population. Over many generations, populations come to "anticipate" recurring features of their environment through the structures encoded in their genomes. In this way, Darwinian processes allow populations to track environmental regularities even in the face of ongoing change.

5.2 Individual-Level Adaptation: The Predictive Strategy

A second strategy shifts adaptation from the level of populations to that of individual systems. Instead of relying solely on generational turnover, some systems construct and maintain internal models that allow them to anticipate environmental change within their own lifetimes. These systems do not merely react to perturbations; they simulate possible futures and select actions accordingly.

For such predictive adaptation to be possible, a system must build a model that captures relevant regularities in its environment. A central challenge is how such a model can be constructed and maintained as input continually changes. Storing each experience independently would lead to an unmanageable proliferation of unconnected details. Instead, effective models arise through the reuse of structure across contexts. When patterns recur in different situations, systems can abstract and stabilize those regularities, integrating them into compact representations that support generalization and prediction.

This process need not begin with fully formed abstractions. Early in learning, systems may encode relatively specific exemplars tied to particular contexts. As experience accumulates, patterns in the environment are encountered repeatedly, often in slightly different forms. Components of these patterns that recur across situations are reused, reinforced, and gradually stabilized, while idiosyncratic details fade. Over time, this iterative process of exposure, reuse, and reinforcement allows shared structure to stand out against background variation. In this way, a predictive world model can be bootstrapped from initially local and concrete representations into a more abstract and flexible form, with abstraction emerging as the residue of repeated reuse.

Abstraction, on this view, is not constructed so much as filtered. It consists of the features of experience that persist across variation, what remains invariant as contexts change. In this sense, abstraction emerges through a process more akin to sieving than design: structure is retained not because it is chosen, but because it survives repeated reuse.

Such dynamics are not merely theoretical possibilities. In biological nervous systems, learning mechanisms reinforce patterns of activity that repeatedly co-occur, allowing stable features and relationships to be extracted from streams of sensory input. In artificial learning systems, including deep neural networks, internal representations likewise emerge through the reuse of learned features across many examples and tasks. In both cases, reuse enables the system to construct a model that is not a mere record of past inputs, but a structured compression of environmental regularities.

Once such a model exists, it can be used to generate possible futures internally. By simulating the likely consequences of different actions, the system can select behaviors that help maintain its organization under changing conditions. This capacity to build and update a reusable world model, to generate counterfactual scenarios, and to choose actions on the basis of predicted outcomes provides a functional characterization of intelligence within the present framework.

5.3 Prediction at Two Scales

Although the Darwinian and predictive strategies are often treated separately, one as the domain of evolutionary biology and the other as the domain of cognitive science, they share a deep structural similarity. Both rely on variation, evaluation against environmental constraints, and retention of structures that prove effective. In the evolutionary case, these processes unfold across populations and generations. In the predictive case, they are internalized within individual systems and unfold across moments of experience and action.

From the standpoint of persistence, these are not fundamentally different kinds of processes, but different scales and implementations of a common adaptive logic. Both enable systems to track environmental regularities and adjust their organization accordingly. Both allow structured models of the world, whether encoded in genomes or internal representations, to be refined over time through a process analogous to hypothesis testing and selection.

Systems that possess neither of these strategies may exist briefly, especially in stable physical regimes. However, they lack general means of compensating for environmental variation and are therefore unlikely to accumulate deep, multi-layered organization over extended periods. The persistence of complex systems in a changing universe thus points toward these adaptive mechanisms as central features of their continued existence.

6. Evolution and Intelligence as the Same Adaptive Logic at Different Scales

The preceding discussion has treated population-level evolution and individual-level predictive modeling as two distinct strategies for persistence. However, their relationship is deeper than a mere similarity of function. They can be understood as different implementations, at different scales and timescales, of a common underlying adaptive logic.

Both processes involve the construction of a model that captures regularities in the environment. In evolutionary systems, this model is distributed across the genomes of a population. Genetic variation generates alternative "hypotheses" about how to maintain organization under prevailing conditions. Environmental interaction serves as a test, and differential reproduction retains variants that prove more effective. Over generations, the population's genetic structure comes to reflect, in compressed form, the statistical regularities of its environment.

In predictive organisms, a comparable process unfolds within the lifetime of a single system. Internal representations encode regularities extracted from experience. Alternative possible actions or internal states can be generated and evaluated through simulation, and outcomes that support continued viability are reinforced. Here, the model is maintained in neural or computational structure rather than in a gene pool, and selection operates over simulated or behavioral outcomes rather than over differential reproduction. Yet the abstract pattern: variation, evaluation under environmental constraint, and retention of effective structure remains the same.

The primary differences between these implementations concern integration and timescale. Evolutionary adaptation is distributed across many individuals and unfolds over generations. It does not involve a single, temporally unified system maintaining a moment-to-moment model of the world. Predictive cognition, by contrast, is localized within individual organisms and operates on much shorter timescales. It supports real-time model updating, counterfactual simulation, and action selection within a single integrated system.

From the standpoint of persistence, these differences are secondary to the shared logic. Both strategies enable systems to track environmental structure and adjust their organization accordingly. Both allow complex organization to be refined through iterative interaction with a changing world. In this sense, predictive intelligence can be viewed as the internalization and acceleration of an evolutionary process, bringing model construction and selection into the lifetime of an individual system.

Recognizing this structural continuity helps to dissolve the apparent divide between evolutionary and cognitive explanations. They are not competing accounts of adaptation but complementary manifestations of the same general solution to the problem of persistence. Whether implemented across populations and generations or within individuals and moments, adaptive model-building emerges as a central principle underlying the endurance of complex systems in a changing universe.

7. Objections and Clarifications

A framework this general invites understandable skepticism. The goal of this section is not to defend every detail, but to clarify the scope of the claims and address several natural objections.

7.1 Is This Merely Tautological?

One might worry that the central thesis reduces to a truism: systems that persist must have mechanisms that allow them to persist. On this reading, the argument would amount to little more than a relabeling of survival.

However, the claim advanced here is not definitional but mechanistic. It does not assert merely that persistent systems have persistence-enabling properties; it proposes that, in a changing universe with entropy-driven degradation of structure, there are only limited classes of processes capable of sustaining deep hierarchical organization over time. Replication with selection and predictive modeling are offered as two such general strategies. The thesis therefore constrains the space of viable mechanisms rather than restating the fact of survival.

A useful comparison is with the claim that any flying system must generate lift. This is not a tautology about flight but a physical constraint on how flight can occur. Similarly, the present argument seeks to identify the structural requirements for persistence under environmental change.

7.2 What About Crystals, Stars, and Other Nonliving Structures?

Many physical systems exhibit organized structure without replication or internal modeling. Crystals, flames, convection cells, and stars can maintain coherent patterns for extended periods. Do such cases undermine the argument?

These systems illustrate an important distinction between passive persistence in stable regimes and adaptive persistence across variable conditions. A crystal maintains its structure because the surrounding physical conditions continually favor the same lattice arrangement. A star remains stable as long as the balance between gravitational collapse and nuclear fusion holds. These are examples of structures that are continuously regenerated by underlying physical processes within relatively narrow parameter ranges.

Such systems lack general mechanisms for coping with significant environmental change. When conditions move outside the regime that supports them, they do not adapt; they simply dissolve or transform into different structures. By contrast, the focus of this paper is on systems that maintain complex, multi-level organization across shifting and unpredictable environments. For these, passive stability is insufficient, and adaptive mechanisms become necessary.

Physical structures such as crystals persist because they occupy local minima of free energy; they are close to thermodynamic equilibrium. Their stability reflects the fact that, under given conditions, no nearby configuration is energetically favored. By contrast, living and cognitive systems persist in states that are far from equilibrium. Their organization is not a static minimum but a dynamically maintained pattern that requires continuous energy flow and regulation. The persistence problem addressed in this paper therefore concerns the maintenance of organized structure in far-from-equilibrium systems, where stability cannot be reduced to simple energetic minimization.

7.3 Is This Just a Restatement of Darwinism?

Another concern is that the framework merely redescribes natural selection in more abstract terms. While Darwinian evolution is indeed one of the central mechanisms discussed, the present argument situates it within a broader physical problem: how complex organization persists at all in a changing universe.

Darwinian processes are presented here as one general solution to this problem, operating at the level of populations across generations. The complementary strategy of predictive modeling, operating within individual systems, is not reducible to classical Darwinian evolution, even though it shares a similar logic of variation, evaluation, and retention. By placing both under the umbrella of persistence, the argument highlights their common structure and clarifies why both might be expected to arise in sufficiently complex, dynamic environments.

7.4 What About Robustness and Error-Correction Mechanisms?

Many biological and engineered systems employ mechanisms that maintain stability without apparent reliance on variation and selection at the moment of perturbation. Examples include DNA repair pathways, protein-folding chaperones, homeostatic feedback loops, and redundancy in functional subsystems. These processes resist disruption and correct errors, thereby contributing significantly to persistence.

Such mechanisms, however, do not replace adaptive strategies; they complement them. Robustness mechanisms typically operate within a limited range of environmental conditions, buffering systems against small or expected perturbations. When conditions move outside this tolerance range, robustness alone is insufficient, and systems must adapt through changes in structure or behavior. Moreover, in biological systems, many robustness mechanisms are themselves products of evolutionary selection, and in engineered systems, they arise from design processes that similarly explore and refine alternatives. Robustness can thus be understood as a local persistence strategy constructed by broader adaptive processes.

Rather than constituting a third independent solution to the persistence problem, robustness mechanisms function as components within adaptive systems, extending the range over which existing organization can be maintained while reducing the frequency with which deeper structural change is required.

7.5 Does This Framework Explain the Origin of Complexity?

The argument addresses the persistence of complex systems, not the detailed mechanisms by which the first instances of such systems arise. Questions about the origins of life, cognition, or genetic coding involve additional historical and chemical contingencies that are beyond the scope of this paper.

The present claim is conditional: given the existence of a complex, hierarchically organized system, its continued survival in a changing environment requires adaptive mechanisms of the kinds described. How the earliest such systems came into being remains an important and open scientific question, but it does not undermine the more general constraints on how complexity can be maintained once it exists.

8. Implications: Life, Intelligence, and a Generalized Anthropic Principle

The perspective developed in this paper reframes several familiar questions about life and intelligence. Rather than treating these as isolated biological or psychological phenomena, it places them within a more general physical context: the problem of maintaining complex organization in a changing universe.

From this standpoint, life is not merely a chemical curiosity but a particularly effective solution to the persistence problem. Through replication with variation and selection, living systems maintain and refine structured organization across generations, allowing populations to track environmental regularities over long timescales. Intelligence, in turn, represents a further internalization of this adaptive logic. By constructing and updating internal models, organisms can anticipate change within their own lifetimes, shifting part of the adaptive burden from population dynamics to real-time prediction and decision-making.

This framing suggests a broader, non-observer-centric form of anthropic reasoning.

Wherever complex systems persist over long periods in a changing environment, mechanisms such as replication, learning, or prediction should be expected to play a central role in that persistence. These mechanisms are not arbitrary embellishments but natural responses to environments that vary over time. In contrast, in regions where conditions are either too chaotic or too static to support the accumulation of hierarchical structure, such mechanisms, and the forms of complexity they sustain, may be rare or absent.

The argument does not imply that life or intelligence are inevitable in every possible universe or environment. Rather, it establishes a conditional expectation: wherever persistent, deeply organized systems exist in a changing world, mechanisms that track and respond to environmental structure are likely to be present. Biological evolution and predictive cognition can thus be understood as two major expressions of a more general principle governing the endurance of complexity.

This perspective also helps explain why processes that build internal models of the world, whether in nervous systems, social institutions, or artificial learning systems, play such a central role in the most enduring forms of organization we know. As the number of interacting levels in a system increases, so does the need for mechanisms that compress environmental structure into reusable form and guide action accordingly. The persistence of such systems depends not only on their material components, but on their ability to represent and anticipate the conditions under which those components can continue to function.

In this way, life and intelligence cease to appear as anomalies superimposed on an otherwise indifferent physical world. Instead, they are predicted wherever long-lived complex systems exist, as mechanisms that sustain organized structure under continual change.

The perspective developed here resonates with several existing scientific traditions, including work on dissipative structures in non-equilibrium thermodynamics and predictive approaches in biology and cognition. These connections are meant to be suggestive rather than technical: the present argument operates at a conceptual level, identifying broad structural constraints on persistence without committing to any specific formalism.

9. The Subjective Side: Consciousness and Agency (Tentative Reflections)

The discussion so far has remained at the level of objective description, treating adaptive modeling and selection as functional mechanisms that support persistence in a changing universe. For systems like human beings, however, these processes are accompanied by subjective experience. It is therefore natural to ask how the persistence-oriented framework developed here might relate to familiar philosophical questions about consciousness and agency.

The account offered in this section is intentionally tentative. It does not aim to solve the so-called "hard problem" of consciousness or to provide a reductive explanation of subjective experience. Instead, it proposes a way of situating consciousness and the experience of free will within the same persistence-oriented framework that has been applied to life and intelligence.

One plausible hypothesis is that consciousness corresponds to the ongoing stabilization of a predictive world model under continuously changing input. Organisms that rely on internal models must integrate new sensory information with prior expectations in a way that preserves coherence over time. From the outside, this process can be described in terms of neural dynamics, representational updating, and error correction. From the inside, it may be experienced as the unified, moment-to-moment flow of awareness: a stable sense of a world and a self persisting through change.

On this view, subjective experience is not an additional feature layered on top of adaptive processing, but the first-person aspect of the very processes that enable a system to maintain a coherent model of its environment. The continuity of consciousness mirrors the continuity required of the model itself; both must remain sufficiently stable to guide action while remaining flexible enough to incorporate new information.

A parallel suggestion can be made about agency and the experience of free will. Predictive systems generate and evaluate possible futures, selecting actions based on their anticipated consequences. Functionally, this is a process of model-based decision-making under uncertainty. Subjectively, it may be experienced as deliberation, choice, and the sense that one could act in more than one way. From this perspective, "free will" can be understood as the first-person perspective on the control layer of a predictive system: its capacity to select among internally simulated possibilities in the service of continued viability.

It is important to distinguish this proposal from the claim that all adaptive processes have a subjective aspect. Population-level evolution, for example, implements a form of distributed adaptation across generations, but it lacks the temporally unified, centrally integrated model maintenance characteristic of individual cognitive systems. The processes that underlie subjective experience, if they exist, appear to require a bounded system capable of real-time integration, model stabilization, and action selection. Evolutionary processes provide the conditions for such systems to arise but do not themselves constitute subjects of experience.

These reflections do not resolve the metaphysical question of why subjective experience should accompany certain physical processes. They instead offer a principled way to connect the subjective side of human life with the functional mechanisms that support persistence. If consciousness and agency are to be located within the natural world, the integrative, model-based processes central to adaptive survival provide a natural place to look.

10. Conclusion

This paper began from a simple but easily overlooked observation: in a universe that evolves through time, highly organized, hierarchically structured systems are not the most probable outcomes of physical processes. For any system with many interacting components, the space of possible configurations is vast, and the subset corresponding to coherent, functionally integrated organization is comparatively small. Environmental fluctuations and internal perturbations are therefore far more likely to disrupt such organization than to enhance it. The continued existence of complex systems is thus not something to be taken for granted, but a phenomenon that calls for explanation.

Framed in this way, life, intelligence, and other forms of deep organization can be understood as consequences of mechanisms that support persistence under change. When processes such as replication, learning, or prediction arise, by whatever means, they enable organized systems to persist and to accumulate hierarchical structure over time.

The central claim of this paper is that there are only limited general strategies by which such persistence can be achieved. Two have been identified. The first operates across populations and generations: replication with variation and differential survival, the familiar Darwinian process of evolution. The second operates within individual systems: the construction and continual updating of internal models that allow possible futures to be simulated and actions to be selected accordingly. Although these strategies are often treated as belonging to separate explanatory domains, they share a common adaptive logic of variation, evaluation under environmental constraint, and retention of effective structure.

From this perspective, evolution and intelligence are not isolated anomalies but scale-dependent implementations of a single underlying solution to the persistence problem. Both enable systems to track environmental regularities and to refine their organization in response to change. Systems that lack such mechanisms may exist transiently, especially under stable conditions, but they are unlikely to accumulate the deep, multi-layered complexity characteristic of living and cognitive systems.

By shifting attention from specific biological or psychological definitions to the more general problem of persistence, this framework provides a unifying lens on the endurance of complexity. It suggests that wherever highly organized systems persist in a changing world, we should expect to find processes that build, maintain, and refine models of environmental structure, whether distributed across populations or embodied within individual agents. The existence of such mechanisms is not an incidental feature of complex systems, but a reflection of the fundamental constraints imposed by life in a dynamic universe.