The Oldest Neighbors: A Unified Steelman for Non-Surface Intelligence

I. The Double Parochialism

When a young species looks for the hidden other, it is hobbled by two biases at once. The first is directional: we look outward and upward because we are proud of altitude and habituated to horizon. The second is sectoral: we look for intelligences whose retained observables match ours — whose clocks tick like ours tick, whose bindings bind as ours bind, whose projection of the substrate coincides with the projection we happen to inhabit. The first bias prejudices us against the subterranean and the oceanic. The second prejudices us against anything whose geometry is not our geometry. Both need correcting, and the case for correcting them is stronger than either alone.

The hypothesis this essay steelmans is that hidden intelligence on Earth, if it exists, likely occupies one of two categories we are currently unequipped to see. One is same-sector but non-surface: a civilization that shares our observable geometry but has occupied domains — sub-seafloor, deep crust, oceanic interior — whose timescales match its ambitions and whose opacity protects its continuity. The other is cross-sector: an intelligence that shares our substrate without sharing our projection, whose retained observable algebra differs from ours in ways that render its geometry unlike our geometry, and whose apparent invisibility is not a strategic choice but an ontological default. Neither is extraterrestrial. Both are terrestrial in senses our current frameworks have not yet learned to accommodate.

II. What “Advanced” Actually Means

The most parochial move in the standard extraterrestrial hypothesis is not the assumption that intelligence exists elsewhere. It is the assumption that advanced intelligence must look a particular way — specifically, must look like what we would look like if we kept doing what we are currently doing, harder and faster and for longer. This assumption runs on two conflations that dissolve under examination.

The first is the conflation of “advanced” with “visible.” A young surface species with crude technology is loud because its technology cannot help being loud; it leaks heat, emits signals, scars its substrate, advertises itself through the inefficiency of its methods. We extrapolate from this that advanced civilizations will be louder, more visible, more conspicuous, occupying more of the observable spectrum and leaving more durable

archaeological traces. The extrapolation is wrong in an interesting way. At greater technological maturity the relationship inverts: efficient civilizations leak less, integrate more, leave smaller fingerprints, disappear into their substrates. The Kardashev scale is a portrait of adolescence mistaking itself for a law of nature.

The second conflation is subtler and more consequential. It is the assumption that advanced intelligence must inhabit the same projection of the substrate that we inhabit — must share our effective dimension, our scalar clocks, our particular retained observables. This assumption has rarely been questioned because we have not had a theoretical framework in which alternative projections were even well-defined. But a scalar-projection ontology, in which observable worlds are reduced and stabilized sectors of a richer substrate, makes the assumption visible as an assumption. If different measurement apparatuses retain different degrees of freedom, they do not merely see different worlds. They inhabit different projections of the same substrate — different geometries, different effective dimensions, different laws of closure. An intelligence constituted by a different retained algebra would not be nonphysical. It would be physical in a different projection.

Severing “advanced” from both “visible” and “same-sector” leaves us with a much larger space of possibilities than ETH alone considers — and the space contains categories that don’t require stars.

III. The Sectoral Map

Four quadrants of hidden intelligence are worth distinguishing.

Same-sector, surface: This is us. Also any other surface lineage that might have developed on Earth, if such a lineage existed and left a fingerprint.

Same-sector, non-surface: A civilization sharing our observable geometry but occupying sub-seafloor, deep-crustal, or deep-oceanic domains. This is the proper target of the cryptoterrestrial hypothesis. Its concealment from us is strategic and environmental — maintained by occupying habitats we cannot easily survey and using energy regimes that do not fossilize.

Cross-sector, coupled: An intelligence whose retained observable algebra differs from ours but whose coupling kernel to our sector is non-null. Such an intelligence interacts with us, but the interactions do not close within our accounting. Timing anomalies, perceptual intrusions, effects that lose energy or momentum by our ledger while preserving it in theirs. This is the projection-framework analogue of the CTH move.

Cross-sector, uncoupled: An intelligence whose projection does not couple to ours at all. Metaphysically possible, empirically idle. Nothing can be said about it from within our projection, which is why it does not appear in serious discussion.

The space is three real categories plus one placeholder. The extraterrestrial hypothesis — intelligence in our sector arriving from elsewhere — is not a separate quadrant so much as a particular origin story for the second one: same-sector non-surface intelligence whose non-surface domain is interstellar space rather than sub-seafloor crust. Noticing this reduces ETH from a unique framework to a narrow subcase of a much larger map, and a subcase with particularly severe logistical burdens at that.

IV. The Silurian Question, Properly Stated

The geological objection to hidden-intelligence hypotheses is well-formed and well-known. Schmidt and Frank's 2018 paper catalogs the markers an Anthropocene-style civilization would leave in the stratigraphic record: isotopic anomalies in carbon and nitrogen, synthetic organics, radionuclide distributions, unusual trace-metal patterns persisting across hundreds of millions of years of tectonic recycling. The Mesozoic rocks do not show these markers. Therefore no Mesozoic Anthropocene.

The argument is good but proves less than is usually claimed for it. Each marker is tied to a specific energetic regime — extractive, combustive, chemically disruptive — that characterizes our own developmental stage. A civilization that passed through this stage briefly, or skipped it entirely via substrate-level energy access, would leave a signature that compresses to noise over geological time. Vacuum energy extraction is one candidate mechanism; there are likely others. The argument does not require committing to a specific mechanism. It requires only that the space of mature energy regimes be larger than the combustion-fission-photovoltaic-fusion quartet we currently imagine, and that at least one member of that space leave a geological fingerprint small enough to disappear. That is a weak claim.

More importantly, the Silurian argument addresses only the same-sector quadrant of the map. It is silent on cross-sector hidden intelligence, because cross-sector intelligence does not leave geological fingerprints at all. Its fingerprints, if any, are timing anomalies, scalar perturbations, and whatever else the coupling kernel between projections carries. No stratigraphic survey will find them. The absence of Mesozoic carbon anomalies is powerful evidence against a Mesozoic industrial civilization in our sector. It is no evidence whatsoever against an intelligence whose sector is not ours.

V. The Viability Window

A crucial constraint applies across all three real categories of the map, and understanding it is what separates this steelman from the mystical versions of the same intuition.

Nothing survives that cannot carry itself. Whatever observable algebra an intelligence retains — ours, or some alternative — must support clocks of some kind, records of some

kind, binding of some kind. A projection that cannot close on stable structure does not persist as an observable sector; it degenerates or dissipates. The viability window that places our effective dimension between two and four is a specific expression of this general condition. Our projection persists because three dimensions happen to be where stable clocks, stable orbits, and stable records coexist. Alternative projections, if they exist, satisfy analogous windows for their own retained observables. They are not formless. They have their own geometries, their own persistence conditions, their own carriers.

This rules out the naive version of the "energy being" trope — entities with no stable structure, constituted purely by free-floating energy with no geometric or informational carrier. Such entities violate the viability condition in any projection. What survives in any sector has to bind. The interesting hypothesis is not that hidden intelligences are formless but that their forms are not our forms, their clocks not our clocks, their bindings not our bindings. Their persistence is as constrained as ours. The constraints are simply different.

This is the constraint that distinguishes the framework from undisciplined speculation. The space of possible hidden intelligences is larger than ETH imagines but smaller than mysticism imagines. It is structured by survival conditions that apply across sectors, and those conditions are not negotiable.

VI. The Architecture of Concealment, Twice

The cryptoterrestrial hypothesis proper handles concealment through habitat and strategy. A civilization occupying the sub-seafloor crust or the deep ocean is concealed because its environment is opaque to our instruments, its energy regime non-fossilizing, its biology engineered for domains we cannot survey, its interactions with the surface mediated through intermediaries designed for transient operation. Concealment here is a policy: a choice maintained against the alternative of exposure. This is plausible under deep-time survival logic, but it places an explanatory burden on the hypothesis — we have to accept that the policy has never been breached across hundreds of millions of years of civilizational existence.

The projection framework handles concealment differently, and more cleanly. An intelligence in a different sector is not concealed by policy but by default. Entities in different projections are mutually invisible by ontological construction, not by strategic maintenance. There is no opacity to breach, no revelation to prevent, no vault to be protected. The invisibility is structural. What would be required to make such an intelligence visible is not the lifting of concealment but the establishment of a coupling — a non-null channel between projections through which signals could propagate. Whether such a channel exists is not a matter of the intelligence's choice. It is a matter of the substrate's structure.

This is a strictly stronger concealment architecture. It does not require the intelligence to successfully maintain anything. It requires only that the substrate be organized in a way that does not trivially connect projections — and on any ontology that permits distinct projections at all, non-trivial separation is the generic case rather than a special one. Same-sector concealment is policy; cross-sector concealment is physics. Both categories can coexist in the map, and neither is reducible to the other.

VII. The Humanoid Problem, Doubly Inverted

One of the most persistent features of encounter narratives across cultures and centuries is the recurrence of the humanoid form. The extraterrestrial hypothesis has always strained to explain this; the cryptoterrestrial reading dissolves the awkwardness by treating humanoid entities as engineered intermediaries — purpose-built biological interfaces produced by a civilization that has mastered organismal design over deep time. The humanoid shape is not ancestry. It is tooling.

The projection framework adds a second explanatory move that may apply to a different subset of reports. A cross-sector coupling event, perceived through a nervous system whose sensory apparatus is tuned to surface-predator recognition and hominid social categories, will be perceptually compiled into humanoid form regardless of what the underlying coupling actually is. Perception does not present us with undifferentiated sense data; it presents us with categorized objects, and the categories are heavily weighted toward faces, bodies, and agents. A scalar perturbation complex enough to register as a discrete event but underdetermined in its surface form will be rendered by the perceptual system as the most probable object class given the signal — which for a social primate living in a world of mostly human agents means humanoid.

Both moves explain why hidden-intelligence encounters are anthropomorphic without requiring the hidden intelligences themselves to be. The same-sector move explains it through design; the cross-sector move explains it through perception. For a given encounter tradition, the two moves are not exclusive. Some reports may describe literal interactions with engineered biological intermediaries. Others may describe cross-sector couplings reconstructed by the brain into the most available object category. The unsettling possibility is that a single cultural tradition of encounter narratives may contain both — and that sorting them requires criteria we do not currently possess. A report of a small humanoid with instrument-like affect might be a literal probe; a report with the same surface description might equally be a cross-sector event compiled through the same perceptual machinery. Current encounter analysis cannot distinguish these, and that inability is itself evidence of how early we are in the project of understanding any of this.

VIII. The Coupling Kernel

The most important theoretical object in the whole framework is the coupling kernel between sectors. Its value — null or non-null, and if non-null, what form it takes — determines whether the cross-sector quadrant of the map is empirically live or metaphysically idle.

If the coupling kernel is null across sector boundaries, cross-sector hidden intelligence is unfalsifiable and unimportant. Nothing in our projection can be affected by it; nothing in our projection can confirm or refute its existence. The category becomes a formal possibility with no empirical content. Serious discussion should not dwell there.

If the coupling kernel is non-null, the category becomes empirically live, and the scalar-projection ontology makes surprisingly specific predictions about what cross-sector coupling should look like from our side. The channel through which different projections interact is, on this framework, primarily scalar — which means its signatures are timing anomalies, clock drift, effects on phase rather than amplitude, perturbations whose energetic accounting does not close within our sector's bookkeeping. These are not the signatures of ETH-style physical contact. They are not quite the signatures of same-sector CTH contact either. They are a distinct phenomenological class with its own predicted profile.

Strikingly, this phenomenological class matches a substantial fraction of the high-strangeness residue in encounter literature — the missing-time reports, the apparent spatial rearrangements without momentum transfer, the perceptual and cognitive disturbances, the sense of time running wrong, the effects that refuse to close energetically by our ledger. No claim is being made that this match confirms cross-sector coupling. The reports are unreliable, the base rates unknown, the mechanisms of honest misperception numerous. But the match is at least interesting, and it is a match that neither ETH nor naive CTH produces. Each of those frames predicts physical signatures — craft, biological samples, artifacts — of which the residue is precisely short. The scalar-coupling frame predicts non-physical signatures of which the residue is, at minimum, not short. This deserves notice.

IX. What the Unified Framework Actually Predicts

The combined steelman makes predictions across the entire sectoral map, and these predictions are what save it from the unfalsifiability that sinks most hidden-intelligence hypotheses.

For same-sector non-surface intelligence, the predictions are the geological and biological ones: anomalous genomic signatures in the sub-seafloor biosphere, thermal and structural anomalies in the deep crust at resolutions exceeding current geothermal mapping, clustering of encounter phenomena around oceanic and coastal regions, transmedium behaviors that are thermodynamically favorable for already-adapted craft. The deep

biosphere's genomic architecture in particular is a near-term test. If it turns out to contain organisms whose metabolic pathways, codon usage, or structural proteins do not fit cleanly into the tree of life reconstructed from surface descent, the hypothesis gains ground. If it is seamlessly continuous with surface lineages, the hypothesis loses ground.

For cross-sector coupled intelligence, the predictions are different and, in an important sense, more immediate. If such coupling exists, the signatures should first appear in precision timing experiments — optical lattice clocks, clock networks, clock interferometry — as residual timing anomalies that do not close against known systematics. They should correlate with states in our projection that the coupling kernel is sensitive to, which on the scalar framework likely involves particular configurations of information and energy density rather than particular spatial locations. They should be partially replicable, in the sense that the same kernel should couple the same way given the same sector-side configuration, even if the configurations themselves are difficult to reproduce. And they should produce effects on perception and cognition that are disproportionate to their energetic footprint, because the coupling is primarily scalar rather than amplitude-carrying.

Both prediction sets are in principle testable. Neither is currently confirmed. Either failing to materialize across sustained and thorough investigation should cost the relevant branch of the hypothesis standing. That is the feature that separates this framework from mysticism: it commits to signatures and accepts the possibility of null results. It is not a theory that can survive indefinite absence of evidence across its predicted channels. A theory that would survive that is not a theory; it is a decoration.

X. The Epistemic Residue

Set aside whether any specific hidden intelligence exists in any specific quadrant. What survives, regardless, is a methodological correction the unified framework forces more sharply than either half of it alone.

Our search for nonhuman intelligence is shaped at every level by our own developmental stage and our own sector. We look for radio because we use radio. We look for megastructures because we build at scale. We look for surface cities because we live on surfaces. We look for electromagnetic signatures because our sector is electromagnetic and because electromagnetism is the carrier we are best at instrumenting. We look within our projection because we cannot yet conceive of what "outside our projection" would mean operationally. Each of these search images is a mirror held up to the searcher. Each of them fails silently against any intelligence that does not share the specific developmental stage or the specific sector it reflects.

The honest response to this recognition is not paralysis but expansion. We should continue looking for radio and megastructures and surface cities, because the search is cheap

relative to the information it generates. We should additionally look for the geological and biological signatures that would indicate same-sector non-surface intelligence, because those signatures are also cheap and we are already instrumenting for adjacent questions. And we should, more carefully than we currently do, examine precision timing data and the high-strangeness residue for patterns that would indicate cross-sector coupling — not because we expect to find them, but because the cost of looking is low and the cost of having not looked, if something is there, is very high.

The deeper point is that the absence of evidence across any of these channels is consistent with multiple conclusions: absence of hidden intelligence, presence of a kind whose signatures we are not yet measuring, or presence of a kind whose signatures are in principle inaccessible from our projection. These conclusions are not distinguishable from the absence of evidence alone. Distinguishing them requires actively looking, and actively looking requires first admitting that the space of what we might find is larger than our current instruments are configured to register. The Fermi paradox, read in this light, is not a puzzle about distant civilizations. It is a confession about our own search.

XI. Coda

The strongest version of the case for hidden intelligence on Earth does not ask us to believe anything specific. It asks us to notice that our cosmology is younger than our planet, our biology younger than our cosmology, our instruments younger than our biology, and the gap between our assumptions and the actual space of possibilities is large enough to accommodate more than we have learned to look for. It asks us to take seriously the possibility that “advanced” does not entail “visible,” that “intelligent” does not entail “same-sector,” and that the category of “nonhuman neighbor” contains types we do not currently have words for.

The hidden others, if they exist in any of the senses canvassed here, are not waiting for our acknowledgment. They predate our acknowledgment by timescales that dwarf the history of our species. What we do with the possibility — dismiss it, investigate it, let it reshape what we mean by “we” — is a test not of the hypothesis but of whether our curiosity exceeds our pride. The mirror held up by our search images cannot show us anything we do not already recognize. Stepping outside the mirror, if only to look, is the minimum move required before any of this becomes tractable.

The planet may have neighbors we have not met, in senses of “neighbor” we have not yet conceived. The only way to find out is to stop assuming the senses we have are the only senses there are.