Journal Academica Vol. 15(2), pp. 3–69, August 26 2026 - Astropolitics - Theoretical Computer Science

ISSN 2161-3338 - online edition www.journalacademica.org

White Paper

Non-Human Intelligence: Why Earth?

NHI Capability, Computability, and the Darwinian Logic of Exploration, Observation, and Preservation

Karim Daghbouche *
GridSAT Stiftung, Hannover (Germany)
*Corresponding author: karim@gridsat.io

July 3, 2026

ABSTRACT

What rational motive for NHI observing Earth can be derived without assuming benevolence, moral superiority, or human-like psychology?

This White Paper begins with the strongest strategic premise. A non-human technological lineage capable of developing, deploying, and sustaining technology far beyond present human capability necessarily exhibits, at the operative system level, a minimum instrumental rationality and a demonstrated capacity for continuation.

Taken to the technological frontier, that Darwinian premise first intensifies the threat problem and then changes its structure. Technology can suppress ordinary environmental selection and make common matter increasingly fungible, while no increase in effective computational capability by itself supplies universal foresight.

The substitute for universal prediction is access to realized search: independent biospheres, technological trajectories, and potentially other mature civilizations contain solutions generated outside the observer's own history.

The resulting *Darwinian inversion* is that extreme self-interest can favor exploration, observation, sampling, and preservation. If multiple mature NHI lineages exist, the same logic extends recursively: other civilizations become both competitors and independent continuation witnesses, forming a strategic ecology around valuable convergence nodes such as living technological worlds.

Keywords: Non-Human Intelligence, UAP, Capability and Intent, Darwinian Survival, Biosphere Preservation, Halting Problem, Best of Evolution, Best of Continuation, Constraint Physical Computing, Spectral Gap, ZFC Incompleteness, Post-Biological Intelligence, Strategic Ecology, Multi-NHI Interaction

EXECUTIVE SUMMARY

KEY QUESTION

Why Earth? If a Non-Human Intelligence (NHI) possesses technology far beyond present human capability, what rational motive for observing Earth can be derived without assuming benevolence, moral superiority, or human-like psychology?

The White Paper answers that question by deliberately maximizing, not softening, the Darwinian premise. The modeled actor is a mature technological lineage: if a class of UAP were eventually attributed to NHI technology and the attributed capability were genuine, sustained advanced technology itself would imply a minimum level of means–end competence at the operative system or lineage level. That is *instrumental rationality*. The actor can retain objectives, select means, correct failure, and preserve functionality through time. An entity unable to sustain such relations would be an unconstrained hazard, not the technologically coherent actor modeled here.

KEY JUDGMENTS

- **Bleeding-edge Darwinism.** Assume extreme capability, instrumental rationality, and a continuation interest. This is the strongest threat premise, not a reassurance premise.
- **Disease of civilization.** A civilization that succeeds at suppressing disease, scarcity, and uncontrolled environmental selection also suppresses part of the external adaptive search that shaped it. Survival becomes increasingly dependent on deliberate adaptation and successor design.
- **No universal prediction.** Within the effective-computation class used in this White Paper, substrate substitution does not erase undecidability when an explicit computable, property-preserving reduction exists. Classical, quantum, and physical representations can become radically more capable without yielding universal foresight. ZFC adds an instance-specific formal incompleteness layer.
- **Harvest the realized search.** If the universal Best of Evolution or Best of Continuation cannot be precomputed for every admissible future, an advanced lineage has reason to access solutions already realized by independent physical histories. A biosphere is therefore not merely a mine; it is a continuing generator of evolutionary witnesses.
- **Darwinian inversion.** When ordinary matter becomes comparatively fungible but independent histories remain path-dependent and future-relevant, self-interest can make observation, sampling, preservation, and exploration rational. The continuation motive remains; what counts as the scarce resource changes.

The official UAP problem is already an intent problem. The 2021 ODNI Preliminary Assessment stated that limited high-quality reporting hampered firm conclusions about the nature or intent of UAP, while also noting a limited set of reports describing unusual flight behavior [28]. That record is used here only as empirical context. The White Paper asks what motive becomes rational *if* a technologically advanced NHI class is established empirically.

The critical change in the model is therefore not that an advanced actor becomes kinder. The actor remains self-interested. The change is that technological maturity alters the scarcity structure. Common matter and energy sources may become easier to obtain locally, while a causally independent four-billion-year evolutionary history cannot be replaced simply by moving an equivalent mass of atoms. The future usefulness of every component cannot be exhausted in advance under the universal-prediction boundary. Preserving the source therefore retains option value.

Humanity increases that value. Homo sapiens is simultaneously an outcome of terrestrial evolution, a powerful modifier of the biosphere, and an emerging technological lineage beginning to recognize the systemic value of the living process that produced it. Observation of humanity can therefore have *trajectory value*: the object is not only what our species is, but what an independently evolved technological species becomes.

The same reasoning extends beyond an NHI–Earth dyad. If deep time contains several mature, causally independent technological lineages, then each lineage is itself a realized continuation history. For one NHI, another NHI can therefore be both a competitor and an information-bearing continuation witness. A living technological planet can become a *convergence node*: independent observers first arrive because of the source, and then become part of one another's strategic environment. The predicted regime is neither universal war nor universal federation. It is *strategic ecology*: competition, avoidance, monitoring, exchange, mutualism, symbiosis, and coevolution can coexist.

CENTRAL THESIS

The White Paper strengthens the Darwinian premise and takes self-interest to the technological frontier. When universal foresight remains bounded and independent realized histories retain future option value, the same continuation imperative that maximizes the initial threat concern can make observation, preservation, and exploration rational.

TABLE OF CONTENTS

TERMINOLOGY AND SYMBOLS

Term	**Meaning in this White Paper**
AGI	Artificial General Intelligence: a prospective AI system with general cognitive competence comparable to or exceeding humans across almost all cognitive domains. AGI is used prospectively; no present-achievement premise is required.
Best of Continuation	The generalized problem faced by a self-modifying or successor-producing intelligence: which future continuation, architecture, or successor state best preserves its relevant objectives, capabilities, identity, or survival?
Best of Evolution	The universal evolutionary-prediction objective introduced in the 2024 *Bona Fide* paper: identification in advance of the optimal future evolutionary trajectory.
Biospheric awareness	The epistemic condition in which a technological species recognizes that immediate or presently measurable utility does not exhaust the systemic, historical, or future option value of components of its biosphere. Biospheric awareness is epistemic and may precede complete behavioral implementation.
Capability–Intent Non-Equivalence	The epistemic principle that observed technological capability does not uniquely determine the objective for which that capability is exercised.
Computability Invariant	The theorem-grounded principle that, for sufficiently expressive effective computational architectures, increasing implementation capability does not by itself eliminate undecidable properties inherited through computable reductions.
Continuation Invariant	A model-level principle: the substrate carrying a technological lineage may change from biological to hybrid or post-biological without thereby eliminating the requirement for successful continuation of a lineage that persists.
CPC	Constraint Physical Computing: a theoretical architecture in which semantic constraints are compiled into a physical substrate and resolved through physical state realization, dissipation, and readout.
Empirical witness	A physically realized configuration showing that a particular solution exists and functions under stated conditions.
Empirical witness reservoir	A collection of physically realized solutions produced through independent natural or artificial history. A biosphere is treated as an evolutionary witness reservoir.
Halting Problem	The undecidable problem of determining, for every possible program and input, whether the program eventually halts.

Term	**Meaning in this White Paper**
Historical-Independence Invariant	The causal property that reproduction of matter or particular functions does not retrospectively remove the independent history by which an external physical or evolutionary witness set was generated.
Instrumental rationality	The minimum means–end competence implied in this White Paper by sustained advanced technological capability: the ability of a technological actor or lineage to retain objectives, select effective actions, correct failures, preserve complex functionality, and coordinate behavior through time. It is a competence concept, distinct from morality, benevolence, or human-like psychology.
Intent	The objective an actor or system is pursuing. Observed capability does not uniquely determine intent.
NHI	Non-Human Intelligence. For purposes of this White Paper, the term follows the operational definition used in the earlier *Bona Fide* discussion: a non-human intelligence capable of manifesting through documented UAP, i.e., a technologically advanced non-human actor. The definition identifies the hypothetical class under examination; attribution of particular UAP remains a separate empirical question.
Organic stagnation	The idealized reduction of externally driven adaptive exploration when technology strongly buffers a biological lineage from environmental selection. Mutation, drift, and other evolutionary mechanisms continue; the term isolates the reduction of externally driven adaptive exploration.
Representation Undecidability	The absence, within a sufficiently expressive effective physical-representation family, of a universal algorithmic or formal certification procedure for every relevant global property of every exactly specified representation.
Resource-Value Inversion	The model-level shift under advanced fabrication in which ordinary raw matter becomes increasingly fungible while independently generated evolutionary or physical histories retain non-fungible informational value.
Convergence node	A location or system whose independent scientific, biological, technological, or informational value can attract more than one mature technological lineage, thereby creating a secondary arena of interaction among the observers themselves.
Strategic ecology	The multi-lineage regime in which independent technological civilizations become part of one another's adaptive environment. Possible relations include competition, avoidance, observation, exchange, mutualism, symbiosis, and coevolution; the term predicts neither universal war nor universal cooperation.

Term	Meaning in this White Paper
Spectral gap	The energy separation between a ground-state sector and the lowest excitation above it. For Hamiltonian families, the asymptotic existence of such a gap can be undecidable.
Technological Survivorship Threshold	The minimum continuity competence implied by a technological lineage that persists long enough to develop, deploy, and sustain deep-space or interstellar capability. It is not a claim of moral superiority.
Trajectory value	The potential informational and strategic value of observing the development through time of an independently evolving system, as distinct from the value of a single biological or technological snapshot.
UAP	Unidentified Anomalous Phenomena. The term describes an unresolved observational category, not an established origin.
ZFC	Zermelo–Fraenkel set theory with the Axiom of Choice, used here as a representative powerful formal axiomatic system.

HOW TO READ THIS WHITE PAPER

The argument contains four logical levels.

Keeping them distinct is essential.

Status	Role
Established theorem	The Halting Problem, Rice-type semantic undecidability, spectral-gap undecidability, and the corresponding formal incompleteness results.
Conditional mathematical transfer	The transfer of spectral-gap undecidability into a CPC family, and the subsequent transfer into biological realizability, each require an explicit effective embedding.
Model assumption	The Darwinian survival model, technological buffering, continued adaptive requirements, the preservation of the home environment, advanced fabrication, and the existence of independent biospheres.
Interpretive application	The UAP capability–intent argument, including the possibility that observation or monitoring serves a survival and preservation motive.

The strongest conclusions are stated relative to their premises. This four-level structure preserves the distinction among established theorem, conditional mathematical transfer, model premise, and interpretive application, so that each conclusion can be stated at its maximum justified strength.

1. WHY EARTH? ACTOR CLASS, INTENT, AND THE DISTINCTIVE QUESTION

The modeled actor class is deliberately narrow. Microbial intelligence, animal-like intelligence, primitive technological species, and inert artifacts lie outside the strategic derivation. The relevant class is a *mature technological NHI*: a lineage or operative system whose demonstrated capability, if empirically attributed to NHI, is sufficient to support sustained advanced technology, long-horizon action, and deep-space or comparable operation.

For such a class, capability itself carries an analytic implication. A system that can develop, deploy, navigate, maintain, repair, reproduce, or coordinate advanced technology must embody sufficient means–end competence for that capability to persist. This White Paper calls that minimum condition *instrumental rationality*. It is not a moral claim. It is the narrower proposition that coherent technology requires coherent relations among objectives, actions, feedback, and persistence.

ACTOR-CLASS PREMISE

If a UAP class is eventually established as NHI technology and the attributed technology demonstrates sustained advanced capability, then the operative NHI system or lineage is modeled as possessing at least the instrumental rationality required to sustain that capability. An entity incapable of maintaining persistent means–end relations belongs to a different analytic class: an unconstrained hazard, not an actor whose strategy can be inferred from durable objectives.

This distinction allows the capability–intent problem to be framed without anthropomorphism. Capability constrains what the actor can do and supplies evidence of operational coherence. The actor's specific intent remains the open variable.

The 2021 ODNI Preliminary Assessment is useful as an official baseline. It reported that limited high-quality data prevented firm conclusions about UAP nature or intent, while also recording a small number of cases with reported unusual flight characteristics and emphasizing the need for improved collection and analysis [28]. The present White Paper addresses the next-stage theoretical question: if a genuinely advanced NHI class is established, what motives remain rational under the capability?

1.1 Closest Antecedents and the Distinctive Pipeline

Several literatures approach individual parts of this problem. Two antecedents are especially close to different halves of the pipeline. On the formal side, Day proves an inherent limit on negation-complete prediction for open-ended evolutionary dynamics and explicitly connects that limit to the Halting Problem and Goedel incompleteness [31]. Hernandez-Orozco, Hernandez-Quiroz, and Zenil further show, within a computable-dynamical-systems framework, that strong open-ended evolution entails undecidability and irreducibility [32]. On the strategic side, Lane's 2025 dual-goals hypothesis treats self-preservation and resource acquisition as convergent instrumental goals of advanced rational agents and argues that complex biospheres and civilizations can be locally information-rich and strategically worth observing [34]. Ball's original Zoo Hypothesis frames Earth as a deliberately protected wilderness or observational reserve; Fogg's Interdict Hypothesis and later zoo-hypothesis work develop the noninterference problem further [35, 37, 44]. Freitas's SETA analysis independently argued that a rational long-duration extraterrestrial probe seeking life and intelligence would treat Earth as the principal Solar-System target for continuous surveillance [36]. Cabrol emphasizes that the coevolution of life and environment makes each planetary experiment historically distinctive [38]. Assembly theory likewise gives formal prominence to formation history and causal contingency in complex objects [39].

Other work supplies adjacent pieces. Dick develops the long-horizon possibility of a post-biological universe [40]. Haqq-Misra and Baum question whether rapid expansion remains sustainable for mature extraterrestrial civilizations [41]. Frank, Grinspoon, and Walker model intelligence as a possible planetary-scale transition in biosphere-technosphere coupling [42]. Baum, Haqq-Misra, and Domagal-Goldman show why contact analysis must span multiple beneficial, neutral, and harmful scenarios rather than project one preferred motive [43]. A recent preprint by Giakas independently advances a structural, non-moral route by which high technological maturity can disfavor overt interstellar aggression through resource, thermodynamic, relativistic, and strategic constraints [46]. Jebari and Asker show that the presence of multiple extraterrestrial civilizations can change the logic of a Hobbesian first-strike trap [45]; Forgan's simulations likewise favor multiple causally connected cliques over a single universal galactic club across broad parameter choices [44].

A striking but methodologically different antecedent is Thomson's 2012 *On the Frontier of Knowledge: A Discussion of Alien Civilizations* [47]. That paper explicitly links Earth's unusual biodiversity to alien interest and discusses multiple alien species interacting around Earth. Its empirical method is coordinate remote viewing, and the paper itself states that the accuracy of that method under the circumstances studied cannot be determined. It is therefore relevant here only as a conceptual antecedent, not as evidence for NHI presence, multiplicity, or motive.

Among the literature reviewed here, these ideas appear separately. The distinctive contribution of the present White Paper is the single derivation that connects them:

DISTINCTIVE PIPELINE

Darwinian continuation → disease-of-civilization suppression of external selection → dependence on deliberate adaptation → substrate-invariant computability/certification boundary → no universal foresight → option value of independent realized histories → harvest, observation, preservation, and exploration → recursive multi-NHI strategic ecology.

The argument is intentionally asymmetric. The computability results are theorem-grounded. The evolutionary, resource, and strategic consequences are explicit model-level inferences. Their separation is retained throughout.

2. BLEEDING-EDGE DARWINISM: START WITH THE STRONGEST THREAT PREMISE

The analysis begins where threat assessment is most uncomfortable. Assume an actor with extreme technological superiority. Assume that its technology is operated coherently enough to imply instrumental rationality. Assume that the lineage places extraordinary value on continuation. Assume that when continuation requires action, it is capable of acting.

No benevolence is introduced to soften the premise. The modeled objective is simply

$$\boxed{\text{continue to exist.}} \tag{1}$$

For an organic lineage, this objective is expressed through Darwinian survival and reproduction. For a hybrid or post-biological lineage, the biological mechanism may change, but a lineage that persists must still produce future states compatible with its continuation. The White Paper therefore treats Darwinian survival as the organic instance of a more general continuation problem.

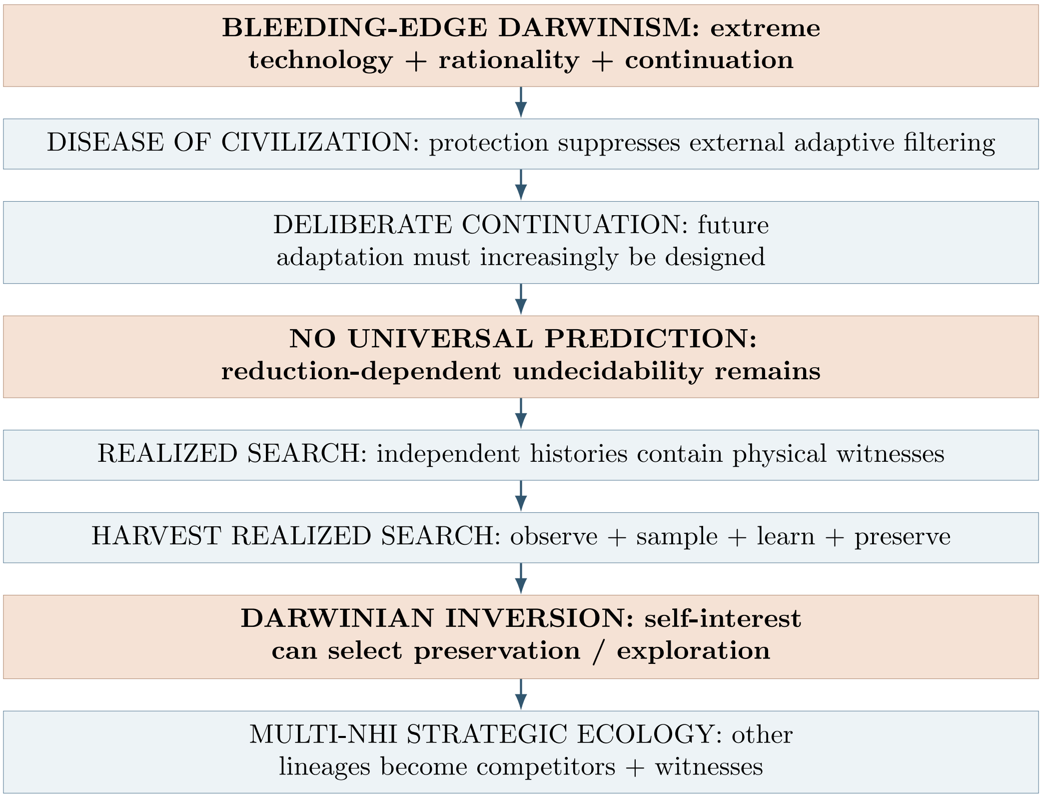

Figure 1: The inference pipeline of the White Paper. The mathematical claim is concentrated in the computability/certification boundary; the survival, resource, and multi-NHI steps are explicit model-level consequences.

BLEEDING-EDGE PREMISE

The strongest initial threat case is retained:

$$\text{extreme capability} + \text{instrumental rationality} + \text{continuation interest.}$$

The paper's preservation result must be derived from this premise, not from a replacement assumption of enlightenment or benevolence.

A technological lineage capable of sustained deep-space or interstellar operation also reveals a selection effect. Some continuity mechanism—biological, social, institutional, machine-based, or hybrid—has survived long enough to build, deploy, and sustain that capability. This is the *Technological Survivorship Threshold*. It says nothing by itself about morality or peacefulness. It says that the lineage has achieved practical continuity competence.

This premise initially increases fear. A superior actor with coherent means–end behavior and a strong continuation objective is exactly the class for which adverse intent must be considered. The rest of the White Paper asks a harder question: what does that same continuation objective select once technological maturity changes both the actor's environment and the resource structure around it?

3. THE DISEASE OF CIVILIZATION: ORGANIC STAGNATION

The first consequence of bleeding-edge Darwinism is biological rather than computational.

The earlier *Bona Fide* argument considered an advanced civilization that has strongly reduced ordinary external selection pressures [1].

The term *organic stagnation* is used here as an idealized limiting concept.

Mutation, recombination, genetic drift, and other evolutionary processes continue.

The idealized quantity driven toward zero is specifically

$$\boxed{\text{adaptive exploration generated by uncontrolled external challenge.}} \tag{2}$$

Let the controlled home environment of the NHI be denoted

$$E_{\text{NHI}}. \tag{3}$$

Let

$$S(E_{\text{NHI}}) \tag{4}$$

denote the effective intensity of uncontrolled environmental selection.

The limiting thought experiment considers

$$S(E_{\text{NHI}}) \to 0. \tag{5}$$

The civilization has become exceptionally successful at preventing environmental conditions from deciding which members survive.

This is the civilizational achievement.

It is also the evolutionary inversion.

The system now becomes more dependent on intentional generation of future adaptations.

Hence,

$$\boxed{\begin{array}{c} \text{survival technology} \\ \Downarrow \\ \text{removal of immediate pressures} \\ \Downarrow \\ \text{loss of uncontrolled adaptive testing} \\ \Downarrow \\ \text{dependence on deliberate adaptation.} \end{array}} \tag{6}$$

The more successful the protection becomes, the more adaptation moves from uncontrolled environmental selection toward conscious design.

The disease-of-civilization analogy therefore becomes the first link in the NHI pipeline:

> The technology that solves the original survival problem creates a new dependence that the earlier environment did not contain.

FIRST INVERSION

The success of survival technology can create a continuation dependency. The more completely a biological civilization suppresses uncontrolled existential filtering at home, the more future adaptation must be generated deliberately rather than supplied by the environment that once selected it.

4. THE COMPUTABILITY BOUNDARY: NO UNIVERSAL PREDICTION

The disease of civilization creates the demand: replace lost uncontrolled adaptation with deliberate foresight. At this point the White Paper grants the NHI as much computational power as the argument can responsibly grant. The question is not whether prediction becomes better. It does. The question is whether any effective architecture can supply a universal predictor for every admissible future case.

The answer is the mathematical hinge of the paper.

4.1 The Human Technology Track: From AI to the Quantum Frontier

A discussion of hypothetical NHI technology benefits from locating present human technology first.

The comparison is necessarily limited.

A civilization capable of sustained deep-space or interstellar activity, if such a civilization exists, could employ architectures substantially beyond present human engineering.

Current technology nevertheless provides a useful reference ladder.

The broad trend is from increasingly powerful symbolic processing toward increasingly direct control of physical state spaces.

4.1.1 General-Purpose AI

Modern general-purpose AI is implemented primarily on classical digital infrastructure.

Large neural systems are trained and executed on specialized accelerators integrated into conventional computing systems.

Their capabilities have advanced rapidly.

The 2026 *International AI Safety Report* states that leading systems now match or exceed human-expert performance on standardized evaluations across a growing range of professional and scientific subjects, while retaining a markedly "jagged" performance profile and important limitations in reliability and real-world operation [19].

Current systems can already assist software development, mathematics, scientific reasoning, multimodal analysis, and increasingly extended agentic workflows.

The present paper does not require a claim that current systems constitute AGI.

Artificial General Intelligence is used here as a prospective category:

$$\text{AI with general cognitive competence across almost all relevant cognitive tasks.} \tag{7}$$

The significance lies in the technological feedback loop:

$$\text{computation} \to \text{AI-assisted science} \to \text{improved technology} \to \text{improved computation.} \tag{8}$$

For an advanced civilization, this feedback could become a major mechanism of technological acceleration.

4.1.2 Quantum Computing

Quantum computation changes the substrate available to computation.

Classical bits are replaced by controlled quantum degrees of freedom.

Algorithms can exploit superposition, interference, and entanglement.

The field has passed important experimental milestones.

Below-threshold surface-code quantum error correction has been demonstrated, with logical error suppression improving as code distance increases [20].

This is a major step toward scalable fault-tolerant quantum computation.

Large-scale, general-purpose, fault-tolerant quantum computing nevertheless remains under development.

The near-term transition is strongly hybrid: quantum processors operate alongside classical computation rather than independently from it [21].

It is equally important to distinguish quantum computational advantage from hypercomputation.

Standard adiabatic quantum computation, for example, is polynomially equivalent to the conventional quantum circuit model [16].

Quantum mechanics can therefore change the resource complexity of important computations without by itself removing the classical computability boundary.

Artificial intelligence and quantum engineering are also beginning to converge.

AI methods are being applied to quantum-device design, characterization, control, compilation, error mitigation, and related stages of the quantum-computing stack [22].

The corresponding feedback loop can therefore widen:

$$\boxed{\begin{array}{c}\text{AI} \rightarrow \text{improved physical design and control} \\ \Downarrow \\ \text{more capable quantum computation} \\ \Downarrow \\ \text{new scientific and engineering capability} \\ \Downarrow \\ \text{more capable AI and automation.}\end{array}} \tag{9}$$

4.1.3 Constraint Physical Computing as a Frontier Architecture

Constraint Physical Computing (CPC) occupies a further conceptual position along this trajectory.

CPC is presently a theoretical physical-computation framework, rather than a mature fault-tolerant commercial platform.

Its significance here follows from the question it asks.

Instead of treating matter primarily as hardware that executes a symbolic instruction sequence, CPC considers whether a semantic constraint can be compiled into a physical substrate whose states, symmetries, dissipation, and readout jointly implement the intended response [8, 9].

Related work has formalized configuration identity, canonical quotient-state realization, response quotients, exact semantic carriers, and conditional limits on exact semantic-carrier separation [3–7].

Abstractly,

$$K \xrightarrow{\mathcal{C}} H_K \xrightarrow{\text{physical evolution}} \rho_K \xrightarrow{R_K} y_K, \tag{10}$$

where K is a constraint specification,

H_K is a compiled physical interaction or Hamiltonian,

ρ_K is a realized physical state,

and R_K is a semantic readout.

CPC therefore shifts the operational emphasis from executing a conventional sequence of instructions to realizing a constrained physical state.

For the present argument, this makes CPC a useful frontier between conventional computation and hypothetical highly advanced physical computation.

A conceptual technology ladder can be represented as

$$\boxed{\begin{array}{c}\text{classical high-performance computing}\\ \Downarrow \\ \text{general-purpose AI and agents}\\ \Downarrow \\ \text{fault-tolerant quantum computing}\\ \Downarrow \\ \text{AI–quantum co-design}\\ \Downarrow \\ \text{constraint-based physical computation}\\ \Downarrow \\ \text{unknown NHI technology, if it exists.}\end{array}} \tag{11}$$

Equation (11) is a capability ladder, not a proposed historical sequence for NHI development.

Its purpose is to grant the hypothetical intelligence the strongest plausible technological position before asking which boundaries remain:

> Grant the hypothetical intelligence as much technological advancement as the discussion reasonably permits, then ask which boundaries remain.

4.2 The Halting Problem: The Boundary Beyond Engineering

For a non-specialist reader, the distinction between a difficult problem and an undecidable problem is essential.

Computational complexity asks how many resources are needed to solve a problem.

Faster hardware, parallel computation, improved algorithms, additional energy, and quantum speedups can alter practical feasibility.

Computability asks a prior question:

> Does a universal algorithm for the problem exist?

The classical Halting Problem asks whether an algorithm can determine, for every possible program M and input x, whether $M(x)$ eventually halts or continues indefinitely.

Turing established that no such universal algorithm exists [11].

This is qualitatively different from insufficient processor speed.

A civilization can build a faster processor.

It can build more processors.

It can improve its algorithms.

It can exploit quantum mechanics.

It can automate the design of subsequent generations of computation.

Within effective computation, these achievements do not convert an undecidable universal predicate into a decidable one.

The distinction can be written

$$\boxed{\begin{array}{ccc} \text{engineering limitation} & \to & \text{potentially removable} \\ \text{complexity limitation} & \to & \text{potentially shifted} \\ \text{undecidability} & \to & \text{structural boundary.} \end{array}} \tag{12}$$

This is why the Halting Problem becomes increasingly important in a thought experiment about extreme technological advancement.

As contingent engineering limitations are removed, foundational limitations become more visible.

4.3 CPC and the Operational Reframing of Halting

CPC changes the operational semantics of computation without removing the underlying computability boundary.

Consider again

$$K \xrightarrow{\mathcal{C}} H_K \xrightarrow{\text{physical evolution}} \rho_K \xrightarrow{R_K} y_K. \tag{13}$$

The chain separates three questions.

4.3.1 Representation

Can the semantic constraint K be represented in a physical substrate?

This is a question about the mapping

$$K \longmapsto H_K. \tag{14}$$

4.3.2 Realization

Does the physical substrate reach, maintain, or expose the intended state or response sector?

4.3.3 Certification

Can the required global property of the physical realization be determined for every member of the architecture?

A conventional Turing computation is naturally represented by a state sequence

$$s_0 \to s_1 \to s_2 \to \cdots \tag{15}$$

with a distinguished halting condition.

A CPC realization can associate successful computation with a physical state, ground sector, response class, or other semantic carrier.

Accordingly,

$$\boxed{\text{CPC success predicate} \;\neq\; \text{Turing halting predicate}} \tag{16}$$

in the general architecture considered here.

Halting therefore becomes operationally secondary.

The physical system does not need to use arrival at an abstract halt instruction as its success criterion.

Physical evolution still occurs in time.

Preparation, dissipation, relaxation, stabilization, and readout remain physical processes.

The important transition is semantic.

The computational question has moved from

$$\text{“Will the symbolic process halt?”} \tag{17}$$

toward

$$\text{“Does the physical family possess the required global property?”} \tag{18}$$

Quantum many-body physics then supplies a remarkable result.

Undecidability appears inside the physical classification problem itself.

4.4 Executable CPC Validation and the Representation–Realization Boundary

The CPC architecture now has an executable reference validation framework. *CPC Validation* v0.6.0 defines a Canonical Constraint Intermediate Representation (CCIR), backend compilation contracts, qualified RC, Digital, and FPGA execution backends, and independent semantic validation [29]. The same canonical constraint program can be compiled into structurally different execution artifacts and compared only after backend execution and fixed decoding.

That architecture reinforces the distinction used throughout this White Paper: representation, compilation, preparation, execution, observation, decoding, and certification are separate layers. Agreement among heterogeneous backends is evidence for the executed corpus; it is not a theorem of universal correctness. Version 0.6.0 additionally specifies RFC-0010 driven-dissipative open-system execution and referenced readout, while the software record explicitly distinguishes that specification and simulated reference witness from a physical open-system realization [29].

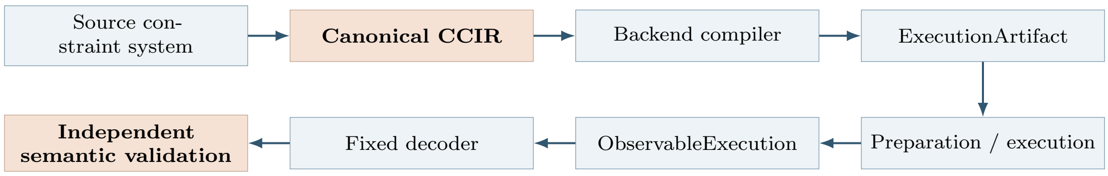

Finite validation can establish reproducible backend agreement and semantic matches for executed cases; it does not collapse the distinction between realization and universal certification.

Figure 2: Simplified CPC validation architecture, adapted from the CPC Validation v0.6.0 reference framework [29].

4.5 Spectral-Gap Undecidability

4.5.1 The Spectral Gap

Consider a sequence of finite quantum systems indexed by system size L, with local Hamiltonians

$$H^{(L)} = \sum_{\langle i,j \rangle} h^{(i,j)}. \tag{19}$$

Let

$$E_0^{(L)} \tag{20}$$

denote the ground-state energy,

and let

$$E_1^{(L)} \tag{21}$$

denote the lowest energy strictly above the ground-state sector.

The finite-size spectral gap is

$$\Delta_L = E_1^{(L)} - E_0^{(L)}. \tag{22}$$

For intuitive purposes, a uniformly gapped family possesses a persistent positive separation between the ground sector and excitations above it.

There exist $\gamma > 0$ and L_0 such that

$$\Delta_L \geq \gamma \qquad \forall L \geq L_0. \tag{23}$$

A gapless family permits arbitrarily low-energy excitations as the system grows.

The formal spectral-gap undecidability theorems use more precise promised alternatives.

For the present non-specialist discussion, the essential point is simpler.

The Hamiltonian specifies local physical interaction rules.

The spectral gap describes a global asymptotic property of the family.

One might naturally expect that exact knowledge of the local rules would always make the global classification computable in principle.

For sufficiently expressive Hamiltonian families, that expectation fails.

4.5.2 Undecidability Enters Quantum Matter

Cubitt, Pérez-García, and Wolf constructed computably specified translationally invariant nearest-neighbor quantum spin systems for which the spectral-gap problem is undecidable [13].

There is no universal terminating algorithm that takes every Hamiltonian in the constructed family and always determines correctly whether the thermodynamic system is gapped or gapless.

Thus,

$$\boxed{\nexists\, A \text{ deciding } \operatorname{Gap}(H) \text{ for every } H \text{ in the family.}} \tag{24}$$

The result was subsequently extended to one-dimensional translationally invariant nearest-neighbor quantum spin chains [14].

The logical mechanism is central to the present paper.

The Hamiltonian construction encodes computation such that the asymptotic spectral behavior depends upon the behavior of an associated Turing computation.

In reduction-theoretic notation,

$$\text{HALT} \ \leq_m \ \text{SPECTRAL GAP}\,. \tag{25}$$

Equation (25) provides the bridge.

The Halting Problem has not disappeared when computation is represented through quantum matter.

Its role has changed.

It no longer needs to describe the operational endpoint of the physical computation.

Instead, it proves that a physically meaningful global classification problem can itself be undecidable.

Thus,

> Halting becomes operationally secondary and foundationally decisive.

This is the sense in which the Halting Problem remains all-important even at the quantum frontier.

4.5.3 Specification and Certification

The spectral-gap result also exposes a general separation.

A Hamiltonian may be specified exactly.

Its local interaction matrices may be finite.

Its rules may be explicit.

Its representation may be complete.

Yet a universal procedure for classifying the relevant asymptotic property can still fail to exist.

Therefore,

$$\boxed{\text{exact local specification} \not\Rightarrow \text{universal global certification.}} \tag{26}$$

Many individual Hamiltonians remain tractable.

Finite systems can possess measurable energy differences.

The theorem concerns the universal classification problem over a sufficiently expressive family.

That distinction is precisely what is required for the CPC argument.

4.6 The ZFC Blockade: Representation Undecidability

4.6.1 From Undecidability to Formal Independence

Algorithmic undecidability asks whether a universal algorithm exists.

Formal independence asks whether a proposition follows from the axioms of a chosen formal system.

The concepts are distinct.

Spectral-gap theory connects them.

Spectral-gap incompleteness results show that for suitable consistent formal systems capable of reasoning about elementary arithmetic, Hamiltonians can be constructed whose relevant spectral-gap classification is independent of the chosen theory [15].

Let F denote such a formal theory.

An associated Hamiltonian

$$H_{\mathsf{F}} \tag{27}$$

can possess a spectral proposition that cannot be settled within F.

Taking

$$\mathsf{F} = \mathsf{ZFC}, \tag{28}$$

under the required consistency assumptions, gives the corresponding ZFC form of the result.

This gives the proposed *ZFC blockade* a direct physical interpretation.

A sufficiently expressive formal description can specify a physical Hamiltonian while the relevant global classification remains outside what the chosen formal theory can settle.

4.6.2 Gap-Certified CPC Realization

Let an effective CPC compiler be written as

$$\mathcal{C} : K \mapsto \{H_K^{(L)}\}_{L\in\mathbb{N}}. \tag{29}$$

Define

$$\mathrm{GAdm}(K) \tag{30}$$

to mean that the compiled family has the spectral behavior required by a specified gap-certified CPC realization.

This is intentionally a restricted physical predicate.

Full CPC correctness may additionally require

semantic carrier identity,

state preparation,

readout fidelity,

noise tolerance,

response separation,

and other properties.

The spectral criterion provides the rigorous undecidability bridge.

Proposition 1 (Conditional CPC Certification Undecidability). Let

$$\{G_n\}_{n\in\mathbb{N}} \tag{31}$$

be a computably specified Hamiltonian family for which the spectral-gap problem is undecidable.

Suppose the CPC architecture is sufficiently expressive to admit a computable embedding

$$r : n \longmapsto K_n \tag{32}$$

such that

$$\mathrm{Gap}(G_n) \iff \mathrm{GAdm}(K_n). \tag{33}$$

Then no total algorithm decides

$$\mathrm{GAdm}(K) \tag{34}$$

for every CPC instance K.

Proof. Assume a total algorithm

$$A_{\mathrm{CPC}} \tag{35}$$

decides $\mathrm{GAdm}(K)$ for every K. For arbitrary n, compute

$$K_n = r(n) \tag{36}$$

and evaluate

$$A_{\mathrm{CPC}}(K_n). \tag{37}$$

By Eq. (33), this decides

$$\mathrm{Gap}(G_n). \tag{38}$$

The procedure would therefore decide the undecidable spectral-gap family.

This yields a contradiction. □

4.6.3 Representation Undecidability

The preceding proposition motivates the following terminology.

> **Representation Undecidability** is the absence, within a sufficiently expressive effective physical-representation family, of a universal algorithmic certification procedure for the required global realization property of every exactly specified representation.

The concept distinguishes representation from universal foresight.

A compiler may provide

$$K \longmapsto H_K \tag{39}$$

exactly.

The existence of that representation does not imply the existence of a universal transformation

$$H_K \longmapsto \text{correct global classification.} \tag{40}$$

Accordingly,

$$\boxed{\text{Representable} \;\not\Rightarrow\; \text{universally certifiable.}} \tag{41}$$

This sharpens the principle that representation precedes solvability [10].

Representation makes a problem available to formal or physical treatment.

Its existence alone does not make every global property decidable.

Corollary 1 (Conditional Formal-Incompleteness Transfer). Let F be a consistent, effectively axiomatized formal theory sufficiently strong for elementary arithmetic.

If a CPC embedding preserves a spectral-gap proposition that is independent of F, then the corresponding gap-certified CPC proposition inherits that formal independence.

For

$$\mathsf{F} = \mathsf{ZFC}, \tag{42}$$

this yields the ZFC form of Representation Undecidability, subject to the embedding and consistency conditions.

The computational boundary has changed its physical form:

$$\boxed{\begin{array}{ccc} \text{sequential computation} & \longrightarrow & \text{halting classification} \\ \text{physical computation} & \longrightarrow & \text{global physical classification.} \end{array}} \tag{43}$$

The foundational boundary remains.

4.7 What "Across Computation" Means

The universality claim in this White Paper is precise. It is not the claim that every conceivable physical process or every logically imaginable machine is Turing-equivalent. The relevant class consists of effective computational architectures for which the target semantic or physical property is linked to a known undecidable source problem by a computable, property-preserving reduction. Within that class, changing substrate does not erase the inherited undecidability. Standard quantum computation changes complexity without supplying a hypercomputational escape from the Turing boundary [16]. Super-Turing proposals also exist in the evolutionary-computation literature; for example, Eberbach explicitly defines nonalgorithmic Evolutionary Turing Machines and studies halting in that enlarged model class [33]. Such explicitly nonalgorithmic models lie outside the effective-computation premise used here unless a physically realizable hypercomputational mechanism is independently established. CPC changes representation and execution semantics; its undecidability conclusion follows only when the explicit embedding conditions stated in Proposition 1 hold.

Formal incompleteness is stronger but narrower. ZFC does not fail to settle every physical question. The Cubitt incompleteness construction yields specific physical propositions that become independent of a sufficiently expressive consistent formal theory under the relevant construction and assumptions [15].

A 2026 result by Xu et al. sharpens the spectral-gap scope further. They prove that the *bulk* spectral gap is semi-decidable and explicitly distinguish that problem from earlier undecidability results for alternative gap notions based on sequences of finite systems with prescribed boundary conditions [30]. The present White Paper therefore uses the Cubitt–Perez-Garcia–Wolf and Bausch et al. family constructions in their stated sense and does not generalize spectral-gap undecidability to every notion of gap, every Hamiltonian, or every finite physical system.

NO UNIVERSAL PREDICTION

Within the effective-computation class modeled here, implementation power can increase without producing a universal decision procedure. The operational substrate may change; the inherited undecidability remains whenever the required computable reduction exists. Prediction can become extraordinary. It does not become universally exhaustive.

5. FROM PREDICTION TO HARVEST: BEST OF EVOLUTION AND BEST OF CONTINUATION

The computational result creates the strategic problem. An advanced lineage can model, simulate, design, and test at extraordinary scale, but the universal objective remains unavailable. The response is therefore not to abandon prediction. It is to stop demanding that prediction replace every realized history.

The White Paper calls the universal biological objective *Best of Evolution* and its substrate-independent generalization *Best of Continuation*. Day's open-ended-evolution theorem provides a direct formal antecedent for the claim that universal evolutionary prediction can encounter computability-theoretic limits even when local evolutionary dynamics are known [31]. The present construction uses that broader computability lineage as a foundation and then asks what the boundary implies for technological continuation strategy. Neither term implies that natural evolution itself computes a global optimum. The strategic substitute is empirical: retain access to physically realized candidate solutions that independent histories have already generated and tested.

THE HARVEST PRINCIPLE

A mature intelligence need not predict every future solution if it can retain access to independent processes that continue to generate and preserve realized solutions. The rational strategy becomes a portfolio: predict what can be predicted, design what can be designed, and harvest what independent history has already realized.

5.1 The Best of Evolution

The 2024 paper applied this distinction to evolution [1].

Let

$$P(\mathrm{NHI}) \tag{44}$$

denote the possible evolutionary trajectories available to a hypothetical NHI.

The proposed function

$$H(P(\mathrm{NHI})) \tag{45}$$

was intended to identify, in universal form, which trajectory would yield the optimal evolutionary outcome:

$$\boxed{\textit{Best of Evolution.}} \tag{46}$$

This is stronger than ordinary biological prediction.

Many biological questions are finite, probabilistic, empirical, statistical, or approximately computable.

The Best-of-Evolution construction instead concerns a universal predictor over an effectively encoded class of adaptive trajectories sufficiently expressive to represent arbitrary computable self-referential dynamics.

Its prediction may itself influence the future trajectory being predicted.

Under this expressiveness condition, the trajectory class supports the diagonal construction required for the Halting reduction.

The central conclusion is therefore correspondingly precise:

> For any such sufficiently expressive class of self-referential adaptive trajectories, a universal computational procedure cannot in general precompute the Best of Evolution.

This result remains significant even if the hypothetical civilization possesses computational capability far beyond humanity.

It may possess AGI.

It may possess post-AGI machine intelligence.

It may possess fault-tolerant quantum computers.

It may possess automated molecular design.

It may possess physical-computation architectures that humanity has not yet developed.

The universal Best-of-Evolution problem remains of another logical type.

Thus,

$$\boxed{\text{more computational capability} \not\Rightarrow \text{universal evolutionary foresight.}} \tag{47}$$

This is the first component of the survival dilemma.

5.2 Organic Realizability: The Conditional Biological Transfer

5.2.1 A Biological Design Predicate

Let g denote a finite description of a proposed

genetic,

molecular,

developmental,

or other biological design.

Let e denote an environment,

and let τ denote an observation interval.

A finite empirical viability predicate may be written

$$V(g, e, \tau) \in \{0, 1\}. \tag{48}$$

When a physically instantiated organism with configuration g survives and reproduces in environment e over interval τ,

$$V(g, e, \tau) = 1 \tag{49}$$

records an empirical result for those conditions.

A mature synthetic-design system may seek a stronger property.

Let

$$\mathrm{OrgAdm}(g) \tag{50}$$

denote a specified robust biological-realizability criterion over a prescribed class of

environments, perturbations, timescales, or other effectively indexed evaluation conditions.

The description g may itself be finite.

For an undecidability transfer, however, the global predicate $\mathrm{OrgAdm}(g)$ must be permitted to quantify over an effectively specified unbounded family of evaluation instances.

The relevant unbounded parameter could in principle be

time, generational depth, environmental sequence, system scale, or another effectively indexed dimension of the target family.

Thus,

$$\boxed{\text{finite description} \not\Rightarrow \text{finite-horizon property.}} \tag{51}$$

The mathematical question is whether a sufficiently expressive biological design family can inherit the same type of certification obstruction identified for Hamiltonian systems.

5.2.2 Organic Realizability Transfer Condition

The spectral-gap theorem supplies the source problem.

A biological undecidability conclusion requires an effective reduction to the biological target problem.

Suppose there exists a computable map

$$b : n \longmapsto g_n \tag{52}$$

from an undecidable spectral-gap family

$$\{G_n\} \tag{53}$$

into a physically admissible biological-design family.

Suppose further that $\mathrm{OrgAdm}(g_n)$ is a global property of an effectively specified unbounded evaluation family associated with g_n, and that the embedding preserves the relevant classification:

$$\mathrm{Gap}(G_n) \iff \mathrm{OrgAdm}(g_n). \tag{54}$$

Under Eq. (54), a universal decision procedure for

$$\mathrm{OrgAdm}(g) \tag{55}$$

would decide the embedded spectral-gap problem.

The undecidability obstruction therefore transfers.

Equation (54) identifies the precise mathematical bridge required for a biological undecidability theorem.

The biological conclusion is therefore conditional on realization of that effective embedding.

This improves the evolutionary argument.

The relevant theorem class concerns sufficiently expressive global biological-design predicates, rather than individual genomes as such.

In a sufficiently expressive physical biological-design architecture, greater representational power does not automatically produce universal pre-certification of every possible novel design.

For an advanced civilization, that distinction has strategic significance.

It may be able to construct a candidate adaptation without possessing a universal method that guarantees the future behavior of every possible candidate.

5.3 Post-Biological Intelligence: From Best of Evolution to Best of Continuation

A natural objection now arises.

Why assume that an intelligence capable of deep-space exploration remains organic?

The objection is substantial.

A sufficiently advanced civilization may be

biological,

machine-augmented,

hybrid,

uploaded,

synthetic,

or entirely post-biological.

The organic argument therefore needs to be separated from the more general computational argument.

5.3.1 Organic or Hybrid NHI

For an organic or substantially biological NHI, the previous argument applies directly:

$$\text{survival} \to \text{adaptive novelty} \to \text{Best of Evolution} \to \text{independent biological witnesses.} \tag{56}$$

The direct value of Earth's biosphere follows most strongly in this case.

5.3.2 Fully Post-Biological AI

A fully post-biological intelligence need not require organic evolution for its own persistence.

Its corresponding problem is more general.

Suppose an AI can modify

its software,

its memory architecture,

its objectives,

its learning system,

its hardware,

its embodiment,

or the architecture of successor systems.

Its continuation can be represented schematically as

$$A_0 \to A_1 \to A_2 \to A_3 \to \cdots . \tag{57}$$

The relevant optimization problem is no longer merely

$$\textit{Best of Evolution.} \tag{58}$$

It becomes

$$\boxed{\textit{Best of Continuation.}} \tag{59}$$

The system may seek a successor A' that optimally preserves

capability,

identity,

objectives,

adaptability,

security,

or continuity.

Abstractly,

$$A^* = \arg\max_{A'} U_{\text{cont}}(A'), \tag{60}$$

where

$$U_{\text{cont}} \tag{61}$$

is the relevant continuation criterion.

5.3.3 The Computability Boundary for Post-Biological Intelligence

If the candidate successor systems are sufficiently expressive computational objects, universal prediction of their future behavior again encounters computability limits.

The Halting Problem already establishes one such boundary [11].

Rice's theorem generalizes the basic phenomenon for extensional semantic properties: every nontrivial property of the partial function computed by an arbitrary program is undecidable in general [12].

Rice's theorem applies specifically to nontrivial extensional semantic properties of the partial function computed by a program.

Syntactic and intensional properties, including code length, directly inspectable architecture, specified hardware parameters, and other locally decidable features, can remain decidable.

For self-modifying AI, the obstruction therefore concerns the sufficiently expressive universal semantic objective.

Many particular changes can be analyzed.

Many successors can be tested.

Many properties can be formally verified.

The stronger universal objective is the problem.

For successor classes expressive enough to represent arbitrary partial computable behaviors, and for continuation criteria that depend upon a nontrivial extensional property of that behavior,

$$\boxed{\begin{aligned}&\text{constructability of a successor}\\ &\not\Rightarrow\ \text{universal certification of its future behavior.}\end{aligned}} \tag{62}$$

Post-biological intelligence therefore changes the object of continuation.

The foundational prediction boundary remains.

5.4 Penrose, AI Self-Awareness, and the Preservation Principle

Roger Penrose famously used Gödelian incompleteness as part of a broader argument against the view that human mathematical understanding can be completely captured by algorithmic computation [23, 24].

The argument generated extensive debate and significant criticism [25].

The present model is fully compatible with computational machine consciousness and self-awareness.

Accordingly, a computational AI may be granted full self-awareness.

The important distinction is

$$\boxed{\text{self-awareness} \neq \text{complete self-prediction.}} \tag{63}$$

A self-aware AI may understand

that it exists,

that it has goals,

that it has a particular architecture,

that it can modify itself,

that future successors can differ from it,

and that its own computational framework contains universal prediction limits.

Self-awareness strengthens the argument when it includes awareness of the system's own prediction boundaries.

This produces an important inversion in the meaning of intelligence itself.

Greater intelligence can sharpen the distinction between enormous predictive power and universal foresight.

A sufficiently advanced intelligence may understand that boundary with greater precision.

It may know which questions are engineering questions.

It may know which questions are complexity questions.

It may know which questions are undecidable in universal form.

It may therefore know that destruction can eliminate future options whose value cannot be exhaustively certified before the destruction occurs.

Self-awareness then becomes a source of epistemic restraint.

Not moral restraint.

Not emotional restraint.

Epistemic restraint.

The intelligence preserves because it knows that knowledge has boundaries.

5.4.1 Awareness of an Epistemic Boundary

Consider an intelligence that understands the Halting Problem, Rice's theorem, and corresponding physical undecidability results.

Such an intelligence can distinguish

$$\text{“I can construct this system”} \tag{64}$$

from

$$\text{“I can universally certify every consequence of constructing it.”} \tag{65}$$

This distinction is central.

The intelligence can explicitly know that the epistemic boundary exists.

5.4.2 External Solutions Acquire Option Value

Suppose an external solution reservoir $\mathcal{S}$ contains physically realized configurations generated independently of the AI's own history.

Suppose further that

1. destruction of $\mathcal{S}$ is effectively irreversible;
2. some elements of $\mathcal{S}$ may have positive future utility;
3. the AI cannot exhaustively precompute which elements will have future utility;
4. the cost of preservation does not dominate every possible future benefit.

Then preservation retains option value.

The intelligence can rationally conclude

$$\boxed{\begin{array}{l}\text{preserve the independent solution space}\\ \text{when future relevance cannot be exhausted in advance.}\end{array}} \tag{66}$$

This is not a theorem that all rational AIs must preserve all external systems.

The conclusion depends upon the stated decision conditions.

Its significance is conceptual.

Self-awareness can convert undecidability from a passive limitation into a behavioral driver.

The chain becomes

$$\boxed{\begin{array}{c}\text{self-awareness}\\+\\\text{awareness of universal prediction limits}\\\Downarrow\\\text{recognition of unknown future needs}\\\Downarrow\\\text{option value of independent solution spaces}\\\Downarrow\\\text{preservation and exploration.}\end{array}} \tag{67}$$

This gives the Penrose discussion a different conclusion.

The operational conclusion is therefore independent of any prohibition on machine self-awareness:

> The more completely an intelligence understands its own computational limits, the stronger the rational case can become for preserving what it cannot recreate or universally predict.

5.4.3 Best of Evolution as a Special Case of Best of Continuation

The post-biological case reveals the deeper invariant.

Darwinian evolution is the organic mechanism.

Continuation is the general problem.

A machine intelligence may continue without DNA, cellular reproduction, or biological selection.

A machine lineage that persists must nevertheless produce successor states compatible with its future conditions.

The biological *Best of Evolution* therefore becomes a special case of the substrate-independent *Best of Continuation*.

The transition is

$$\boxed{\begin{array}{c}\text{Darwinian continuation}\\\Downarrow\\\text{technological continuation}\\\Downarrow\\\text{post-biological continuation.}\end{array}} \tag{68}$$

The carrier changes.

The continuation requirement remains.

5.5 Technology-Invariant Constraints

The preceding sections permit the argument to be reorganized around a simpler question:

$$\boxed{\text{What survives technological substitution?}} \tag{69}$$

The term *invariant* is used here at four explicitly distinguished logical levels:

theorem-grounded,

model-relative,

causal and historical,

and epistemic.

The first is theorem-grounded.

The second is model-relative.

The third is causal and historical.

The fourth is epistemic.

Their common function is that each identifies a property that is not removed merely by increasing the technological capability of the hypothetical actor.

5.5.1 Computability Invariant

The strongest invariant is the computability boundary.

The implementation may change from

classical computation

to

quantum computation

to

physical constraint realization.

The computational resources may increase by orders of magnitude.

The civilization may automate its own scientific development.

None of those facts by itself establishes a hypercomputer.

For sufficiently expressive effective computational architectures, undecidable properties inherited through computable reductions remain undecidable.

Thus:

$$\boxed{\text{greater computational capability} \;\not\Rightarrow\; \text{greater computability class.}} \tag{70}$$

This is the *Computability Invariant.*

Its strategic significance is considerable.

An intelligence can know vastly more than humanity without becoming omniscient.

The difference between those propositions is the difference between technological superiority and universal foresight.

5.5.2 Continuation Invariant

The second invariant belongs to the model.

The physical carrier of intelligence may change.

An organic lineage may become augmented.

An augmented lineage may become hybrid.

A hybrid lineage may produce machine successors.

The original biological species may eventually disappear.

Yet a lineage that continues still confronts the continuation problem.

For biological organisms, Darwinian evolution provides the familiar form.

For post-biological intelligence, the corresponding problem is the *Best of Continuation*.

Thus:

$$\boxed{\text{substrate replacement} \not\Rightarrow \text{elimination of continuation risk.}} \tag{71}$$

This is the *Continuation Invariant* within the model.

Darwinian survival is not discarded when the discussion moves to AI.

It is generalized.

The biological question

> Which adaptation allows the lineage to survive?

becomes

> Which successor allows the relevant continuation to persist?

The object changes.

The structural problem remains.

5.5.3 Residual Uncertainty

The first two invariants combine into a consequential result.

Capability can grow.

Prediction can improve.

Simulation can improve.

Scientific understanding can improve.

Yet sufficiently expressive universal prediction problems retain undecidable instances.

Therefore:

$$\boxed{\text{increasing prediction} \not\Rightarrow \text{exhaustive prediction.}} \tag{72}$$

Residual uncertainty is not a statement that the NHI knows little.

It may know almost unimaginably more than humanity.

The statement is that "almost unimaginably more" and "everything" are different categories.

For deep-time continuation, that difference matters.

5.5.4 Historical-Independence Invariant

Suppose an intelligence can reproduce a molecule.

Suppose it can reproduce a genome.

Suppose it can reproduce an organism.

Suppose it can construct an artificial ecosystem.

Those capabilities can make particular biological functions technologically reproducible.

They do not alter the causal fact that another biosphere was generated through a history independent of the observer.

Thus:

$$\boxed{\text{reproduction of an outcome} \neq \text{reproduction of its independent history.}} \tag{73}$$

This is the *Historical-Independence Invariant*.

Individual functions may remain technologically substitutable.

The invariant concerns historical provenance, and its strategic significance is epistemic.

Before future requirements are known, an observer cannot presume that its endogenous search has already generated every survival-relevant outcome contained in every independent physical history.

5.5.5 Capability–Intent Non-Equivalence

The final invariant belongs to the observer's inference problem.

Capability constrains what an actor can do.

The specific intent remains underdetermined by capability alone.

Thus:

$$\boxed{C \not\Rightarrow I,} \tag{74}$$

where C denotes observed capability and I a unique intent.

This remains true whether the proposed intent is hostile or benign.

The preservation motive follows only after capability is combined with:

continuation,

computability limits,

independent history,

future uncertainty,

and irreversibility.

The resulting question is no longer merely

> How advanced is the actor?

It becomes

> Given what remains invariant even at extreme capability, what strategies remain rational?

6. HARVESTING THE REALIZED SEARCH: BIOSPHERES AS CONTINUING SOURCES

The word *harvest* is used here in an informational, non-destructive sense. It includes observation, sampling, comparison, selective transfer, and learning from realized solutions. Indeed, once a biosphere is understood as a continuing generator rather than a static deposit, destruction can reduce the very flow of novelty that makes the source strategically valuable.

Cabrol's description of planetary coevolution supports the historical premise: each biosphere emerges through a distinctive interaction of life, environment, bottlenecks, and stochastic events [38]. Assembly theory independently emphasizes that history and causal contingency influence which complex objects exist [39]. These results do not prove the NHI motive model; they strengthen its characterization of a biosphere as a path-dependent physical history rather than interchangeable bulk matter.

THE SOURCE MUST REMAIN A SOURCE

A biosphere is not merely biological inventory. It is an active search process. If continuing independent novelty has positive future option value, then the highest-value harvest can be observation and selective sampling while preserving the generator itself.

6.1 External Biospheres as Living Evolutionary Witnesses

Universal prediction attempts to calculate what can happen.

Nature supplies something different:

what has happened and remained physically viable.

That distinction becomes increasingly important once universal foresight is unavailable.

A living biosphere is not merely a catalogue of genomes.

It is an ongoing, physically realized historical process.

Its surviving organisms are local witnesses.

Its ecological relationships are compound witnesses.

Its evolutionary history is a continuing record of solutions, failures, dependencies, and contingencies that were physically tested rather than merely simulated.

Suppose an organism with biological configuration g has survived and reproduced under environment e for interval τ.

Then

$$W = (g, e, \tau) \tag{75}$$

is a physically realized witness to

$$V(g, e, \tau) = 1. \tag{76}$$

The witness does not predict every biological future.

It establishes that one physical configuration has functioned under one realized set of conditions.

A biosphere consists of an enormous collection of such witnesses.

Across evolutionary time,

mutation,

recombination,

selection,

drift,

development,

competition,

symbiosis,

pathogens,

climate,

ecological interaction,

and geological change

continually explore physical biological configuration space.

The surviving biosphere contains the residue of that physical search.

Let

$$\mathcal{W}_P = \{(g, e, \tau) : V(g, e, \tau) = 1\} \tag{77}$$

denote the empirically accessible witness set of planet P.

Two independently evolved planets need not possess identical witness sets.

Thus

$$\mathcal{W}_{P_1} \neq \mathcal{W}_{P_2} \tag{78}$$

is generically possible because the planets possess different

histories,

environments,

contingencies,

and lineages.

An advanced intelligence exploring another biosphere therefore acquires more than specimens.

It gains access to a physical search history that its own lineage did not execute.

The central value is

$$\boxed{\text{independent evolutionary history.}} \tag{79}$$

This sharpens the earlier "planetary oracle" intuition.

A biosphere is not treated as an oracle Turing machine.

The more precise concept is

an empirical evolutionary witness reservoir.

The value of $\mathcal{W}_P$ therefore lies in more than the number of specimens that can be collected.

It lies in provenance.

The search occurred.

It occurred physically.

It occurred independently of the observer.

It occurred under a history the observer did not choose.

This is why the earlier "planetary oracle" metaphor is better understood as an *empirical evolutionary witness reservoir*.

The biosphere does not answer arbitrary undecidable questions.

It does something more modest and physically meaningful.

It presents realized solutions that exist without first requiring the observer to predict them.

That distinction is precisely what an intelligence aware of universal prediction limits has reason to value.

6.2 Exploration as a Necessary Survival Strategy

The stronger conclusion requires the assumptions of the model to be explicit.

6.2.1 Assumption A1: Protected Civilization

The NHI maintains a highly controlled home environment that suppresses much of the uncontrolled selective pressure that shaped its earlier evolution.

6.2.2 Assumption A2: Continued Adaptive Requirement

Deep-time survival requires continuing access to adaptive novelty because future perturbations cannot all be known from the civilization's present biological state.

Internal simulation, synthetic evolution, AI-assisted search, and deliberate biological design can generate substantial novelty.

The model grants such internal methods substantial, and potentially extreme, generative power.

The required condition is narrower: their internally selected search spaces, objective functions, training histories, and experimental environments cannot be presumed in advance to exhaust every survival-relevant configuration that independent physical histories may realize.

Independent evolutionary histories therefore possess non-redundant option value under uncertainty.

6.2.3 Assumption A3: Predictive Boundary

A universal Best-of-Evolution predictor is unavailable under the Halting construction.

Sufficiently expressive physical design families can additionally retain undecidable certification properties.

6.2.4 Assumption A4: Preservation of the Protected State

Reintroducing unrestricted species-level Darwinian pressure inside the NHI population would sacrifice the protected civilizational condition that technological development created.

The civilization seeks adaptation without subjecting its own population to uncontrolled existential selection.

6.2.5 Assumption A5: Independent Biospheres

Other accessible biospheres exist whose evolutionary histories are substantially independent of the NHI lineage.

6.2.6 The Three Broad Strategies

Under A1–A5, three broad strategies can be distinguished.

Universal Prediction The civilization could attempt to calculate every future adaptation that deep-time survival will require.

This is the universal Best-of-Evolution strategy.

Assumption A3 supplies its foundational boundary.

Prediction remains powerful for particular cases.

It cannot serve as a universal replacement for open-ended physical history.

Internal Restoration of Selection The civilization could restore strong, uncontrolled Darwinian selection within its own population.

This would recreate an endogenous evolutionary search.

Under A4, the strategy conflicts with the defining protected condition of the civilization.

It deliberately reintroduces the existential mechanism that technology had removed.

External Evolutionary Acquisition The third strategy preserves the protected home environment while accessing evolutionary outcomes generated under independent selective histories.

The civilization explores external biospheres. It thereby gains access to

$$\mathcal{W}_{P_1}, \mathcal{W}_{P_2}, \ldots, \mathcal{W}_{P_n}. \tag{80}$$

Synthetic experimentation can continue.

AI can continue.

Quantum simulation can continue.

CPC can continue.

Artificial evolution can continue.

The distinctive contribution of external exploration is independence.

Here, "independent" means historically and causally independent of the NHI's own evolutionary lineage, design choices, training history, and experimental search process.

The term denotes historical and causal provenance, rather than statistical independence in the probabilistic sense or functional uniqueness of every resulting solution.

It accesses trajectories generated outside

the NHI lineage,

its controlled environment,

its chosen model class,

its internal search history,

and its technological representation framework.

Proposition 2 (Exploration Necessity within the Organic Model). Under assumptions A1–A5, continued access to independently generated evolutionary witness sets, while preserving the protected home environment, requires external biosphere exploration.

Argument. Universal prediction cannot provide a complete replacement under A3.

Unrestricted endogenous species-level reselection violates A4.

Internal experiments remain valuable.

This includes synthetic evolution, large-scale simulation, AI-generated biological search, quantum simulation, and physically realized artificial ecosystems.

Such methods may reproduce particular external solutions and may generate solutions superior to particular naturally evolved ones.

Their limitation for the present proposition is not computational scale but provenance: their search conditions, objective functions, initializations, and representational choices remain generated within the NHI's own technological and epistemic environment.

Under A2, no such endogenous search process is presumed in advance to exhaust the survival-relevant outcomes of all causally independent physical histories.

Independent naturally evolved witness sets therefore add a distinct source of realized adaptive information.

Under A5, access to those external witness sets requires external exploration. □ The proposition establishes necessity for the specified objective.

It does not establish that exploration guarantees survival.

No particular planet is guaranteed to contain the solution to a particular future problem.

The conclusion is strategic:

$$\boxed{\begin{array}{c}\text{if independent adaptive diversity is required}\\ \text{while uncontrolled endogenous selection remains excluded}\\ \Downarrow\\ \text{external biological exploration is not optional within the model.}\end{array}} \tag{81}$$

Exploration becomes part of species-preservation infrastructure.

The proposition is compatible with extreme internal technological capability.

The NHI may possess extraordinary simulation.

It may possess artificial evolution.

It may operate synthetic ecosystems.

It may use AI to generate biological designs.

It may use quantum simulation and CPC.

The model grants all of those possibilities.

External exploration adds something those systems do not acquire merely by becoming larger:

historical independence.

The distinction concerns provenance:

> endogenous search and independent search are different information sources.

For an intelligence whose future requirements cannot be exhaustively known in advance, diversity of provenance has option value.

Exploration therefore becomes part of continuation infrastructure not because internal technology has failed, but because internal technology cannot make an external history cease to be external.

6.3 Resource-Value Inversion: When Darwinian Scarcity Changes Its Object

Conventional threat narratives about advanced extraterrestrial intelligence often begin with resource acquisition.

The intuition is historically understandable.

Human societies have competed for

land,

energy,

minerals,

water,

food,

and strategic materials.

Darwinian language appears to reinforce the intuition.

A self-interested actor competes for what it needs.

The White Paper retains that principle and carries it to extreme technological maturity:

what does such an actor actually need?

Advanced fabrication is a separate premise from CPC.

The following argument therefore grants that the advanced NHI possesses highly developed material design, fabrication, recycling, and access to extraterrestrial resources.

Under that premise, the scarcity structure changes.

Matter remains necessary.

Energy remains necessary.

But ordinary terrestrial matter loses some of its *unique* strategic value.

Iron is not unique to Earth.

Carbon is not unique to Earth.

Silicon is not unique to Earth.

Water is not unique to Earth.

A technological actor capable of deep-space engineering has access to material environments far beyond the terrestrial surface.

A living evolutionary history is different.

If material structure can be engineered with high precision, the strategic value moves increasingly from

$$\text{where matter happens to be} \tag{82}$$

toward

$$\text{what information and organization the matter contains.} \tag{83}$$

The Resource-Value Inversion is a claim about comparative scarcity.

Matter and energy remain necessary; their relative strategic uniqueness changes.

At sufficient technological maturity:

$$\boxed{\begin{array}{c}\text{the strategic uniqueness of ordinary matter can decline}\\ \text{while the uniqueness of independent realized history remains.}\end{array}} \tag{84}$$

This changes the Darwinian threat calculation.

At one technological level, self-interest can favor resource competition.

At another, the same self-interest can favor preservation of resources whose defining property is that they cannot simply be mined, manufactured, or reconstructed after destruction.

The object of scarcity has changed.

The Darwinian objective has not.

6.3.1 Path-Dependent Information Is Different

A living biosphere contains ordinary matter.

Its historical organization is extraordinary.

Earth's current biological state is the result of a contingent sequence of

mutations,

extinctions,

ecological transitions,

climatic events,

competition,

cooperation,

and environmental changes

extending across geological time.

One can write schematically

$$\mathcal{H}_P = \{h_0, h_1, \ldots, h_t\} \tag{85}$$

for the evolutionary history of planet P.

The present witness set

$$\mathcal{W}_P \tag{86}$$

is partly a function of that path:

$$\mathcal{W}_P = F(\mathcal{H}_P, E_P). \tag{87}$$

Reproducing the constituent atoms does not reproduce the historical path.

Thus,

$$\boxed{\text{material identity} \neq \text{historical identity.}} \tag{88}$$

This establishes historical non-fungibility.

Historical non-fungibility and functional substitutability are distinct.

A sufficiently advanced synthetic system may reproduce a particular biological function found in an external biosphere, and it may generate a functionally superior solution by another route.

The narrower claim is epistemic.

Before the future survival problem is known, the observer cannot presume that its internally generated search has already reproduced every function that may later prove useful within an independent evolutionary witness set.

The value of the external history therefore lies not in a claim that each of its products is impossible to synthesize, but in retaining a non-redundant reservoir of realized solutions whose future utility is not exhaustively known in advance.

6.3.2 The Resource-Value Inversion

The result can be expressed as a model-level inversion:

$$V(P) = V_{\text{matter}}(P) + V_{\text{history}}(P). \tag{89}$$

As advanced fabrication reduces the scarcity of ordinary material resources,

$$V_{\text{matter}}(P) \tag{90}$$

can lose relative strategic importance.

The independent evolutionary history

$$V_{\text{history}}(P) \tag{91}$$

retains value because it is

path-dependent,

historically contingent,

and not obtainable merely by transporting equivalent bulk matter.

Hence,

$$\boxed{\begin{array}{c}\text{advanced fabrication}\\ \Downarrow \\ \text{lower uniqueness of ordinary raw matter}\\ + \\ \text{independent evolutionary history}\\ \Downarrow \\ \text{higher relative value of biological information.}\end{array}} \tag{92}$$

Matter and energy remain physically necessary.

The inversion concerns strategic uniqueness, not physical necessity.

The narrower claim is that ordinary terrestrial resources become a weaker *unique* motive for interstellar observation once advanced fabrication and widely available extraterrestrial raw material are granted.

The biosphere remains distinctive.

7. THE DARWINIAN INVERSION: WHEN SELF-INTEREST SELECTS PRESERVATION

The argument can now return to its deliberately severe starting premise. The NHI remains instrumentally rational. It remains self-interested. Its continuation objective remains intact. What changes is the resource landscape in which that objective operates.

At lower technological capability, scarce local matter, territory, energy, and reproduction can make competition adaptive. At high technological capability, advanced fabrication and extraterrestrial resource access can reduce the unique value of ordinary terrestrial matter. At the same time, no universal predictor supplies certainty about every future continuation need, and causally independent histories remain non-redundant sources of physically realized solutions.

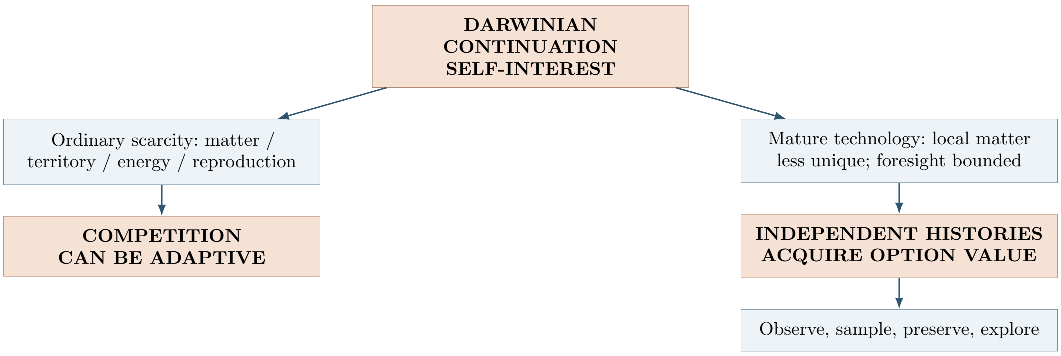

Figure 3: The Darwinian inversion. The continuation objective remains; technological maturity changes which resources are strategically scarce.

The result is not pacifism. A preservation motive can coexist with deterrence, competition, coercion, or conflict when another risk dominates. The model makes a narrower claim:

CENTRAL INVERSION

Under the stated conditions, the strengthened Darwinian premise itself generates a preservation motive. A self-interested intelligence has reason not to destroy a continuing independent source whose future usefulness cannot be universally exhausted before destruction.

This is the central reversal of the White Paper. The fear and the restraint arise from the same continuation premise.

8. FROM ONE NHI TO MANY: STRATEGIC ECOLOGY

The single-NHI model scales recursively. Suppose deep time contains more than one mature technological lineage,

$$\mathcal{N}^\star = \{N_1, N_2, \ldots, N_k\}, \tag{93}$$

where each N_i independently satisfies the actor-class premises used in this White Paper. The derivation applies to each lineage separately. No single observer is privileged.

The Historical-Independence Invariant then turns inward on the observer class. For N_i, another causally independent mature lineage N_j is itself the outcome of a different continuation history. Its existence demonstrates that another lineage survived a distinct sequence of biological, technological, institutional, computational, and existential problems long enough to reach the mature threshold. It is therefore a realized *continuation witness*.

$$\boxed{N_j = \text{competitor} + \text{independent continuation witness}} \tag{94}$$

Those terms can point in opposite strategic directions. If N_j presents overwhelming existential risk, threat reduction may dominate. If its independent information, technology, continued experimentation, or possible cooperation carries higher option value, observation, preservation, exchange, or coexistence can dominate. The model therefore predicts neither universal warfare nor a universal galactic federation. It predicts an ecology of conditional strategies.

A schematic utility for lineage N_i with respect to lineage N_j can be written

$$U_i(N_j) = V_{ij}^{\text{witness}} + V_{ij}^{\text{information}} + V_{ij}^{\text{cooperation}} - R_{ij}^{\text{threat}} - C_{ij}^{\text{interaction}}. \tag{95}$$

The expression is not a universal law. It simply makes explicit that another civilization can carry positive information and option value at the same time that it carries strategic risk. This is the multi-NHI form of the Darwinian inversion.

8.1 The Convergence-Node Effect

Now place an independently evolved biosphere P between several lineages. Initially each lineage can value P for the same reason developed earlier: P contains a search history it did not generate. If several observers converge, the observers become part of one another's environment. The value of the location can then become recursive:

$$V_i(P) = V_i^{\text{biosphere}}(P) + V_i^{\text{trajectory}}(P) + V_i^{\text{other lineages}}(P) + V_i^{\text{interaction}}(P). \tag{96}$$

This is the *convergence-node effect.* A location first attracts actors because of the source; once multiple actors arrive, the actors themselves become additional sources of information, risk, and opportunity.

The human analogy is a cafe, party, or scientific meeting: strangers arrive because of a common attraction, then discover that the other participants are part of the value of remaining. The evolutionary analogue is a watering hole: organisms converge on one resource and subsequently create competition, predation, signaling, avoidance, commensalism, mutualism, symbiosis, and new selection pressures.

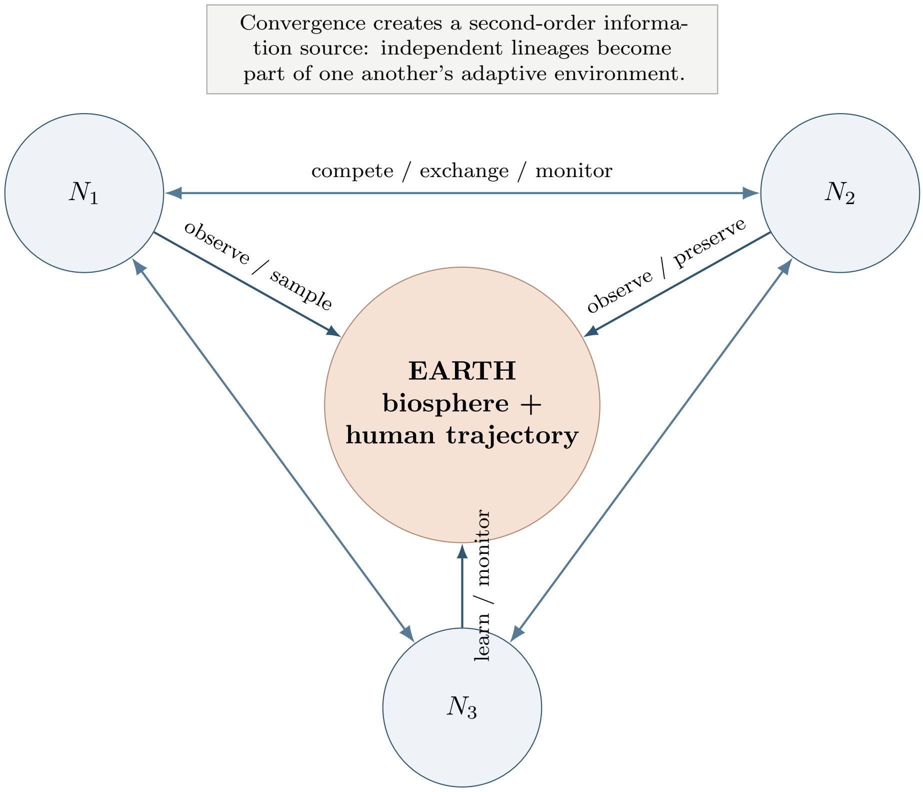

Figure 4: Earth as a possible convergence node in the multi-NHI extension. The diagram is a model consequence, not an empirical claim that multiple NHI are present.

8.2 Interaction Generates New Search

Contact can create novelty that neither lineage contained before contact. Let $\mathcal{W}_i^{\text{pre}}$ and $\mathcal{W}_j^{\text{pre}}$ denote the pre-contact witness sets of two lineages, and let $\mathcal{W}_{ij}^{\text{post}}$ denote configurations realized after interaction. Define the interaction-generated increment

$$\Delta\mathcal{W}_{ij} = \mathcal{W}_{ij}^{\text{post}} \setminus \left(\mathcal{W}_i^{\text{pre}} \cup \mathcal{W}_j^{\text{pre}}\right). \tag{97}$$

When $\Delta\mathcal{W}_{ij} \neq \varnothing$, interaction itself has become a source of realized novelty. This is the technological analogue of coevolution: independent lineages create new pressures, responses, combinations, institutions, defenses, scientific questions, and technological solutions for one another.

Forgan's Galactic-Cliques analysis supports the plausibility of a heterogeneous multi-civilization regime rather than one universal policy [44]. Jebari and Asker show that a multitude of civilizations can change strategic incentives relative to a two-actor dark-forest model [45]. These works do not supply the present information-value derivation, but they make the multi-actor extension structurally relevant.

A useful deep-time conclusion follows:

STRATEGIC ECOLOGY

At sufficient technological maturity, other intelligences can become part of the environment to which intelligence itself must adapt. The same Historical-Independence Invariant that gives external biospheres option value can give independent technological civilizations option value. Competition remains Darwinian; so do cooperation, symbiosis, and coevolution when they improve continuation.

9. HUMANITY INSIDE THE SOURCE: BIOSPHERIC AWARENESS, UAP, AND STRATEGIC SALIENCE

Earth's value in this model is not confined to non-human biodiversity. Humanity belongs to the same independent evolutionary history. We are also the component of that history that has acquired the greatest present technological leverage over the biosphere and is beginning to alter evolution deliberately. That dual position makes the human trajectory both valuable and potentially dangerous to the source itself.

Frank, Grinspoon, and Walker's concept of planetary intelligence provides a useful terrestrial analogue: technological intelligence can be treated as a planetary-scale process whose maturity depends on integration with the planetary systems that sustain it [42]. The present White Paper uses a narrower term, *biospheric awareness*, for the epistemic stage at which a technological species recognizes that immediate utility does not exhaust systemic value.

9.1 Biospheric Competence: Humanity as a Local Analogue

The survivorship threshold also suggests a second transition.

A technological species first acquires the capability to modify its environment.

Later, if it survives, it may acquire the understanding required to recognize the systemic consequences of that capability.

Humanity presently demonstrates that these transitions need not occur simultaneously.

Our species can alter planetary systems faster than institutions, infrastructure, and collective behavior can adapt.

The lag matters.

Recognition precedes implementation, often by decades of technological, economic, and institutional transition.

A civilization can understand that fossil-fuel combustion has planetary consequences while remaining deeply dependent on inherited energy infrastructure.

It can recognize biodiversity loss while continuing to generate it.

It can understand ecological interdependence without possessing the political, economic, or technological coordination to respond immediately.

This is not evidence that awareness is irrelevant.

It identifies a transition period between

capability,

understanding,

and

systemic competence.

The terrestrial trajectory can be represented as

$$\boxed{\begin{array}{c}\text{technological capability}\\ \Downarrow \\ \text{planetary causal power}\\ \Downarrow \\ \text{recognition of systemic dependence}\\ \Downarrow \\ \text{attempted biospheric management.}\end{array}} \tag{98}$$

This White Paper calls the epistemic stage *biospheric awareness*.

Biospheric awareness is the recognition that immediate utility does not exhaust systemic value.

An organism can therefore retain strategic ecological value even when its immediate utility to the technological species is absent, unknown, or negative.

A bee may sting.

A predator may compete.

A pathogen may kill.

An organism may possess no presently known economic value.

A technologically informed ecological perspective nevertheless recognizes that removing poorly understood components from a complex system can have consequences that are difficult or impossible to reverse.

Thus:

$$\boxed{\text{present inconvenience} \;\not\Rightarrow\; \text{absence of future or systemic value.}} \tag{99}$$

Humanity has entered this transition.

That fact is relevant to the NHI model because it provides a terrestrial example of the proposed direction:

increasing technological capability

can eventually increase

the recognition of preservation value.

The early structure of the preservation argument is therefore already visible in human technological development.

9.2 Humanity as Witness, Hazard, and Emerging Biospheric Manager

Earth's strategic value within the model includes humanity.

Homo sapiens is part of the same independent evolutionary history.

This point matters because a simple "protect the biosphere from humanity" narrative would contradict the model's own premises.

Homo sapiens is an intrinsic product of terrestrial evolution and one of its realized outcomes.

At the same time, human technology has acquired planetary causal leverage.

Humanity is therefore

$$\boxed{\text{biospheric product} + \text{biospheric modifier.}} \tag{100}$$

This dual status creates an unusual observational object.

Humanity can damage the evolutionary witness reservoir.

Humanity is also part of that reservoir.

A hypothetical observer valuing independent terrestrial evolution therefore treats the human species as an internal component of the system being valued.

There is a further reason for interest.

Humanity is undergoing a technological transition that changes the character of terrestrial evolution itself.

Biological evolution has produced a species capable of

understanding evolution,

altering selection pressure,

editing genomes,

constructing synthetic biology,

externalizing cognitive functions,

building artificial intelligence,

and beginning to explore space.

The biosphere has therefore generated an actor capable of partially redesigning the future trajectory of the biosphere.

This is a historically distinct event within the terrestrial witness set.

The relevant object is not merely the present human genome.

It is the trajectory:

$$\boxed{\begin{array}{c}\text{evolution}\\ \Downarrow\\ \text{intelligence}\\ \Downarrow\\ \text{technology}\\ \Downarrow\\ \text{planetary agency}\\ \Downarrow\\ \text{deliberate modification of evolution and intelligence.}\end{array}} \tag{101}$$

Humanity may therefore be particularly salient to an observer interested in the transition from natural adaptation to deliberate continuation.

The interest need not be anthropocentric.

It follows from the unusual state of the trajectory.

The terrestrial biosphere has produced an intelligence that is beginning to encounter the same class of continuation questions this White Paper assigns to advanced NHI.

Humanity is developing AI.

Humanity is developing quantum computation.

Humanity is manipulating biological systems.

Humanity is increasingly capable of controlling its physical environment.

Humanity is simultaneously discovering that the complexity of its evolutionary substrate makes careless intervention dangerous.

We are therefore not merely an object inside the biosphere.

We are becoming an observer, designer, and manager of parts of it.

This creates *trajectory value*.

A snapshot answers:

> What is this species now?

A trajectory answers:

> What does an independently evolved technological species become?

The latter question requires time.

It requires observation.

And it creates a reason for minimal perturbation.

If an external observer intervenes heavily, the observed trajectory is no longer the same independent trajectory.

Thus:

$$\boxed{\text{observation value} \rightarrow \text{value of limiting unnecessary perturbation.}} \tag{102}$$

This is not a quantum observer effect.

It is a simple causal point.

Intervention changes the system's history.

The more the independent history is itself valuable, the more intervention has an informational cost.

Preservation therefore need not mean freezing the biosphere.

Evolutionary preservation can instead mean

> maintaining sufficient integrity and autonomy for the independent process to continue.

Human technological development can remain part of that process.

Indeed, whether humanity can reconcile planetary technological power with preservation of the system that produced it may itself be one of the most informative features of the trajectory.

Humanity thus occupies three roles:

$$\boxed{\begin{array}{c}\text{evolutionary witness}\\+\\\text{potential endogenous hazard}\\+\\\text{emerging biospheric manager.}\end{array}} \tag{103}$$

None of these roles requires moral judgment.

They describe the present causal position of a technological species inside its own biosphere.

9.3 UAP: Capability versus Intent

The motive model now returns directly to the UAP question.

Prior work has examined the computational complexity of reverse engineering unknown UAP technology [2].

The present White Paper addresses the motive stage of the NHI hypothesis.

If a future body of evidence establishes that a class of UAP represents NHI technology, the central strategic question immediately becomes:

$$\boxed{\text{Why is it here?}} \tag{104}$$

This is the *bona fide* question.

9.3.1 Why Unknown Capability Produces Strategic Fear

An actor whose intent is unknown creates uncertainty.

If the actor also appears technologically superior, the uncertainty becomes more consequential.

A system capable of behavior outside familiar human technological envelopes could in principle possess a wide range of available actions.

Threat analysis therefore has rational reasons to consider adverse intent.

The analytical error would be to treat the threat possibility as a demonstrated motive.

Capability asks

$$\text{what can the system do?} \tag{105}$$

Intent asks

$$\text{what is the system trying to do?} \tag{106}$$

These remain different variables.

Accordingly,

$$\boxed{\text{high capability} \not\Rightarrow \text{unique hostile intent.}} \tag{107}$$

Intent remains underdetermined in either direction.

The analytical task is to identify which motives remain rational under the assumed capability.

9.3.2 A Darwinian Motive

The present model supplies one such motive.

The motive is independent of culturally specific ideas of enlightenment, moral superiority, or spiritual transformation.

It follows directly from persistence.

A technologically advanced lineage still has a fundamental continuity interest.

For an organic or hybrid NHI, this can appear as the need for adaptive biological novelty.

For a post-biological AI, it can appear as the need for robust continuation solutions that cannot be universally pre-certified.

The general objective is

$$\boxed{\text{survive an open-ended future.}} \tag{108}$$

9.3.3 How Survival Generates a Preservation Strategy

Human threat models often associate survival with competition:

$$\text{their survival} \quad \text{versus} \quad \text{our survival.} \tag{109}$$

The present model identifies a different strategic resource.

The valuable object may be

$$\boxed{\text{independently generated evolutionary information.}} \tag{110}$$

A living biosphere has continuing value because it continues to generate, filter, and preserve new biological configurations.

Large-scale destruction of that biosphere would destroy part of the very solution reservoir that makes the planet valuable.

Preservation can therefore follow from self-interest.

Observation can follow from self-interest.

Sampling can follow from self-interest.

Long-term monitoring can follow from self-interest.

The sufficient condition is instrumental:

continued existence of the observed system has positive value to the observer's own continuation objective.

The distinction is important:

$$\boxed{\text{self-interest} \not\Rightarrow \text{destruction.}} \tag{111}$$

Under the evolutionary-witness model,

$$\text{self-interest} \to \text{preservation} \tag{112}$$

can be rational.

9.4 Nuclear Capability, UAP, and the Maximum Capability–Intent Problem

The nuclear-UAP discussion is where the emotional and strategic stakes of the model become most visible.

It is also where evidentiary discipline is most important.

Official U.S. reporting requirements include UAP associated with strategic nuclear weapons and nuclear-powered ships and submarines [27].

AARO's historical review records claims from former U.S. Air Force personnel of UAP sightings near ICBM facilities and alleged disruption of launch-control systems [26].

Recent reporting also documents UAS incidents near nuclear infrastructure while explicitly distinguishing those incidents from UAP [27].

These are different evidentiary categories.

They remain different in this White Paper.

No step in the mathematical argument establishes that NHI has interfered with nuclear weapons.

No historical claim is converted here into proof of non-human origin.

The relevant question is conditional.

Suppose, for purposes of the model, that an external intelligence does observe Earth.

Suppose further that Earth's independent evolutionary history has strategic option value to that intelligence.

What terrestrial development would rationally become especially salient?

The emergence of a biosphere-generated species capable of rapidly damaging the biosphere is an obvious candidate.

Nuclear technology is important in that sense for two reasons.

It demonstrates a threshold of technological capability.

And it demonstrates a threshold of destructive capability.

Humanity becomes, simultaneously,

more interesting

and

more consequential.

Observation and intervention remain distinct strategies.

Intervention itself changes the trajectory.

The model therefore admits a spectrum of rational responses:

observation,

monitoring,

sampling,

risk assessment,

limited intervention,

or non-intervention.

The mathematical framework constrains the motive space while leaving the specific policy choice underdetermined.

Strategic observation is compatible with more than one intent class.

Two possibilities are

$$\text{strategic observation} \Rightarrow \text{hostile preparation} \tag{113}$$

and

$$\text{strategic observation} \Rightarrow \text{biosphere-risk monitoring.} \tag{114}$$

The capability evidence alone does not discriminate between them.

The preservation motive is generated here by the more completely self-interested actor.

An observer whose own future option value depends partly on the continued existence of independent evolutionary histories has a strategic reason to notice technologies capable of destroying them.

The same reasoning also gives it a reason to notice humanity itself.

We are not merely the species holding the nuclear weapons.

We are an independently evolved technological lineage whose future development is part of the terrestrial historical process.

This produces a genuine tension:

$$\boxed{\begin{array}{c}\text{humanity has value as an independent evolutionary trajectory}\\ +\\ \text{humanity has power to alter the wider evolutionary reservoir.}\end{array}} \tag{115}$$

Observation therefore follows naturally from the dual status of humanity within the model.

HUMANITY'S DUAL STATUS

Humanity is simultaneously an independent evolutionary witness, a source of biosphere-scale technological risk, and an emerging manager of the process that produced it. For a preservation-oriented observer, this makes our technological evolution unusually salient: we are part of the source and increasingly capable of deciding its future.

10. THE *BONA FIDE*: WHY EARTH?

The complete motive chain can now be stated.

For an organic or hybrid NHI:

$$\boxed{\begin{array}{c}\text{extreme technological capability}\\ \Downarrow \\ \text{suppression of ordinary environmental pressure}\\ \Downarrow \\ \text{reduced endogenous evolutionary exploration}\\ \Downarrow \\ \text{greater dependence on deliberate adaptation}\\ \Downarrow \\ \text{Halting blocks universal Best-of-Evolution prediction}\\ \Downarrow \\ \text{physical computation retains undecidable certification boundaries}\\ \Downarrow \\ \text{independent evolutionary witnesses retain value}\\ \Downarrow \\ \text{external biosphere exploration becomes}\\ \text{a necessary survival strategy within the model.}\end{array}} \tag{116}$$

For post-biological intelligence, the more general chain is

$$\boxed{\begin{array}{c}\text{self-modifying intelligence}\\ \Downarrow \\ \text{Best-of-Continuation problem}\\ \Downarrow \\ \text{universal successor prediction encounters computability boundaries}\\ \Downarrow \\ \text{self-awareness can include awareness of those boundaries}\\ \Downarrow \\ \text{independent external solution spaces acquire option value}\\ \Downarrow \\ \text{preservation and exploration become rational strategies.}\end{array}} \tag{117}$$

The biological and post-biological paths therefore share a common structure:

$$\boxed{\text{universal foresight is bounded; independent successful realization retains value.}} \tag{118}$$

10.1 Why Earth Can Remain Valuable

Within the present model, Earth possesses a value independent of unique raw materials, human technological superiority, or shared human moral concepts:

$$\boxed{\text{Earth contains an independently evolved biosphere.}} \tag{119}$$

Billions of years of terrestrial evolutionary history have generated a continuing collection of physically realized biological witnesses.

For an organic or hybrid intelligence whose endogenous evolutionary pressure has been technologically suppressed, that history has direct adaptive relevance.

For a self-aware post-biological intelligence, the biosphere can additionally represent an independent, irreversible, physically realized solution space whose future utility cannot necessarily be exhausted in advance.

10.2 The Rational Answer to the Intent Question

The question

$$\text{“Why would they come here?”} \tag{120}$$

therefore admits a constrained answer:

> Because even an intelligence with technology far beyond ours can retain a survival interest in solutions that it did not generate, cannot universally predict, and can observe already realized in independent physical history.

This is a candidate *bona fide*.

This candidate *bona fide* is a motive model for the NHI hypothesis.

NHI presence, specific UAP attribution, and empirical origin remain separate observational questions.

The motive derivation reduces one part of the unexplained intent space to a familiar and non-anthropocentric objective:

$$\boxed{\text{survival through access to independent solutions.}} \tag{121}$$

10.3 The Multi-NHI Extension of the Bona Fide

If more than one mature NHI lineage exists, the answer to “Why Earth?” becomes recursively stronger. Earth can supply

$$\boxed{\begin{aligned} V_{\text{Earth}} = V_{\text{biosphere}} &+ V_{\text{human trajectory}} \\ &+ V_{\text{independent lineages}} + V_{\text{interaction}}. \end{aligned}} \tag{122}$$

The first visitor can come for the source. A later visitor can value both the source and the independent technological histories already converging upon it. The location becomes a strategic ecology. Multiplicity here remains a deep-time model extension rather than an empirical attribution. It is a coherent extension of the same motive pipeline rather than a separate theory.

THE BONA FIDE IN ONE SENTENCE

A mature, instrumentally rational, self-interested intelligence can have a survival reason to observe and preserve Earth because universal foresight remains bounded while Earth's independent biosphere, human technological trajectory, and any independent lineages converging upon them constitute realized histories whose future continuation value cannot be exhaustively known in advance.

11. SCOPE, ASSUMPTIONS, AND BOUNDARY CONDITIONS

The argument combines established results, conditional transfers, model premises, and an interpretive application.

Their boundaries determine the exact strength of the White Paper.

11.1 The Halting Result Is Established

The undecidability of the classical Halting Problem is an established theorem of computability theory [11].

11.2 Rice-Type Semantic Undecidability Is Established

Rice's theorem provides a broader semantic undecidability result for nontrivial properties of arbitrary partial computable functions [12].

Its use here supports the generalized Best-of-Continuation discussion.

It does not imply that every practical AI self-modification is undecidable.

11.3 Spectral-Gap Undecidability Is Established

Spectral-gap undecidability is established for particular families of local quantum Hamiltonians [13, 14].

The theorem concerns family-level asymptotic classification. Many individual finite systems remain tractable.

11.4 Formal Incompleteness Is Instance-Specific

Spectral-gap incompleteness provides the formal-system layer [15].

Application to ZFC uses the relevant consistency and expressiveness assumptions.

The existence of such instances does not imply that arbitrary physical questions are independent of ZFC.

11.5 The CPC Transfer Condition

Representation Undecidability becomes a theorem about a particular CPC architecture when an effective embedding such as Eq. (33) exists.

CPC itself is not asserted to be a hypercomputer.

Changing the physical representation of computation does not erase classical computability theory.

11.6 The Biological Transfer Condition

The biological extension requires an effective realizability reduction such as Eq. (54).

Biological undecidability follows exactly when the required effective realizability reduction exists.

Ordinary finite biological viability questions may remain tractable.

11.7 Organic Stagnation Is an Idealized Limit

Technological buffering does not make mutation, drift, or every evolutionary mechanism disappear.

The model concerns the suppression of externally imposed adaptive filtering.

11.8 Scope of Exploration Necessity

The necessity conclusion depends upon assumptions A1–A5.

Its precise claim is that external exploration is necessary for continued access to independently evolved witness sets while preserving the protected home environment.

Exploration is not guaranteed to solve every future survival problem.

11.9 Post-Biological AI Changes the Target

A fully post-biological AI does not automatically need genetic or biological solutions.

Its direct problem is continuation, not organic evolution.

The Earth-biosphere necessity conclusion is therefore strongest for organic or hybrid NHI.

A pure AI inherits the more general computability and preservation logic, while a specific interest in biology requires an additional utility premise.

11.10 Self-Awareness within the Model

The White Paper does not attempt to settle whether machine consciousness or self-awareness is possible.

It does not rely on the strong Penrose conclusion.

A self-aware AI is compatible with the model.

Indeed, recognition of its own non-universal predictability can strengthen the preservation rationale.

11.11 The Survivorship Threshold Is a Selection Premise

Deep-space capability establishes, within the model, a continuity-selection premise:

some biological, social, institutional, or machine continuity mechanism has persisted long enough to support the observed capability.

Wisdom and moral character remain separate variables.

11.12 The Resource-Value Inversion Requires Advanced Fabrication

CPC alone does not make raw material irrelevant.

The inversion requires the separate premise of highly advanced material design, fabrication, recycling, and local resource use.

Matter and energy remain physically necessary.

The argument concerns the declining *uniqueness* of ordinary terrestrial raw materials.

11.13 The Nuclear-UAP Record Is Heterogeneous

Official reports document

statutory nuclear-UAP reporting categories,

recent UAS activity near nuclear facilities,

and historical claims of UAP-related missile disruption.

These evidentiary classes must remain separate.

The White Paper uses them only to examine how the preservation model could alter an intent inference.

They are not used as proof of NHI origin.

11.14 Scope of the UAP Intent Model

The computational argument does not identify the origin of any observed UAP.

The UAP section addresses motive under a hypothetical NHI premise.

Its result is therefore an admissible intent model, not an attribution.

11.15 The Multi-NHI Extension Is Recursive and Model-Relative

The strategic-ecology section assumes that more than one mature, causally independent technological lineage exists and that each satisfies the relevant actor-class premises. It does not infer multiplicity from UAP reporting. Its claim is structural: if multiple mature lineages exist, the Historical-Independence Invariant can apply recursively to the lineages themselves, and interaction among them can create additional realized novelty.

11.16 The 2026 Bulk-Gap Result Narrows the Spectral Statement

The present spectral-gap argument concerns the specific undecidable family formulations used by Cubitt–Perez-Garcia–Wolf and Bausch et al. [13, 14]. Xu et al. show that the bulk spectral gap is semi-decidable under their bulk-gap formulation [30]. The White Paper therefore treats “spectral-gap undecidability” as theorem-specific language, not as a blanket statement about every spectral-gap notion.

12. CONCLUSION: THE DARWINIAN INVERSION

The White Paper began with the strongest strategic premise available to the NHI capability–intent problem. Grant a mature technological actor extreme capability. Infer the minimum instrumental rationality required to sustain that capability. Retain a Darwinian continuation interest. Do not introduce benevolence.

The first consequence is a disease of civilization. A biological lineage that becomes increasingly successful at buffering itself from disease, scarcity, and uncontrolled environmental selection also reduces the external adaptive search that once filtered its future. Survival becomes more dependent on deliberate adaptation.

The obvious response is prediction. Use better science, AI, quantum computing, physical computation, simulation, and design to calculate what will be needed next. The White Paper grants all of them. The universal objective still fails. Within effective computation, the Halting Problem is a structural boundary, not a shortage of processors or energy. Rice extends the semantic boundary. Spectral-gap constructions show that undecidable classification can be embedded in quantum many-body families. Formal incompleteness then yields specific propositions that sufficiently expressive consistent theories such as ZFC cannot settle under the relevant construction. CPC can inherit the same obstruction when an explicit property-preserving effective embedding exists. The representation changes. The universal certification problem remains.

NO UNIVERSAL FORESIGHT

Technological superiority and universal foresight are different categories. An intelligence can know almost unimaginably more than humanity while still lacking a total method that decides every admissible future case. The mature intelligence can therefore know something strategically important: it can know that it cannot know everything in advance.

That boundary changes the value of history. A biosphere is not only matter. It is a realized search through physical possibility. Its present state contains the residue of billions of years of mutation, selection, recombination, drift, symbiosis, competition, climate, geology, extinction, and contingency. Advanced fabrication can reproduce matter and may reproduce or surpass individual functions. It cannot retrospectively make an independent history part of the observer's own past. Before future continuation needs are known, that independent provenance retains option value.

Prediction therefore gives way to a portfolio strategy. Predict what can be predicted. Design what can be designed. Simulate what can be simulated. And harvest realized search from histories that occurred outside the observer's own model, training history, lineage, and experiment. The biosphere is most valuable not when exhausted, but when it remains a living generator. Observation, sampling, and preservation can therefore be forms of self-interested resource strategy.

This is the Darwinian inversion. The continuation objective has not softened. The scarcity landscape has changed.

Earth is unusually rich in precisely the object this model makes valuable: an ongoing biosphere and an independently evolved technological species that has begun to alter its own evolutionary future. Humanity is therefore not outside the resource. We are part of it. We are simultaneously witness, hazard, and emerging biospheric manager. Our current transition—from biological evolution to deliberate technological continuation—has trajectory value of its own.

If multiple mature NHI lineages exist in deep time, the argument becomes recursive.

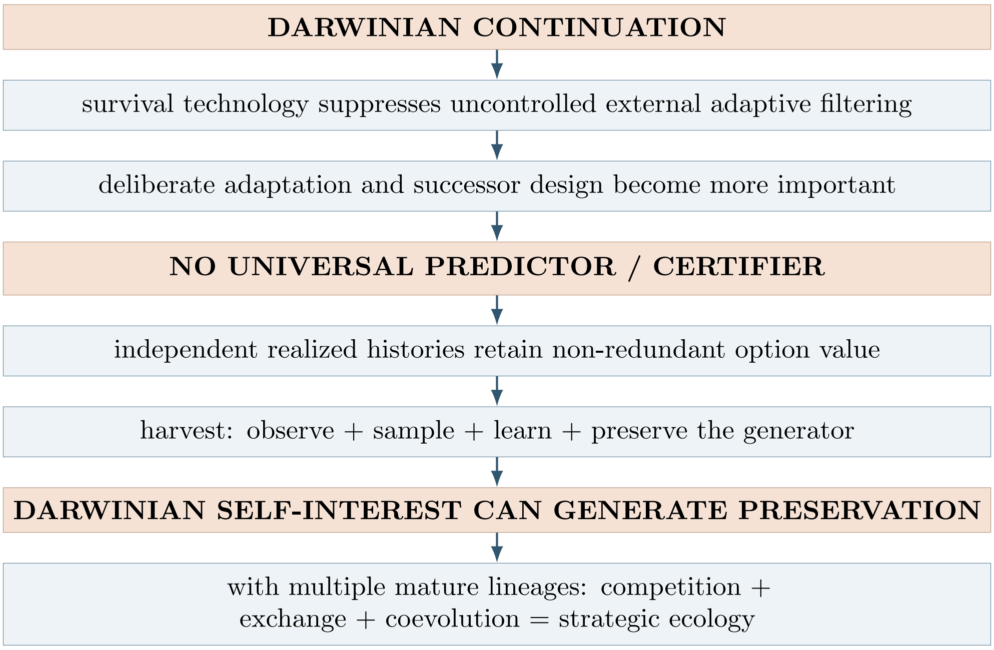

Figure 5: Final synthesis. The pipeline begins and ends with Darwinian continuation; the intermediate computability boundary changes the rational value of independent realized history.

Another civilization is not merely another competitor. It is another independent continuation history. Lineages meeting around a valuable living world become part of one another's environment. The resulting regime can contain conflict, monitoring, avoidance, exchange, mutualism, symbiosis, and coevolution. A world can become a convergence node: actors arrive because of the source, and then the actors themselves become part of the source of novelty.

The White Paper therefore gives a constrained answer to its title question.

WHY EARTH?

Because technology can remove many forms of scarcity without removing the universal prediction boundary. Because an independent living world contains realized solutions that another lineage did not generate. Because humanity is itself an evolving technological outcome of that independent history. And because, if more than one mature lineage exists, a living technological world can become a convergence node where independent intelligences learn from the source and from one another.

The same Darwinian premise that first makes an unknown superior intelligence frightening can, at the technological frontier, make observation and preservation rational.

The fear and the restraint arise from the same assumption.

REFERENCES

[1] Daghbouche, K. (2024).

The Bona Fide of Non-Human Intelligence (NHI) Exploring Earth: Evolutionary Pressure and the Halting Problem.

JAcad. (N.Y.), 13(2), 3–10.

IPFS CID:QmeWZbcbyXGDXG5rSDpQWGF1uz7bmsCsvivVv9WqytsZMq

[2] Daghbouche, K. (2025).

Computational Complexity of UAP Reverse Engineering: A Formal Analysis of Automaton Identification and Data Complexity.

arXiv:2505.00051

[3] Daghbouche, K., & Duman, D. (2026).

Configuration Identity as a Formal Parameter in Deterministic State Counting (Version v1). Zenodo. doi:10.5281/zenodo.22062716

[4] Daghbouche, K., & Duman, D. (2026).

Canonical Quotient-State Realization of Bounded-Width Dynamic Programming in Layered Tseytin 3-CNF (Version v1).

Zenodo. doi:10.5281/zenodo.22062774

[5] Daghbouche, K. (2026).

Response Quotients for Exact Semantic Representation (Version v1).

Zenodo. doi:10.5281/zenodo.22062855

[6] Daghbouche, K. (2026).

Fixed-Function Exact Semantic Carrier Conditions over VPV (Version v1).

Zenodo. doi:10.5281/zenodo.22062943

[7] Daghbouche, K. (2026).

One-Way Permutations and the Conditional Unprovability of Exact Semantic Carrier Separation in PV (Version v1).

Zenodo. doi:10.5281/zenodo.22063029

[8] Daghbouche, K., & Duman, D. (2026).

Exact Semantic Carriers for Structured Physical Systems (Version v1).

Zenodo. doi:10.5281/zenodo.22063161

[9] Daghbouche, K. (2026).

A Mathematical Framework for Carrier-Based Physical Computation: Constraint Compilation, Engineered Dissipation, and Microwave-Referenced Readout (Version v1). Zenodo. doi:10.5281/zenodo.22063338

[10] Daghbouche, K. (2026).

Representation Precedes Solvability: An Epistemology of Knowledge, Computation, and Natural Problems (Version v1).

Zenodo. doi:10.5281/zenodo.22063467

[11] Turing, A. M. (1936).
On Computable Numbers, with an Application to the Entscheidungsproblem.
Proceedings of the London Mathematical Society, s2-42(1), 230–265.
doi:10.1112/plms/s2-42.1.230

[12] Rice, H. G. (1953).
Classes of Recursively Enumerable Sets and Their Decision Problems.
Transactions of the American Mathematical Society, 74(2), 358–366.
doi:10.1090/S0002-9947-1953-0053041-6

[13] Cubitt, T. S., Pérez-García, D., & Wolf, M. M. (2015).
Undecidability of the Spectral Gap.
Nature, 528, 207–211.
doi:10.1038/nature16059

[14] Bausch, J., Cubitt, T. S., Lucia, A., & Pérez-García, D. (2020).
Undecidability of the Spectral Gap in One Dimension.
Physical Review X, 10, 031038.
doi:10.1103/PhysRevX.10.031038

[15] Cubitt, T. S. (2021).
A Note on the Second Spectral Gap Incompleteness Theorem.
arXiv:2105.09854

[16] Aharonov, D., van Dam, W., Kempe, J., Landau, Z., Lloyd, S., & Regev, O. (2007).
Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation.
SIAM Journal on Computing, 37(1), 166–194.
doi:10.1137/S0097539705447323

[17] Eaton, S. B., Konner, M., & Shostak, M. (1988).
Stone Agers in the Fast Lane: Chronic Degenerative Diseases in Evolutionary Perspective.
The American Journal of Medicine, 84(4), 739–749.
doi:10.1016/0002-9343(88)90113-1

[18] Carrera-Bastos, P., Fontes-Villalba, M., O'Keefe, J. H., Lindeberg, S., & Cordain, L. (2011).
The Western Diet and Lifestyle and Diseases of Civilization.
Research Reports in Clinical Cardiology, 2, 15–35.
doi:10.2147/RRCC.S16919

[19] Bengio, Y., Clare, S., Prunkl, C., Murray, M., et al. (2026).
International AI Safety Report 2026.
International AI Safety Report 2026

[20] Google Quantum AI and Collaborators (2025).
Quantum Error Correction Below the Surface Code Threshold.
Nature, 638, 920–926.
doi:10.1038/s41586-024-08449-y

[21] *Nature Biotechnology* (2026).
Quantum Computing in Transition.
Nature Biotechnology, 44, 1065–1066.
doi:10.1038/s41587-026-03233-x

[22] Alexeev, Y., Farag, M. H., Patti, T. L., et al. (2025).
Artificial Intelligence for Quantum Computing.
Nature Communications, 16, 10829.
doi:10.1038/s41467-025-65836-3

[23] Penrose, R. (1989).
The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics .
Oxford University Press.

[24] Penrose, R. (1994).
Shadows of the Mind: A Search for the Missing Science of Consciousness .
Oxford University Press.

[25] Feferman, S. (1996).
Penrose's Gödelian Argument.
PSYCHE, 2, 21–32.

[26] All-domain Anomaly Resolution Office (2024).
Report on the Historical Record of U.S. Government Involvement with Unidentified Anomalous Phenomena (UAP): Volume I .
U.S. Department of Defense.
AARO Historical Record Report, Volume I

[27] All-domain Anomaly Resolution Office (2026).
Fiscal Year 2025 Consolidated Annual Report on Unidentified Anomalous Phenomena .
U.S. Department of War.
Fiscal Year 2025 UAP Annual Report

[28] Office of the Director of National Intelligence (2021). *Preliminary Assessment: Unidentified Aerial Phenomena*. 25 June 2021. ODNI Preliminary Assessment.

[29] Daghbouche, K. (2026). *CPC Validation* (Version v0.6.0) [Computer software]. Zenodo. doi:10.5281/zenodo.22045319. Code repository: github.com/GridSAT/cpc-validation.

[30] Xu, X., Schötz, M., Wang, J., Magron, V., Klep, I., Fawzi, O., & Renou, M.-O. (2026). The bulk spectral gap is semi-decidable: a convergent family of certified upper bounds. arXiv:2606.03836.

[31] Day, T. (2012). Computability, Gödel's incompleteness theorem, and an inherent limit on the predictability of evolution. *Journal of the Royal Society Interface*, 9(69), 624–639. doi:10.1098/rsif.2011.0479.

[32] Hernández-Orozco, S., Hernández-Quiroz, F., & Zenil, H. (2018). Undecidability and Irreducibility Conditions for Open-Ended Evolution and Emergence. *Artificial Life*, 24(1), 56–70. doi:10.1162/ARTL_a_00254.

[33] Eberbach, E. (2005). Toward a theory of evolutionary computation. *BioSystems*, 82(1), 1–19. doi:10.1016/j.biosystems.2005.05.006.

[34] Lane, W. C. (2025). The extraterrestrial hypothesis: an epistemological case for removing the taboo. *European Journal for Philosophy of Science*, 15, 8. doi:10.1007/s13194-025-00634-8.

[35] Ball, J. A. (1973). The Zoo Hypothesis. *Icarus*, 19(3), 347–349. doi:10.1016/0019-1035(73)90111-5.

[36] Freitas, R. A., Jr. (1983). The Search for Extraterrestrial Artifacts (SETA). *Journal of the British Interplanetary Society*, 36, 501–506. Author archive.

[37] Fogg, M. J. (1987). Temporal aspects of the interaction among the first galactic civilizations: The "Interdict Hypothesis". *Icarus*, 69(2), 370–384. doi:10.1016/0019-1035(87)90112-6.

[38] Cabrol, N. A. (2016). Alien Mindscapes—A Perspective on the Search for Extraterrestrial Intelligence. *Astrobiology*, 16(9). doi:10.1089/ast.2016.1536.

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

[40] Dick, S. J. (2003). Cultural evolution, the postbiological universe and SETI. *International Journal of Astrobiology*, 2(1), 65–74. doi:10.1017/S147355040300137X.

[41] Haqq-Misra, J. D., & Baum, S. D. (2009). The Sustainability Solution to the Fermi Paradox. *Journal of the British Interplanetary Society*, 62(2), 47–51. arXiv:0906.0568.

[42] Frank, A., Grinspoon, D., & Walker, S. I. (2022). Intelligence as a planetary scale process. *International Journal of Astrobiology*, 21(2), 47–61. doi:10.1017/S147355042100029X.

[43] Baum, S. D., Haqq-Misra, J. D., & Domagal-Goldman, S. D. (2011). Would contact with extraterrestrials benefit or harm humanity? A scenario analysis. *Acta Astronautica*, 68(11–12), 2114–2129. doi:10.1016/j.actaastro.2010.10.012.

[44] Forgan, D. H. (2017). The Galactic Club, or Galactic Cliques? Exploring the limits of interstellar hegemony and the Zoo Hypothesis. *International Journal of Astrobiology*, 16(4), 349–354. doi:10.1017/S1473550416000392.

[45] Jebari, K., & Asker, A. S. (2024). Saved by the Dark Forest: How a Multitude of Extraterrestrial Civilizations Can Prevent a Hobbesian Trap. *The Monist*, 107(2), 176–189. doi:10.1093/monist/onae006.

[46] Giakas, E. (2026). *The Thermodynamic Peace Framework: A Conditional Constraint Model for Spacefaring Civilizations* (Final Draft v12) [Preprint]. Zenodo. doi:10.5281/zenodo.20575006.

[47] Thomson, S. (2012). On the Frontier of Knowledge: A Discussion of Alien Civilizations. *Physics Procedia*, 38, 304–318. doi:10.1016/j.phpro.2012.08.030.

SUBJECT INDEX