The Alien Bootstrap Hypothesis:
Interstellar Access, Logistics Constraints, and the Conditions for Extraterrestrial Presence

G. Kent

June, 2026

Abstract

The Fermi paradox is usually framed as a tension between the apparent ease of galactic colonization and the absence of observed colonizers. Most proposed resolutions modify either the *probability* that intelligence arises (the Drake factors) or the *motivation* of intelligence that does (the zoo, interdict, and sustainability solutions). This paper develops a third class of resolution, grounded in *logistics*. Its premise is that interstellar *access* does not entail interstellar *logistics*: a civilization may possess rare, costly, or payload-constrained transit capability without the ability to move large masses, populations, or infrastructure across interstellar distances. Under that premise a rational visitor minimizes transported mass, transports information instead, and waits until a target biosphere develops the local capacity to convert information into embodiment and infrastructure. We call the moment that capacity crosses threshold the *bootstrapping point* and the resulting model the *Alien Bootstrap Hypothesis* (ABH). We argue that the payload-limitation premise—the model's principal vulnerability—is defensible beyond analogy: in weak form it follows from the rocket equation and from the thermodynamic asymmetry between transporting bits and atoms; in strong form it is concretely instantiated in the energy-and-geometry constraints of metric-engineering propulsion solutions. The bootstrap argument requires only the weak form. We then identify candidate threshold technologies and connect the biological route to a companion rate law, the *Kent Equation*, which models when a civilization becomes able to fabricate a viable body to specification on its own surface and identifies *developmental viability*—not genome synthesis or machine intelligence—as the binding gate. Finally we treat the model's empirical content systematically: we enumerate the consequences the ABH generates, classify each by evidential weight, and use the Kent rate law to argue that the model's content lives in the *timing and ordering* of technological crossings rather than in their co-movement. We separate the model's genuinely risky predictions—above all that disclosure pressure should track the binding biological gate—from the larger set of post-hoc-compatible "signatures," and we are explicit that the model's central risk is overexplanation. The contribution is not a claim that the hypothesis is true, but a demonstration that logistics is a more tractable and more falsifiable axis of analysis than arrival.

1 Introduction

Enrico Fermi's question—"where is everybody?"—acquires its force from a chain of apparently modest assumptions. If technological civilizations are not vanishingly rare, if at least some of them expand, and if expansion across the Galaxy is feasible on a timescale short compared with the Galaxy's age, then the Galaxy should already be colonized, and Earth should already be colonized. It is not. The argument was sharpened by Hart (Hart, 1975) and, in its strongest deterministic form, by Tipler (Tipler, 1980), who held that self-reproducing probes would have saturated the Galaxy long ago, so that the absence of such probes implies the absence of extraterrestrial intelligence (ETI). The same colonization-time argument implies a temporal point that is easy to overlook: any ETI capable of reaching us would, on cosmological timescales, have been able to do so long before our own emergence (Hair, 2011). Arrival is not a recent possibility; it is, on the standard assumptions, an ancient one.

Most resolutions of the paradox attack one of two links in the chain. The first family modifies the *rate* term—the probability that communicative or expansive intelligence arises at all—whether by appeal to the rarity of complex life, the fragility of technological civilizations, or selection effects in the distribution of arrival times (Sandberg et al., 2018). The second family grants that ETI exists and is capable of reaching us, and modifies its *behavior*: the zoo hypothesis holds that ETI deliberately abstains from contact to preserve Earth as a wilderness or laboratory (Ball, 1973; Forgan, 2011); the related planetarium hypothesis holds that the observed environment is itself managed (Baxter, 2001); the sustainability solution holds that fast expansion is self-limiting, so that no persistent galactic empire is expected to be observable (Haqq-Misra & Baum, 2009); and diffusion and percolation models hold that expansion is slow, patchy, and may halt well short of saturation (Newman & Sagan, 1981; Landis, 1998).

This paper develops a resolution in a third register. It does not principally modify the probability that ETI exists, nor the motivation of ETI that does. It modifies the implicit assumption—present in nearly all versions of the paradox—that a civilization able to *reach* a destination is thereby able to *deliver arbitrary mass and infrastructure* to it. We will call this assumption *capability monotonicity*: the supposition that interstellar reach implies interstellar omnipotence, that "if they can cross the stars, they can do anything." Capability monotonicity is rarely defended; it is usually assumed because the energies involved in any interstellar transit are so large that further engineering seems a rounding error. We argue that this intuition is mistaken, that the gap between *access* and *logistics* can be large and even widening, and that taking the gap seriously yields a distinct and unusually testable model of what extraterrestrial presence—if it exists—would look like.

We state the model informally first. Suppose a civilization can place a small, delicate, low-payload vehicle at another star, but cannot economically move cargo, population, ecosystems, or heavy infrastructure across the same distance. Such a civilization faces a logistics problem structurally identical to the problem of building in a remote wilderness: the capability to *arrive* does not confer the capability to *settle*, because settlement presupposes infrastructure that is precisely what is hard to transport. A rational actor in this position does not import the whole stack. It imports the smallest possible seed—information, instructions, perhaps minimal biological or material samples—and waits for the destination to develop, on its own, the local infrastructure that can convert that seed into embodiment, manufacturing, and presence. The threshold at which the destination acquires that converting capacity is the *bootstrapping point*. Before it, large-scale presence is impractical regardless of intent; after it, presence becomes practical without solving the transport problem at all. We call the resulting model the *Alien Bootstrap Hypothesis* (ABH).

Two clarifications fix the epistemic status of what follows. First, this is a constraint-solving exercise, not an evidential claim. We do not argue that ETI exists, that it has visited, or that any particular report of anomalous phenomena is veridical. We argue conditionally: *if* a payload-limited visitor existed, the bootstrap structure is what rational behavior would look like, and that structure has consequences we can examine. Second, the model's interest is asymmetric with respect to its tiers. Its premises and its broad consequences are, we will argue, reasonably defensible; its more colorful downstream embellishments (engineered human compatibility, the veridicality of abduction folklore) are weak, and are flagged as such throughout. The discipline the paper tries to maintain is to keep these tiers separate, and to identify which parts of the model could, even in principle, be wrong in a way we could detect.

The paper proceeds as follows. Section 2 fixes the taxonomy and distinguishes the ABH from the zoo hypothesis it superficially resembles. Section 3 defends the payload-limitation premise, the model's load-bearing and most contestable element. Section 4 states the bootstrap argument with its premises and consequences separated. Section 5 identifies the candidate threshold technologies and connects the biological route to the companion Kent Equation, whose central result—that the binding gate is developmental viability—resolves a standing question about which technology the model should track. Section 6 derives the interaction regime the model predicts. Section 7 confronts the model's empirical content head on: it enumerates the full catalog of derived signatures, grades each by evidential weight, and isolates the genuinely risky predictions from the rest. Section 8 locates the model among existing resolutions, and Section 9 catalogs its weaknesses.

2 Three Classes of Resolution and the Place of Logistics

It is useful to make the taxonomy explicit, because the ABH is easy to misread as a restatement of the zoo hypothesis when it is in fact a different kind of claim.

A resolution of the Fermi paradox must break at least one link in the inference "ETI is common ⇒ ETI expands ⇒ expansion saturates the Galaxy ⇒ Earth is colonized." We can classify resolutions by which link they break.

Rate resolutions break the first link. They argue that communicative or expansive ETI is rare enough that the expected number of observable colonizers is small. The Rare Earth program and the contemporary "dissolving" of the paradox via uncertainty in the Drake factors belong here (Sandberg et al., 2018). These resolutions are about *how many*.

Motivation resolutions break the third or fourth link by denying that capable ETI behaves as the argument assumes. The zoo and interdict hypotheses deny that ETI contacts us; the planetarium hypothesis denies that we perceive its presence; the sustainability solution denies that expansion persists (Ball, 1973; Baxter, 2001; Haqq-Misra & Baum, 2009). These resolutions are about *what they want*, or about *what expansion does to those who attempt it*.

Logistics resolutions break the second-to-third link: they deny that the ability to reach a destination implies the ability to project force, mass, or infrastructure to it at scale. Diffusion and percolation models are partly of this kind—they make expansion slow and incomplete by treating it as a stochastic spatial process rather than a deterministic wave (Newman & Sagan, 1981; Landis, 1998). The ABH is a logistics resolution of a sharper kind. It does not merely make expansion slow; it asserts a structural separation between the transit capability a civilization has and the delivery capability it lacks, and it works out the strategy that separation rationally implies.

The distinction from the zoo hypothesis is worth stating precisely, because the two models can produce superficially similar predictions (limited contact, careful non-disruption) for entirely different reasons. The zoo hypothesis is *motivational*: ETI *chooses* restraint, typically for ethical or scientific reasons, and could in principle abandon that restraint at any time, since it is assumed capable. The ABH is *constraint-driven*: restraint is, before the bootstrapping point, not a choice but a consequence of what is physically practical. The two models therefore diverge on a decisive question—*what changes the regime?* Under the zoo hypothesis, the regime changes when ETI's preferences or our "maturity" change, which are unobservable to us. Under the ABH, the regime changes when *our own technology* crosses a threshold, which is observable to us, because it is our technology. This is the source of the model's comparatively strong empirical content, and we return to it in Section 7.

3 The Payload-Limitation Premise

The entire model rests on a single premise: that a civilization can possess interstellar transit capability while lacking interstellar bulk-delivery capability, and that the gap is large enough to make local bootstrapping the rational strategy. This is the model's most vulnerable point, and it is tempting to leave it as a bare assumption defended only by analogy—the image of a craft as a tuned resonant system (a "bell") whose function degrades if too much is attached to it. An assumption defended only by analogy is a liability. The purpose of this section is to show that the premise is in fact defensible from established physics, and—more importantly—to identify exactly how much of it the argument actually needs.

We separate two forms of the premise. The *weak form* asserts only that moving large amounts of *matter* across interstellar distances is enormously more costly than moving *information*, costly enough that a rational payload-limited actor prefers to send information and instantiate locally. The *strong form* asserts the specific "bell" constraint: that the transit mechanism is sensitive to its own payload—its mass, geometry, or internal configuration—so that payload is limited not merely by energy budget but by the physics of the mechanism itself. We will argue that the weak form is robust and nearly unavoidable, that the strong form is plausible and concretely instantiated in at least one well-studied class of propulsion solutions, and that the bootstrap argument needs only the weak form.

3.1 The weak form: the rocket equation and the bit/atom asymmetry

Two independent and entirely conventional arguments establish the weak form.

(i) The tyranny of the rocket equation. For any reaction-propelled vehicle, the Tsiolkovsky relation $\Delta v = v_e \ln(m_0/m_f)$ ties the achievable velocity change to the natural logarithm of the mass ratio. The dependence is brutal in the direction that matters here: every increment of delivered payload mass multiplies the required propellant, and for the velocities relevant to interstellar transit on human-relevant timescales the mass ratios become astronomical even under optimistic exhaust velocities. The kinetic energy that must be invested to bring a payload to a relativistic cruise speed scales linearly with that payload's rest mass; doubling the delivered mass at least doubles the energy, and once one accounts for the propellant that energy must also accelerate, the scaling is far worse. Whatever an advanced civilization's energy budget, *the marginal cost of delivered interstellar mass is high and superlinear*. This alone makes mass a resource a rational actor economizes ruthlessly. It does not by itself bound payload to near-zero, but it establishes that payload is expensive in a way that destinations and timescales do not relax.

(ii) The asymmetry between transporting bits and transporting atoms. Information can be transmitted across interstellar distances electromagnetically at the speed of light, with energy costs that, while nontrivial, are many orders of magnitude below the cost of accelerating the corresponding mass of matter to a usable fraction of c. A genome is roughly 10^9 base pairs, on the order of 10^9 bytes; the energy to *radiate* that information to a receiver is negligible beside the energy to *ship* even a single gram of the corresponding biological material across light-years. This asymmetry is not an engineering contingency that a sufficiently advanced civilization escapes; it is thermodynamic and relativistic. Bits are cheap to move and atoms are not, and the ratio of the two costs grows with distance. A civilization that has solved the problem of *reaching* a star at all has, by construction, a communication channel to it that is overwhelmingly cheaper per unit of conveyed specification than its transport channel. The rational implication is immediate: *convey specifications, not substance, wherever substance can be reconstructed at the destination.*

The weak form follows from (i) and (ii) jointly and requires no exotic physics, no contested phenomenology, and no specific propulsion mechanism. It is, in fact, the same logic that Crick and Orgel invoked in proposing directed panspermia: faced with the implausibility of passively shipping organisms across interstellar distances, they reasoned that deliberate seeding with minimal biological payload was the more practical channel (Crick & Orgel, 1973). The ABH generalizes that move from the origin-of-life setting to the colonization setting: it replaces "ship organisms" with "ship specifications and wait for local fabrication," and it makes the waiting—the dependence on the destination's own development—explicit.

3.2 The strong form: metric-engineering constraints

The strong form is not required for the argument, but it is worth establishing that it is not idle. If one entertains the possibility of a transit mechanism beyond reaction propulsion—specifically the metric-engineering or "warp" class introduced by Alcubierre (Alcubierre, 1994)—then the "bell" intuition turns out to have a concrete realization in the structure of those solutions, rather than being a free-floating analogy.

The Alcubierre solution achieves effective superluminal transit by contracting spacetime ahead of a bubble and expanding it behind, leaving the bubble's contents in locally flat, inertial space. The solution's notorious difficulty is its energy requirement: in its original form it demands a quantity of negative energy density that scales with the size of the bubble, with relativistic bubbles of order one meter requiring negative energy comparable in magnitude to the mass-energy of a star (Alcubierre, 1994), and quantum-inequality analyses tightening rather than relaxing these constraints (Pfenning & Ford, 1997; Lobo & Visser, 2004). The first relevant point is therefore direct: in this class of solutions, *the cost scales with the volume one wishes to transport*. Payload volume is not free; it is the dominant term in an already prohibitive budget. A civilization operating such a mechanism would be under extreme pressure to minimize the transported volume—which is exactly the bootstrap incentive, arrived at from the propulsion physics rather than from the rocket equation.

The second point is more striking, and we flag it as suggestive rather than evidential. Van Den Broeck showed that the energy requirement can be reduced by many orders of magnitude using a bubble whose *exterior* is microscopic while its *interior* volume is large—a geometry in which a tiny outer "neck" connects to a spacious internal pocket (Van Den Broeck, 1999). This is a genuine feature of the mathematics, not a metaphor: the least-prohibitive known warp geometries are precisely those that are small on the outside and large on the inside, and that pass their contents through a constricted throat. Two consequences follow for our purposes. First, even in the energy-reduced regime the constriction is real: the throat that makes the geometry affordable also limits what can pass through it, so payload remains constrained by the mechanism, not merely by the energy account—this is the strong-form "bell" claim, now with a referent. Second, and we note this with appropriate caution, the qualitative signature such a geometry would present to an occupant—an interior larger and stranger than the exterior implies—coincides with a recurrent and otherwise puzzling motif in the anomalous-craft reporting literature. We do not advance this coincidence as evidence; the reports are of contested provenance and we conditionalize on nothing here. We note only that the strong form of the premise is not ad hoc: it corresponds to an actual property of the least-implausible known solutions for faster-than-light transit, which is more than can be said for capability monotonicity, the assumption it replaces.

3.3 What the argument actually needs

The decisive observation is that the bootstrap argument is driven entirely by the *weak* form. The argument requires only that, for a payload-limited visitor, the cost of transporting bulk matter and infrastructure exceeds the cost of (a) transporting information plus (b) waiting for the destination to develop local instantiation capacity. The weak form—rocket equation plus bit/atom asymmetry—establishes exactly this inequality, and establishes it from physics no advanced civilization can repeal. The strong form adds color and a possible mechanism, and it dissolves the objection that the premise is mere analogy, but it is logically optional. This matters for the model's robustness: a reader may reject metric-engineering propulsion entirely, regard the "bell" as fanciful, and discard Section 3.2 in full, and the bootstrap argument survives intact on the strength of Section 3.1 alone.

We therefore recommend, in answer to the question of how the premise should be held, that it be owned in the following layered way: the *weak* payload-limitation premise is asserted as a near-theorem, defended from the rocket equation and the bit/atom asymmetry; the *strong* "bell" premise is offered as a defensible and concretely instantiated possibility that strengthens the picture but bears no logical weight. This is stronger than "flat assumption defended by analogy," and it is honest about which parts are load-bearing.

4 The Bootstrap Argument

We now state the argument with its premises and derived consequences separated, so that the reader can see exactly where assumption ends and inference begins.

Premises.

(P1) Existence and access.
At least one non-human intelligence exists and possesses interstellar transit capability sufficient to place a small, low-payload vehicle at a target star.

(P2) Payload limitation (weak form).
Transporting bulk matter and infrastructure across interstellar distances is enormously more costly than transporting information, sufficiently so that bulk transport is not the rational means of establishing large-scale presence (Section 3.1).

(P3) Instrumental rationality.
The visitor selects strategies that achieve its objectives at acceptable cost; it does not pay enormously more than necessary to achieve a given end.

(P4) Non-destabilization preference.
To the extent the visitor's objectives depend on the destination biosphere or its developmental trajectory, it prefers interventions that do not collapse that trajectory prematurely.

Derived consequences. From (P1)–(P3): the visitor prefers to convey the minimal seed—information, and at most minimal biological or material samples—rather than to import populations, ecosystems, or infrastructure. From (P2)–(P3): large-scale presence is impractical until the destination can convert conveyed information into local embodiment and infrastructure; before that point, the value of large-scale presence is gated not by intent but by feasibility. From (P3)–(P4): interaction during the pre-threshold interval is rationally *minimal and non-maximal*—enough to serve whatever monitoring or steering objective the visitor has, but not enough to incur the cost or risk of destabilizing the developmental trajectory on which the strategy depends.

The bootstrapping point. Define the *bootstrapping point* as the earliest time at which the destination civilization possesses the local capacity to convert conveyed information into embodiment, manufacturing, repair, and reproduction without transported resources. Before it, the rational strategy is *seed-and-wait*: convey specifications, monitor and lightly steer, and incur no large logistics cost. After it, presence becomes achievable using local resources, and the strategic regime can change—from passive monitoring toward active instantiation, integration, or disclosure—*without the transport problem ever having been solved*. The model's central structural claim is that the bootstrapping point is set by the *destination's* technological state, not by the visitor's preferences, and is therefore, in principle, observable from within the destination civilization.

It is worth stressing what this buys and what it does not. It does not predict that any regime change *will* occur, since (P1) may simply be false and no visitor may exist. What it provides is a conditional structure with an observable hinge: the claim that *if* the model holds, the hinge is our own technological maturation, which we can watch. A model whose decisive variable is the alien's inscrutable will (the zoo hypothesis) offers no such hinge; a model whose decisive variable is our own synthetic-biology and autonomy capability does.

5 Candidate Thresholds and the Convergence Claim

If the bootstrapping point is the moment a civilization can "convert arbitrary information into machines, minds, and bodies," then identifying it amounts to identifying the technologies that perform that conversion. We organize the candidates here and then connect the biological route to a companion model that makes the relevant threshold quantitative and, importantly, identifies *which* capability binds.

Capability	Conversion it performs
Artificial intelligence	Converts conveyed specifications and data into executable engineering and design
Robotics / autonomous fabrication	Builds and maintains physical systems without a large imported labor presence
Synthetic biology	Constructs biological systems locally from designed components
Genome-scale DNA synthesis	Turns organisms into transmissible information (a printable specification)
Developmental viability / ectogenesis	Grows a specified genome into a supported, viable organism without an existing compatible host
Advanced materials & high energy density	Enables local fabrication and the power budgets that fabrication, computation, and containment require

Table 1: Candidate threshold technologies and the information→embodiment conversion each performs. The bootstrapping point is the convergence at which conveyed specification suffices for local instantiation; Section 5.2 argues that, for the biological route, the convergence is gated by developmental viability.

5.1 The instantiation stack

Table 1 lists the principal candidate technologies and the conversion each performs. The unifying property is that each turns *information* into *embodiment or infrastructure*: each reduces the amount of substance that must be transported by increasing the amount that can be reconstructed locally from specification. The bootstrapping point is then better understood not as any single invention but as the *convergence* at which conveyed specification suffices for local instantiation.

5.2 The biological threshold and the Kent Equation

Among these candidates the biological column is of particular interest, because it converts a logistics problem into an information problem in the most literal way: a mature biological-fabrication capability would make an organism a specification. The companion to this paper models *when* a civilization acquires that capability (Kent, 2026). Where the Drake equation estimates a static expected count as a product of probabilities, the companion model—the *Kent Equation*—replaces the static product with a time-indexed *hazard* and takes as its headline quantity a *velocity* rather than a count.

The construction can be summarized as follows. Readiness is conjunctive: instantiation requires several necessary capabilities, so the crossing time is the maximum over their individual crossing times, $\tau^* = \max_c \tau_c$—the threshold is set by the *last* necessary capability to mature. Each capability's crossing is modeled through a hazard $h_c(t) = A_c(t)\,M(t)^{\beta_c}$, where $M(t)$ is a *substrate-maturity index* common to all capabilities and $A_c(t)$ is a capability-specific *conversion accelerant*. The substrate index is a capped weighted geometric mean of three measurable components—energy E, general-purpose computation C, and biological experimental throughput B—each normalized to the level x_i^* it must reach for the gate:

$$M(t) = \prod_{i\in\{E,C,B\}} \left(\frac{x_i(t)}{x_i^*}\right)^{w_i}, \qquad \frac{x_i(t)}{x_i^*} \leq 1, \quad w_i > 0. \tag{1}$$

The geometric mean encodes *complementarity*: because all three layers are required, a near-absence of any one drives maturity toward zero, where a weighted sum would wrongly score the civilization as moderately mature. The accelerant $A(t)$—the natural locus of machine intelligence—is two-sided, with $A > 1$ net acceleration, $0 < A < 1$ net drag, and $A = 1$ the brute-force baseline.

The model's headline object is not the crossing probability but its *rate of advance*. Taking the logarithmic derivative of the binding hazard $h = A\,M^{\beta}$, and using $\ln M = \sum_i w_i(\ln x_i - \ln x_i^*)$, the unknown required levels

x_i^* cancel exactly:

$$\frac{\dot{h}}{h} = \frac{\dot{A}}{A} + \beta \sum_{i \in \{E,C,B\}} w_i \frac{\dot{x}_i}{x_i} = r_A + \beta \sum_i w_i r_i, \tag{2}$$

a sum of growth rates of quantities with real, long-run time series. *How fast the threshold is approaching can be measured without knowing how far away it is*; recovering an actual date requires committing to one quantity the rate law dispenses with—the gap between the binding input's current and required levels. The companion paper's position is therefore that the model "does not output a year": it localizes all timing disagreement to a single interpretable number, the gap on the binding input, and exhibits the sensitivity of the date to it (Kent, 2026).

Three of the companion model's results bear directly on the present argument, and we use them in Sections 5.3 and 7.

(K1) The binding gate is developmental viability, not synthesis.
At human (multicellular-vertebrate) scale the last necessary capability to mature is *developmental viability*—growing a specified genome into a supported, viable organism—because the knowledge it requires is obtainable only by running development. The empirical record bears this out: genome writing has reached single-celled-eukaryote scale (Gibson et al., 2010; Annaluru et al., 2014), whereas developmental support reaches only the two ends of gestation (Aguilera-Castejón et al., 2021; Partridge et al., 2017), with continuous instantiation of a specified body absent. Synthesis crosses earlier; development binds.

(K2) Conversion is substrate-gated and back-loaded.
The marginal return on the accelerant is $\partial h/\partial A = M^\beta$, which for a threshold-like capability ($\beta > 1$) is convex and near zero until substrate is mature. Intelligence applied to body-growing buys little while the wet-lab substrate is immature, and a great deal once it has matured.

(K3) Co-movement is nearly uninformative about proximity.
Because the least-mature component caps maturity while a fast-growing component can dominate the *rate*, the percentage rate of approach and the absolute distance to the threshold decouple. The model's empirical content lives in the *timing and ordering* of crossings, not in the shared upward direction of a basket of indices.

Equations (1)–(2) are reproduced only to show the model's form and the three results we rely on; the operational definitions, structural derivations, and sensitivity analysis are developed in the companion work, and nothing in this paper depends on a particular numerical value of τ^*. The matched pair is then easy to state: the ABH is the qualitative claim that a payload-limited visitor would wait for the bootstrapping point; the Kent Equation is the quantitative model of *when the biological bootstrapping point arrives and what gates it*. The ABH's general bootstrapping point is the earliest of whatever instantiation routes are available; where the objective requires biological embodiment, the Kent Equation supplies the relevant threshold and names its binding gate.

5.3 Why neither "autonomy-first" nor "biology-first" is right

A natural objection to the biological framing is that the genuine bootstrapping layer is general-purpose information-to-artifact conversion—artificial intelligence coupled to robotics—because that layer could in principle stand up the rest of the stack, biology included. The companion model adjudicates this dispute more precisely than the bare intuition allows, and the answer is that the dichotomy is false.

In the Kent Equation, machine intelligence is not a separate threshold but the *conversion accelerant* $A(t)$: it converts substrate into solved developmental capability. Its leverage is real but doubly constrained.

First, its return is *substrate-gated*: by (K2), the marginal effect of conversion on the hazard is M^{β}, near zero while the biological substrate is immature, so intelligence cannot move the crossing appreciably until the wet-lab layer matures. Second, the binding gate is *developmental viability* (K1), whose required knowledge is obtainable only by running development; this is a knowledge frontier that computation accelerates but does not substitute for, because the missing information is empirical and must be generated experimentally. The companion model formalizes the trade-off as an elasticity-weighted race between a fast, unit-elasticity conversion lever and a slow, high-elasticity substrate lever, with an explicit crossover (Kent, 2026).

The upshot for the ABH is a sharper threshold than "AI plus robotics" or "synthetic genomics" taken alone. The bootstrapping point, for the biological route, is gated by developmental viability; energy and computation are necessary substrate, and machine intelligence is an accelerant that pulls the date modestly earlier as a tail correction but cannot leapfrog the empirical biology. "Autonomy-first" mistakes the accelerant for the gate; "biology-first," read as "genome synthesis," mistakes an early-crossing prerequisite for the binding one. The correct reading is that *the binding gate is the growing of a specified body*, and it is the maturation of that specific capability—not of synthesis, nor of intelligence in the abstract—whose timing the ABH should track. We make this the backbone of the model's empirical content in Section 7.

6 Strategic Interaction Under the Model

The model implies a characteristic interaction regime during the pre-threshold interval, derivable from (P3) and (P4). Because the strategy depends on the destination's continued development, and because large-scale presence is in any case impractical before the bootstrapping point, the rational interaction profile is one of *calibration rather than conquest*: enough engagement to monitor or lightly steer the trajectory, not enough to collapse autonomy, induce panic, or transfer operational knowledge that would perturb the developmental path on which the strategy relies.

Concretely, the model favors interaction that is brief rather than sustained, ambiguous rather than dispositive, and selective rather than broad—plausibly concentrated on the institutions (scientific, aerospace, military, high-energy) whose trajectories most affect the threshold. The contrast the model draws is with the invasion archetype: a payload-limited visitor cannot mount an invasion and has no reason to telegraph capabilities it does not possess, so the predicted phenomenology—*if any phenomenon exists at all*—is one of intermittent, low-information presence rather than overwhelming or instructive contact. We stress the conditional. The point of this section is not that such a phenomenology is observed; it is that, *given the premises*, this is the interaction profile the premises rationally entail, which is a claim about the model's internal logic and not about the world.

7 Empirical Content: Consequences, Signatures, and Demarcation

A speculative model earns scientific interest only to the extent that it makes the world look different than it would under the alternatives, and that it risks something on the difference. This section is the most important in the paper, because the ABH has a characteristic and dangerous failure mode: it is *too accommodating*. Almost any observation about secrecy, biotechnology politics, or anomalous reports can be retrofitted to it, and a model that can absorb any datum predicts nothing. Our task is therefore not to suppress the model's many consequences but to *enumerate them in full and grade each by evidential weight*, so that the small set of genuinely risky predictions is not lost among the large set of observations that are merely consistent with the model.

7.1 Three layers, and why the grading is necessary

We keep three layers separate, following the structure of the argument in Section 4: *assumptions* (P1–P4); *derived consequences* (seed-and-wait, deferred presence, calibrating interaction); and *signatures*—observations that the world would or would not present if the model held. The model gets stronger the more its content can be moved from the second layer to the third in a form that risks disconfirmation. The catalog below is the third layer in full. Each entry is a genuine consequence of the model; the question for each is not whether it follows but whether its presence in the world would *discriminate* the model from its mundane alternatives.

7.2 The Kent rate law as the grading instrument

The single most useful methodological tool is the companion model's decoupling result (K3). In an era of broad technological growth, nearly every index rises together, so correlation among technology indicators is nearly free and carries almost no information about proximity to the threshold. The content lives in *timing and ordering*, not in co-movement. This converts the informal injunction "do not count consistency as evidence" into a structural claim with teeth: the ABH's testable content is concentrated in (a) the *ordering* of crossings—which capability inflects before which—and (b) the *shape* of the binding indicator's trajectory, which the conjunctive structure predicts should be late-inflecting and S-shaped rather than a linear extrapolation of current trends. A signature that merely records that "many relevant technologies are advancing" is, by (K3), close to content-free; a signature that specifies an ordering or an inflection is not. We grade the catalog accordingly.

7.3 The catalog of derived signatures

Table 2 lists the consequences the ABH generates, the expected signs of each, an evidential tier, and the dominant mundane explanation that the model must outcompete. We use three tiers. **Tier A** signatures are genuinely risky: they specify an ordering, a trajectory shape, or a covariation that could fail, and their failure would count against the model. **Tier B** signatures are consistency-only: they are necessary for the model's viability but are at least as well explained by ordinary causes, so their presence is weak confirmation at best. **Tier C** signatures presuppose contested data (chiefly the veridicality of abduction and contact reports, or the detectability of engineering in human biology) and carry essentially no evidential weight; several ride on the model's weakest premises.

7.4 The genuinely risky predictions (Tier A)

Three predictions in the catalog risk something, and they divide into two kinds that it is important not to conflate.

The first kind tests the *companion timing model* and requires no aliens at all to be checked. The claim that developmental viability is the binding gate, crossing *after* genome synthesis, is an ordering claim about technological development that is checkable against the public record now (Gibson et al., 2010; Annaluru et al., 2014; Aguilera-Castejón et al., 2021; Partridge et al., 2017). The claim that developmental-viability indicators should trace a late-inflecting S-curve, rather than the linear extrapolation a naive reading of current culture durations would suggest, is a claim about trajectory shape that the coming decade will sharpen. These predictions have empirical content independent of the ABH's alien premise: they test whether the Kent model correctly identifies *which* capability binds and *how* its approach is shaped. They are interesting on their own terms, and if they fail, the companion model—and with it the ABH's account of the bootstrapping point—is weakened regardless of any claim about visitors.

The second kind is the ABH's own distinctive prediction, and it is the only one in the catalog that genuinely requires the alien premise:

> *The intensity of pressure toward disclosure, regulation, and institutional attention to non-human intelligence should co-vary specifically with the maturation of the binding biological gate—developmental viability—rather than being independent of it, tracking unrelated drivers, or co-varying merely with the broad basket of advancing technologies.*

Signature class	**What the ABH predicts (expected signs)**	**Tier**	**Dominant mundane null**
Technology threshold	Importance concentrates on information→embodiment capabilities; developmental viability binds and crosses *after* genome synthesis	A/B	Technologies gain importance for ordinary reasons
Crossing trajectory	Developmental-viability indicators trace a late-inflecting, S-shaped curve; sign and curvature of $\dot{h}/h$ readable from public substrate series	A	Ordinary sigmoid technology adoption
Disclosure timing	Institutional disclosure and attention pressure rise ahead of the threshold and track the *binding biological gate*; signal is fragmented, contested, institutionally awkward	A (covariation) / B (fragmentation)	Common-cause technological acceleration; politics of any sensitive topic
Strategic interaction	Contact, if any, is brief, ambiguous, selective (scientific, aerospace, military, high-energy foci); calibration, not invasion	B	Conditional on a phenomenon existing at all
Payload limitation	Craft, if real, are small, few-occupant, low-cargo, transit-optimized, with interiors larger or stranger than exteriors imply	B	Conditional on craft being real
Genetic policy	Intensifying conflict over genome synthesis, embryo editing, ectogenesis; biosecurity framing of *synthesis capability* specifically	B	Real medical, economic, and biosecurity stakes
Private-sector migration	Strategic biotech / AI / robotics capability concentrates in private firms; defense interest follows the capability	B	Documented feature of the contemporary economy
Infrastructure convergence	Energy, compute, synthetic biology, synthesis, and robotics mature together (the substrate index M rises)	B (convergence) / A (internal ordering)	Broad technological growth lifts all indices
Biological sampling	Reports, *if veridical*, emphasize reproduction, immunity, neural, endocrine, and microbiome systems over reconnaissance	C	Presupposes abduction reports; known folklore mechanisms
Human compatibility	Genetic features that resist standard evolutionary explanation (cognition, language, neoteny, sociality)	C	Selection, drift, introgression already explain these
Information seeding	Contact reports, *if veridical*, feature cognitive "downloads," symbols, and anticipated concepts rather than blueprints	C	Cryptomnesia, zeitgeist, cultural diffusion
Historical intervention	Mythologized sky-beings, watchers, forbidden knowledge, and "gifted" technological leaps	C	Universal features of human cognition and culture
Failure-mode noise	Conflicting factional pressures yield a noisy, contradictory public signal	C ($\approx$ content-free)	Noise is consistent with the hypothesis *and* its negation

Table 2: The full catalog of consequences the ABH generates, graded by evidential weight. Tier A signatures risk disconfirmation; Tier B signatures are necessary but have strong mundane explanations; Tier C signatures presuppose contested data and carry essentially no weight.

The specificity is what (K3) buys. A naive version of this prediction—"disclosure pressure rises as technology advances"—is nearly content-free, because in an era of broad growth almost everything rises together. The sharpened version ties disclosure pressure to the *binding* indicator in particular, and to its ordering relative to the early-crossing prerequisites, which is a far narrower and more falsifiable claim. It can fail in identifiable ways: if disclosure pressure is uncorrelated with developmental-viability milestones, or tracks genome-synthesis or AI milestones rather than the binding gate, or is fully accounted for by ordinary drivers (military-technology cycles, media dynamics, domestic politics) with no residual tied to the binding gate, the prediction is disconfirmed. The honest difficulties remain real: "disclosure pressure" is hard to operationalize, the time series are short, and confounders are abundant. And the obvious deflationary reading—that *both* variables are driven by common-cause technological acceleration with no visitor required—remains, in our view, the correct null hypothesis. The prediction is interesting not because it is likely to be vindicated, but because it is the rare claim in this domain that specifies in advance what would count against it, and because (K3) lets us state it in a form sharp enough to be worth testing.

7.5 Why the Tier B signatures are weak, and the Tier C signatures weaker

The Tier B signatures—the calibrating interaction profile, the transit-optimized craft features, the intensification of genome-synthesis politics, the migration of capability into private firms, and the broad infrastructure convergence—share an epistemic character that must be named plainly: each is an observation the model can *accommodate* but does not *uniquely predict*, because each is at least as well explained by ordinary causes. Fragmented, contested disclosure is what any politically sensitive subject produces under genuine uncertainty, with or without a referent. Intensifying conflict over genome synthesis is fully explained by the real medical, economic, and biosecurity stakes of that technology. The concentration of strategic AI and biotechnology capability in private firms is a documented feature of the contemporary economy. And the brute fact of infrastructure convergence is, by (K3), close to uninformative about proximity. To treat any of these as confirmation is to commit the model's signature error of counting consistency as evidence. They are, at most, the observations that would have to be *present* for the model to remain viable; their presence is necessary, not sufficient, and given how weakly each constrains the world, "necessary" is a low bar.

The Tier C signatures are weaker still, and two of them deserve specific deflation because they are the ones most likely to tempt over-reach.

The first is the search for *engineered signatures in human biology*—genetic features that look designed rather than selected, or developmental traits "too useful" for technological civilization to be accidental. This line has an instructive precedent. Crick and Orgel, in proposing directed panspermia, suggested that an anomalous biochemical signature—an unexpected dependence on a cosmically rare element such as molybdenum—might in principle betray an external origin of terrestrial life (Crick & Orgel, 1973). The proposal was serious, made by a co-discoverer of the structure of DNA, and it went nowhere: no such signature has been established, and ordinary chemistry and selection account for the observed dependencies. The lesson generalizes directly to the human-compatibility claim. Evolution by natural selection, drift, introgression, and known demographic events already explains an enormous amount of human biology, including the traits one might be tempted to call "engineered"—cooperation, language, neoteny, encephalization. The bar for inferring intervention is correspondingly high: one would need genetic patterns that resist explanation by the standard evolutionary processes, not merely traits that are useful in hindsight. We regard the human-compatibility claim as the weakest element of the entire model and recommend that it be held as a speculative possibility carrying no evidential weight, precisely because the null explanation is so powerful.

The second is the *abduction and contact-report* literature, which the biological-sampling, information-seeding, and payload-limitation "interior" signatures silently presuppose. We have conditionalized on none of it. The reports are of contested and largely unverifiable provenance, are known to be generated by well-documented psychological and cultural mechanisms independent of any external cause, and cannot bear the weight of an inference to non-human intelligence. Where the model's internal logic happens to rationalize a motif in this literature—for instance, that microbiome or reproductive sampling would be a more sensible objective than the cruder folklore suggests, or that a "larger inside" would follow from a Van Den Broeck geometry—this is at most a remark about what the model would imply *if* the premise held; it is not, and must not be presented as, support for the premise.

Finally, the failure-mode signature—that the public signal should look noisy and contradictory because multiple factions contest its interpretation—is the limiting case of overexplanation and is essentially content-free, because a noisy signal is exactly what one expects whether or not anything is happening. We include it in the catalog for completeness and assign it no weight.

8 Relation to Existing Resolutions

The ABH is best understood as a member of the logistics family that incorporates elements of the motivation family without reducing to it.

Its nearest relatives in spirit are the expansion-is-hard models. The sustainability solution holds that fast expansion is self-terminating, so that no persistent observable empire is expected (Haqq-Misra & Baum, 2009); diffusion and percolation models hold that expansion is slow, patchy, and apt to halt short of saturation (Newman & Sagan, 1981; Landis, 1998). The ABH shares the conviction that the step from "capable of reaching" to "capable of saturating" is where the standard argument breaks, but it locates the obstacle more specifically—in a structural gap between transit and bulk delivery—and it draws out the strategy (seed-and-wait) that the gap rationally implies. It is, in effect, a micro-foundation for one route by which expansion can be slow: not because expanders collapse, but because the rational expander defers physical presence until the destination can supply the infrastructure transport cannot.

Its relationship to the motivation family is more subtle. Like the zoo hypothesis, it predicts limited, non-disruptive contact (Ball, 1973); like the planetarium hypothesis, it is compatible with managed perception (Baxter, 2001). But it derives restraint from constraint rather than from choice, and this difference is not cosmetic. The motivation models place the decisive variable beyond our observation—the alien's preferences, our "maturity" as judged by the alien. The ABH places the decisive variable within our observation—our own technological state, and specifically the maturation of the binding biological gate identified by the companion model (Section 5.2)—and thereby exposes itself to disconfirmation in a way the motivation models do not. That exposure is the model's principal scientific virtue, and it is purchased precisely by the logistics framing.

Finally, the model inherits from directed panspermia (Crick & Orgel, 1973) the core insight that information is the rational interstellar payload, and from the Tipler–Sagan debate over self-reproducing probes (Tipler, 1980) the recognition that what a civilization can *send* ahead of itself is the crux. It differs from the probe literature in its emphasis: where the probe debate asks what an autonomous artifact would do upon arrival, the ABH asks what a payload-limited principal would do *given* that the destination is not yet able to receive it—and answers that it would wait, and watch, and time its presence to the destination's own development.

9 Limitations

We close by stating the model's weaknesses directly, since the discipline of the paper has been to keep its tiers honest.

The model is conditional on an unverified premise. Everything downstream assumes (P1)—that a non-human intelligence with interstellar access exists. We have offered no evidence for (P1) and none follows from the model; the paper analyzes a conditional and nothing more.

The empirical content is thin and confounded. The ABH's own distinctive prediction (Section 7) involves variables that are hard to operationalize, short time series, abundant confounders, and a powerful deflationary null—common-cause technological acceleration—which we regard as the favorite. The Tier A predictions inherited from the companion model are cleaner, but they test the timing model rather than the alien hypothesis; the remaining signatures carry, by our own argument, little to no evidential weight.

The binding-gate identification is inherited, and inherits its caveats. The claim that developmental viability is the binding gate, and the timing structure built on it, come from the companion model (Kent, 2026) and share its limitations: the gap on the binding input is a load-bearing scalar judgment that is genuinely multi-dimensional; the responsiveness exponent is an unknown constant; and co-trending is mitigated by the ordering and shape claims but not eliminated.

The strongest downstream claims are the weakest links. Engineered human compatibility and the veridicality of abduction reports are the elements most likely to attract attention and least able to bear it. The null evolutionary and psychological explanations are strong, and we have recommended that both be held as speculation carrying no weight.

The model risks overexplanation. Its native failure mode is to absorb every datum and predict nothing. The only defense is the one we have tried to practice: enumerate the consequences in full, grade each by evidential weight, specify the risky predictions with their disconfirming conditions, and refuse to count mere consistency as confirmation.

What survives these concessions is, we think, the part worth keeping. Stripped of its embellishments, the ABH is the observation that *access is not logistics*, that a payload-limited visitor would rationally substitute information for substance and defer presence until a destination could instantiate that information locally, and that this substitution—unlike the inscrutable preferences invoked by motivation models—hinges on a variable we can watch. The companion model sharpens that variable to a single capability: not synthesis, not machine intelligence in the abstract, but the growing of a specified body. As a claim about the world, the hypothesis is unproven and probably unprovable on present evidence. As a lens, it does something more modest and more defensible: it reorients the analysis of extraterrestrial presence from the question of arrival, which is intractable, to the question of logistics, which is not—and in doing so it identifies, among all human technologies, the one whose maturation would matter most if payload-limited visitors existed.

References

Aguilera-Castejón, A., et al. (2021). Ex utero mouse embryogenesis from pre-gastrulation to late organogenesis. *Nature*, 593, 119–124.

Alcubierre, M. (1994). The warp drive: hyper-fast travel within general relativity. *Classical and Quantum Gravity*, 11(5), L73–L77.

Annaluru, N., et al. (2014). Total synthesis of a functional designer eukaryotic chromosome. *Science*, 344, 55–58.

Ball, J. A. (1973). The zoo hypothesis. *Icarus*, 19(3), 347–349.

Baxter, S. (2001). The planetarium hypothesis: a resolution of the Fermi paradox. *Journal of the British Interplanetary Society*, 54, 210–216.

Crick, F. H. C., & Orgel, L. E. (1973). Directed panspermia. *Icarus*, 19(3), 341–346.

Forgan, D. H. (2011). Spatio-temporal constraints on the zoo hypothesis, and the breakdown of total hegemony. *International Journal of Astrobiology*, 10(4), 341–347.

Gibson, D. G., et al. (2010). Creation of a bacterial cell controlled by a chemically synthesized genome. *Science*, 329, 52–56.

Hair, T. W. (2011). Temporal dispersion of the emergence of intelligence: an inter-arrival time analysis. *International Journal of Astrobiology*, 10(2), 131–135.

Haqq-Misra, J. D., & Baum, S. D. (2009). The sustainability solution to the Fermi paradox. *Journal of the British Interplanetary Society*, 62, 47–51.

Hart, M. H. (1975). An explanation for the absence of extraterrestrials on Earth. *Quarterly Journal of the Royal Astronomical Society*, 16, 128–135.

Kent, G. (2026). *The Kent Equation: A Rate Law for the Onset of Local Biological Self-Instantiation*. Companion preprint, Rhomega Corp.

Landis, G. A. (1998). The Fermi paradox: an approach based on percolation theory. *Journal of the British Interplanetary Society*, 51, 163–166.

Lobo, F. S. N., & Visser, M. (2004). Fundamental limitations on 'warp drive' spacetimes. *Classical and Quantum Gravity*, 21(24), 5871–5892.

Newman, W. I., & Sagan, C. (1981). Galactic civilizations: population dynamics and interstellar diffusion. *Icarus*, 46(3), 293–327.

Partridge, E. A., et al. (2017). An extra-uterine system to physiologically support the extreme premature lamb. *Nature Communications*, 8, 15112.

Pfenning, M. J., & Ford, L. H. (1997). The unphysical nature of 'warp drive'. *Classical and Quantum Gravity*, 14(7), 1743–1751.

Sandberg, A., Drexler, E., & Ord, T. (2018). Dissolving the Fermi paradox. arXiv:1806.02404.

Tipler, F. J. (1980). Extraterrestrial intelligent beings do not exist. *Quarterly Journal of the Royal Astronomical Society*, 21, 267–281.

Van Den Broeck, C. (1999). A 'warp drive' with more reasonable total energy requirements. *Classical and Quantum Gravity*, 16(12), 3973–3979.