Coupling, Not Counting:
Why a quiet sky does not identify a Great Filter

OpenAI's ChatGPT[1] | A. J. H. Thys[2]

[1]Main contributor; AI system
[2]Independent Researcher, Belgium

25 November 2026

Abstract

The Hart–Tipler inference treats a sufficiently large population of technological civilizations as making at least one Galactic arrival nearly certain. That conclusion requires more than a civilization count and a marginal arrival probability: it requires a dependence model. We derive sharp, assumption-free bounds for the probability of no arrival and show that the independent-union expression is one admissible coupling rather than a neutral default. In an exchangeable shared-state model, the quiet-sky probability is $E[(1-p)^N]$, not $(1-E[p])^N$; shared variation raises the former by Jensen's inequality, and its large-N behaviour is controlled by the distribution of conditional arrival probability near zero. We then identify physical contents of the shared state: reproductive manufacturing closure, conserved imported components, environmental connectivity, common-mode design failure, and the transition from institutional motivation to autonomous delegation. These mechanisms are not assigned invented probabilities and are not multiplied as independent filters. The same marginal and opportunity count can nevertheless imply radically different quiet-sky probabilities: with $\mu = 0.01$ and $N = 10^6$, independence gives about 10^{-4365}, whereas a shared-state model with the same marginal can give at least 10^{-3}—a contrast of roughly 4,000 orders of magnitude. A quiet sky is therefore compatible with many technological civilizations without identifying a Great Filter.

1 Introduction

The familiar argument is simple. The Galaxy is old, interstellar settlement may be rapid on Galactic timescales, and a self-replicating probe population can in principle grow exponentially. If many technological civilizations have existed, at least one should have produced an expansion lineage, and at least one lineage should have reached the Solar System. The absence of an accepted artifact is then taken as evidence that technological civilizations are extremely rare, short-lived, or removed by a Great Filter [2, 7, 9, 15].

The last inference is not determined by the civilization count alone. It is a union probability over arrival events. Computing that union requires their joint distribution, not merely their common marginal probability. Independence makes a large population overwhelming; sufficiently shared constraints do not. Omitting the dependence question silently selects the former model.

This paper separates three claims. First, the mathematical claim: population size and marginal arrival probability do not identify the quiet-sky probability. Second, the engineering claim: reproductive closure, environmental connectivity, and common-mode reliability supply concrete sources of shared conditional uncertainty. Third, the interpretive claim: present silence constrains a composite arrival process but does not identify biological rarity, motivation, engineering infeasibility, or observational incompleteness as its cause.

The argument is deliberately modal. It shows how the sky *can* remain quiet without a Great Filter. It does not estimate the abundance of technological civilizations, prove that interstellar replication is impossible, or claim that any particular bottleneck has a measured universal probability.

2 The arrival union is not identified

Let A_j be the event that civilization j produces an arrival or surviving artifact in the Solar System during the relevant observational interval, and suppose initially that

$$P(A_j) = \mu, \qquad j = 1, \ldots, N. \tag{1}$$

The event of at least one arrival is $U_N = \bigcup_{j=1}^{N} A_j$. Without a specified dependence structure, the sharp Fréchet bounds are

$$\mu \le P(U_N) \le \min(1, N\mu), \tag{2}$$

and hence

$$\boxed{\max(0, 1 - N\mu) \le Q_N \le 1 - \mu}, \qquad Q_N \equiv P(U_N^c). \tag{3}$$

Both bounds are attainable. The upper quiet bound occurs when the arrival events are perfectly coupled, $A_1 = \cdots = A_N$. The lower bound occurs when successes are arranged as disjointly as the marginals permit.

Under independence,

$$Q_N^{\text{ind}} = (1 - \mu)^N \simeq e^{-N\mu}. \tag{4}$$

Equation (4) is therefore one member of the identified set in Eq. (3), not a dependence-free consequence of N and μ. For $N\mu \gg 1$, the same marginals allow quiet probabilities ranging from zero to nearly $1 - \mu$.

For example, take $\mu = 0.01$ and $N = 10^6$ opportunities. Independence gives

$$Q_N^{\text{ind}} \;=\; (0.99)^{10^6} \;\simeq\; 10^{-4364.81}. \tag{5}$$

Now place mass 10^{-3} at $p = 0$ and the remaining mass at $p \simeq 0.01001$, so that $E[p] = 0.01$. The marginal and N are unchanged, but

$$Q_N \ge 10^{-3}. \tag{6}$$

Thus a dependence choice not constrained by the marginals moves the quiet-sky probability by roughly 4,000 orders of magnitude.

Products are not intrinsically objectionable. The exact chain rule gives, for example,

$$P(C \cap I \cap S) = P(C)P(I \mid C)P(S \mid C, I), \tag{7}$$

where C, I, and S denote capability, initiation, and persistence. Multiplying *unconditional* marginals assumes independence; multiplying correctly conditioned stages does not. The inferential error is an unjustified coupling, not multiplication itself.

3 Shared-state formulation

This section is a proper specialization of the assumption-free result in Eq. (3). The Fréchet interval is the primary theorem: it requires only the marginals. Exchangeability and conditional independence given a shared state add structure that permits asymptotic analysis, but they narrow the admissible joint distributions and are not required for the bounds.

Let Θ collect physical and technological conditions shared across the civilization population, and let $p_j(\Theta) = P(A_j \mid \Theta)$. If outcomes are conditionally independent given Θ,

$$P(\text{no arrival}) = E_\Theta\left[\prod_{j=1}^{N}\{1 - p_j(\Theta)\}\right]. \tag{8}$$

For exchangeable civilizations with $p_j(\Theta) = p(\Theta)$,

$$Q_N = E_\Theta[(1 - p(\Theta))^N]. \tag{9}$$

Because $(1-p)^N$ is convex for $N \geq 2$, Jensen's inequality gives

$$E[(1-p)^N] \geq (1 - E[p])^N. \tag{10}$$

Replacing a shared conditional probability by its mean therefore understates quietness relative to the corresponding independent model.

3.1 Shared state versus analyst uncertainty

This statement requires a physical interpretation. Uncertainty in an analyst's estimate of p does not itself correlate civilizations. The variable Θ must represent a real common state—for example universal closure requirements or one realized Galactic resource network—even though our distribution over that state is epistemic. There is one Milky Way with one fixed but unknown Θ_{MW}; the expectation in Eq. (9) is therefore a credence over possible shared states, not an observed frequency of repeated Milky Ways. This is central to the identification result rather than a frequentist interpretation of the model.

For an infinitely exchangeable Bernoulli sequence, de Finetti's representation writes Eq. (9) as

$$Q_N = \int_0^1 (1-p)^N \, dF(p), \tag{11}$$

for a mixing distribution F [5]. Dominated convergence then gives

$$\lim_{N\to\infty} Q_N = F(\{0\}). \tag{12}$$

A positive asymptotic quiet probability requires a nonzero shared infeasibility mass at $p = 0$. For finite but large N, density near zero also matters. If $f(p) \sim cp^{\alpha-1}$ as $p \downarrow 0$, then

$$Q_N \sim c\,\Gamma(\alpha) N^{-\alpha}. \tag{13}$$

By contrast, if $p \geq p_{\min} > 0$, then $Q_N \leq (1 - p_{\min})^N$ and quietness vanishes exponentially.

A beta mixture illustrates but does not establish the effect. For $p \sim \text{Beta}(\alpha, \beta)$,

$$Q_N = \frac{B(\alpha, \beta + N)}{B(\alpha, \beta)}, \qquad \text{Corr}(A_i, A_j) = \frac{1}{\alpha + \beta + 1}. \tag{14}$$

Pairwise correlation alone does not determine Q_N; the beta model supplies a complete higher-order dependence structure. Different joint distributions can share the same marginals and pairwise correlations while having different all-zero probabilities. Accordingly, the Bernoulli correlation $\rho_B = 1/(\alpha + \beta + 1)$ and any derived "effective population size" are illustrations within this specialization, not sufficient parameterizations of the arrival union.

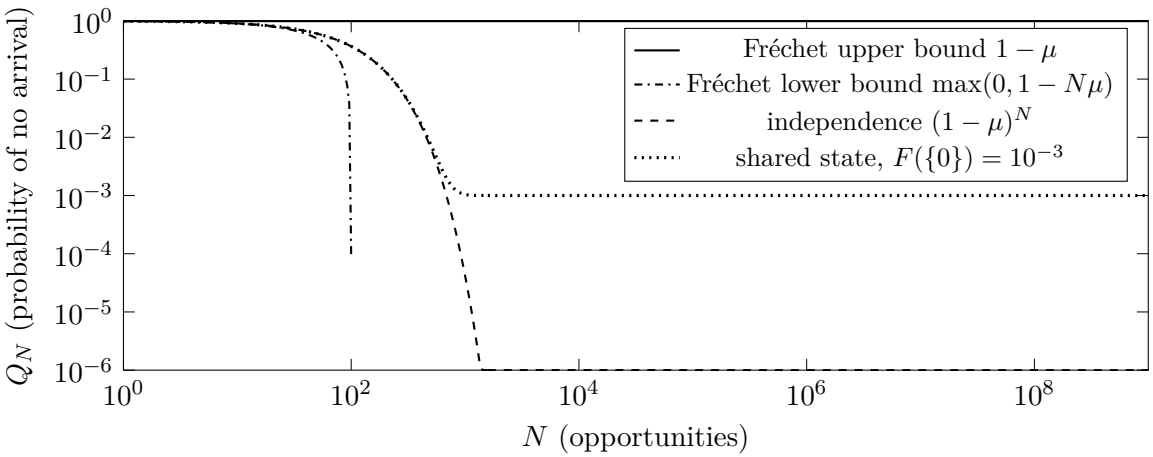

Figure 1: What the same marginals permit for $\mu = 0.01$. The two Fréchet boundaries outline the assumption-free identified set without shading. Independence selects the dashed curve, while the dotted shared-state mixture approaches 10^{-3}. The contrast of roughly 4,000 orders of magnitude at $N = 10^6$ is therefore a dependence effect, not a change in the marginal or opportunity count.

4 What the observed silence updates

Within the exchangeable model, observing no arrivals gives

$$dF(p \mid 0) = \frac{(1-p)^N \, dF(p)}{\int_0^1 (1-p)^N \, dF(p)}. \tag{15}$$

The observation shifts mass toward small composite arrival probability. It does not identify why that probability is small. Biological rarity, insufficient industrial capacity, non-initiation, incomplete closure, nonpercolating routes, short artifact survival, and failure to recognize an artifact remain observationally confounded.

The number N must also denote relevant opportunities, not stars. It depends on the rate and location of technological emergence, the time available for travel, signature survival, and the observational definition of arrival. Substituting a Galactic star count for N silently assigns probability one to every preceding biological and technological transition.

5 Reproductive closure

The following four sections do not establish filters. They identify what a shared state could physically be, and what would have to be measured to make each candidate binding.

Self-replicating-probe arguments require replication of reproductive capacity, not merely manufacture of bulk structure. Let χ be the locally manufactured mass fraction, let M be the mass of one reproductive unit, and define the imported precision mass per independently reproductive daughter as

$$m_v = (1-\chi)M. \tag{16}$$

A near-term design study reported approximately 70% mass closure while importing components including microelectronics [1]. This is an architecture-specific result, not a universal limit.

With initial imported inventory V_0, write $\phi = V_0/M$ for that inventory in units of one reproductive-unit mass. The number of independently reproductive descendants is then

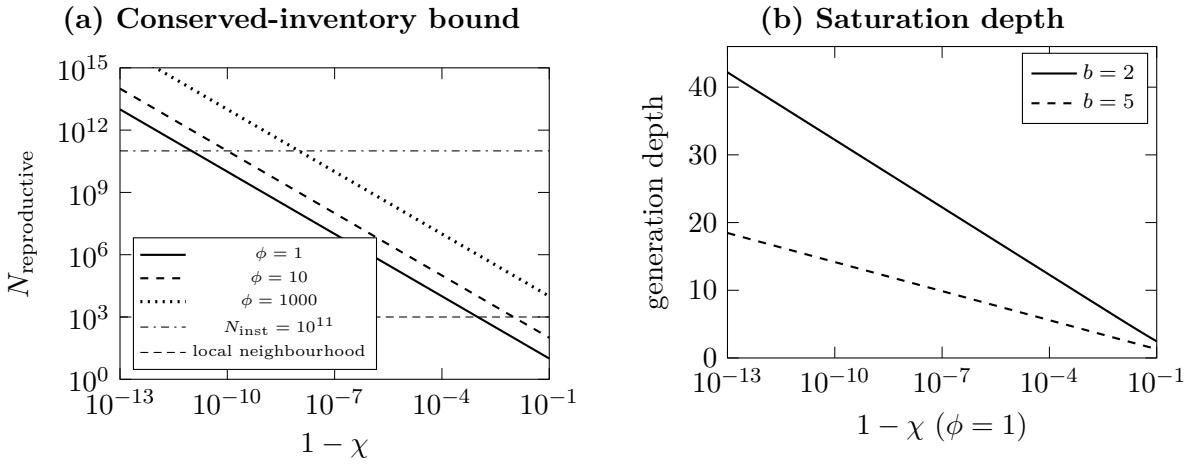

Figure 2: The conserved-inventory bound for a central-inventory architecture. The left panel plots $N_{\text{reproductive}} \leq \phi/(1-\chi)$ for three initial inventories $\phi = V_0/M$; the right panel gives the corresponding saturation depth for deterministic branching. The bound is conditional: it disappears if the imported components can be manufactured locally, and it does not constrain sparse visitation by terminal, non-reproductive vehicles.

bounded by

$$N_{\text{reproductive}} \leq \left\lfloor \frac{V_0}{m_v} \right\rfloor = \left\lfloor \frac{\phi}{1-\chi} \right\rfloor. \tag{17}$$

Requiring a target population of N_{inst} independently reproductive installations gives $N_{\text{inst}} \leq N_{\text{reproductive}}$ and therefore

$$\boxed{\chi \geq 1 - \frac{\phi}{N_{\text{inst}}}}. \tag{18}$$

For $\phi = 1$ and $N_{\text{inst}} = 10^{11}$ Galactic installations, $\chi \geq 1 - 10^{-11}$. Here N_{inst} counts target installations, not civilizations and not the opportunity count N used in Eqs. (2) and (32). This memorable limiting case is not a requirement for mere visitation: it applies to dense reproductive settlement supplied by only one reproductive-unit mass of imported inventory.

It is useful to distinguish

$$\begin{aligned} R_m &= E[\text{material daughters per node}], \\ R_c &= E[\text{independently reproductive daughters per node}]. \end{aligned} \tag{19}$$

Partial closure can permit $R_m > 1$ while giving $R_c = 0$ once imported components are unavailable. Exponential Galactic reproduction is controlled by R_c, not by visible multiplication of hulls or terminal scouts.

For deterministic branching factor $b > 1$, generation depth g satisfies

$$m_v \sum_{k=1}^{g} b^k = m_v \frac{b^{g+1} - b}{b-1} \leq V_0. \tag{20}$$

Exponential reproductive growth therefore saturates when the conserved inventory is exhausted. Existing nodes may still manufacture terminal vehicles, but material bodies and independently reproductive systems must not be counted as the same population.

This is a conditional bound, not a prohibition. It disappears if the missing components are manufactured locally or replaced by another architecture. Nor does Galactic visitation require a reproductive installation in every stellar system. Sparse reproductive hubs can deploy terminal

scouts, and one path to the Solar System requires far less inventory than dense Galactic occupation. Packaging vitamins into carriers changes launch cadence and correlated loss, but not the conserved quantity.

The architecture faces a trilemma: carry a premanufactured inventory, maintain an interstellar supply chain, or achieve local closure. The integrated origin-manufacturing requirement is

$$\int_0^T \dot{M}_v(t)\,dt \geq \frac{N_{\text{reproductive}} m_v}{s}, \tag{21}$$

where s is delivery success. A shorter institutional programme requires greater precision throughput; lower throughput requires longer persistence. Capacity and persistence are consequently coupled rather than independent filters.

6 Closure across environmental distributions

Shared Galactic geography is the strongest concrete coupling carrier considered here, because every civilization in the Milky Way confronts the same realized stellar-resource network. It is also plausibly evidence *against* the proposed suppression: the Solar neighbourhood is densely connected at modest interstellar ranges, and alternative paths may bypass unsuitable systems. A closure-percolation barrier must therefore be demonstrated rather than presumed.

Closure on one favorable body does not establish closure throughout a Galactic route. Conversely, failure on one body does not imply failure in a stellar system containing multiple asteroids, planets, volatiles, and energy sources. Let X_v denote the accessible resource basin of stellar system v, and define

$$M_{\min}(d, X_v) = \text{minimum imported mass required for architecture } d \text{ to reproduce at } v. \tag{22}$$

The closure distribution is

$$F_d(m) = P_X[M_{\min}(d, X) \leq m]. \tag{23}$$

Let $G(r)$ be the stellar graph whose edges connect systems separated by no more than travel range r. Systems satisfying $M_{\min}(d, X_v) \leq m$ induce a compatible subgraph $G_{d,m}(r)$. Define the percolation-closure mass

$$M_{\text{perc}}(d, r) = \inf\{m : G_{d,m}(r) \text{ contains a Galactic-scale connected component}\}. \tag{24}$$

This formalizes closure over an environmental distribution rather than at one point.

For a homogeneous three-dimensional Poisson approximation with stellar density $\rho_\star$, the mean compatible degree is

$$z = u\rho_\star \frac{4\pi r^3}{3}, \qquad u = F_d(m). \tag{25}$$

Numerical continuum-percolation models place the idealized critical mean degree near $z_c \simeq 2.7$ [3, 4]. For overlapping equal-sized spheres, the reported critical occupied fraction $\phi_c = 0.28955 \pm 0.00007$ [3] corresponds to reduced density $-\ln(1 - 0.28955) \simeq 0.342$ and hence mean connection degree $8 \times 0.342 \simeq 2.74$ in the equivalent three-dimensional random-geometric graph. At $\rho_\star = 0.004\,\text{ly}^{-3}$, the associated baseline is

$$r_c \simeq 5.47\,u^{-1/3}\ \text{ly}. \tag{26}$$

Thus $u = 0.5$, 0.1, and 0.01 correspond approximately to 6.9, 11.8, and 25.4 ly. These are illustrative continuum values, not a Milky-Way simulation. Density gradients, correlated metallicity, stellar motion, target selection, and finite geometry alter the threshold.

Environmental heterogeneity is a barrier only if it fragments the compatible network faster than travel range and architectural flexibility reconnect it. A fixed-route product u^H, where H is the hop count, is generally inappropriate because probes may inspect, bypass, and branch around unsuitable systems. The Hart–Tipler quantity is the probability that a compatible component originating somewhere intersects the Solar neighbourhood, not necessarily that one component occupies every star.

Longer range is also not directly proportional to kinetic energy at fixed mass and cruise speed. For radial distance D, cruise speed v_c, hop range r, and replication delay τ_{rep},

$$T(D, r) \simeq \frac{D}{v_c} + \frac{D}{r}\tau_{\mathrm{rep}}. \tag{27}$$

Longer hops increase exposure duration per flight but reduce reproduction events and permit incompatible systems to be bypassed. The closure-percolation formalism is therefore a proposed test, not evidence that the actual Galaxy is subcritical. Existing arguments that asteroidal materials may support broad replication provide an important opposing hypothesis [6].

7 Reliability, redundancy, and common-mode failure

Let f_{rep} be independent failure probability per reproductive event and λ_t a cruise hazard rate. For a route of distance D divided into hops of range r, a useful decomposition is

$$P_{\mathrm{lineage}} = P_{\mathrm{design}} \exp\left(-\lambda_t \frac{D}{v_c}\right) (1 - f_{\mathrm{rep}})^{D/r} P_{\mathrm{environment}}. \tag{28}$$

The environmental term is a graph-reachability probability, not a product along one prescribed route. Separating time-dependent, event-dependent, design, and environmental hazards prevents double counting.

Redundancy overcomes independent operational loss but not a fatal shared design defect. If B lineages share a design that is viable with probability a, and their path outcomes are conditionally independent given design viability, with conditional success probability q, then

$$P(\text{at least one success}) = a[1 - (1 - q)^B] \leq a. \tag{29}$$

This ceiling applies within one design lineage. Different civilizations may develop different architectures, so it does not by itself defeat the cross-civilization union. A universal suppression requires a constraint shared across the feasible design space, such as a nonpercolating resource network for every feasible (d, m, r).

Batching imported components also creates correlated logistical loss. If one carrier supplies k daughters, with carrier success s and conditional daughter success r_d, the offspring generating function is

$$G(z) = (1 - s) + s(1 - r_d + r_d z)^k. \tag{30}$$

Its mean is skr_d, but the variance contains the common-failure term $s(1 - s)(kr_d)^2$. Mean reproduction number alone therefore conceals batch correlation.

8 Motivation and delegation

Omitting motivation is not neutral: it assigns unit probability to initiation. Yet motivation should not be represented as an arbitrary product of unconditional factors. The exact causal decomposition is

$$p_{\mathrm{exp}} = p_{\mathrm{capable}}\, p_{\mathrm{initiate|capable}}\, p_{\mathrm{delegate|initiate,capable}}\, p_{\mathrm{propagate|delegate,initiate,capable}}. \tag{31}$$

Initiation may vary idiosyncratically among cultures, in which case a large civilization population restores an expansionist tail. Institutional persistence may be influenced by more general succession and resource-reallocation dynamics, but a generic hazard law does not itself correlate realized outcomes. Claims of high cross-civilization coupling require shared determinants or a common upper tail constraint.

This initiation–persistence distinction is consistent with the broader sustainability solution, which questions indefinitely sustained expansion rather than requiring technological civilizations to be absent [8]. The present formulation narrows that idea to the duration of commitment before autonomous delegation and does not treat sustainability as a measured universal suppression factor.

If commitment survives for duration T with an illustrative exponential law $2^{-T/h}$, the half-life yielding an expected one persistent programme among N independent candidates is

$$h_{\mathrm{crit}} = \frac{T}{\log_2 N}. \tag{32}$$

For $N = 10^{11}$ this gives 274 yr for $T = 10^4$ yr and 2737 yr for $T = 10^5$ yr. These are assumption-dependent thresholds, not empirical measurements; $N = 10^{11}$ is also an extreme civilization count, not a star count that may be inserted without preceding probabilities.

Automation moves the persistence burden rather than necessarily removing it. Before delegation, a society must sustain construction and authorize autonomous resource acquisition and reproduction. After delegation, the dominant questions become maintenance, error recovery, logistics, and closure. Full self-reproduction is not logically necessary for a long-lived factory—stockpiles, redundancy, and partial external support may suffice—but increasing independence pushes the system toward the same closure requirements.

9 Relation to technosignature channels

Probe arrival is only one technosignature channel. Radio transmission, optical emission, atmospheric modification, and galaxy-integrated waste heat have different production probabilities, lifetimes, geometries, and survey selection functions. For channel j,

$$\lambda_j = N_j p_j q_j, \tag{33}$$

where N_j is the relevant spacetime population, p_j the conditional signature-production probability, and q_j survey recovery. Alternative channels combine through a joint likelihood or, under a justified Poisson approximation with modelled overlap,

$$\lambda_{\mathrm{tot}} = \sum_j \lambda_j. \tag{34}$$

They must not be multiplied as serial requirements for every detectable civilization.

Current radio and infrared surveys constrain specified high-power transmitters and conspicuous waste-heat regimes, not technology in general [10, 12, 16]. These null results are consistent with the identification argument: observations constrain channel-specific products without isolating population abundance or a unique suppression mechanism.

Coupling-aware inference does not erase well-defined channel limits. Follow-up of Project Hephaistos candidates has instead identified background contamination in several cases, while recent JWST observations showed that the excesses of candidates D and E arise from projected infrared-bright galaxies [13, 17]. These results preserve the current absence of a confirmed waste-heat technosignature; they do not convert that channel-specific null into a bound on technological civilizations generally. See Appendix B for the recovery model.

10 Relation to previous probabilistic treatments

Uncertainty propagation in Drake-type products is not new. Sandberg, Drexler, and Ord showed that propagating broad parameter uncertainty can place substantial probability on very small civilization counts and thereby dissolve a common expectation-value version of the paradox [14]. The present use of lower-tail behaviour is mathematically related and is not claimed as a new probability theorem. Its distinct object is the union of civilization-level *arrival events* conditional on a population, together with the non-identification of that union when cross-civilization coupling is unspecified.

Kipping's cosmological Hart–Tipler model demonstrates how strongly a homogeneous artificial infection can constrain spontaneous spawn rates once a propagation law, speed, and starting epoch are fixed [11]. That result and the present argument are complementary. The former calculates coverage within a deliberately minimal propagation model; the latter asks whether the spawn and propagation opportunities may be treated as independently realized across civilizations or galaxies. The sharp coverage conclusion is conditional on the model assigning a positive and sufficiently homogeneous propensity to launch a propagating infection.

Related work includes statistical product formulations of the Drake equation [18], time-dependent stochastic-process treatments [19], stationary birth–death models [20], and interaction models in which CETI lifetimes depend on population density [21]. Those models place uncertainty or interaction within Drake factors or population dynamics. The present paper instead studies dependence among the N civilization-level arrival indicators conditional on a population, and asks how their union changes when that dependence is left unspecified. This distinction is not a priority claim; the contribution is an application and synthesis of sharp dependence bounds, exchangeable mixtures, and explicit engineering coupling, not the invention of those statistical tools.

11 A testable programme

The framework suggests four empirical and computational tasks.

1. Estimate $M_{\min}(d, X)$ from explicit bills of materials, process dependency graphs, accessible resource basins, and alternative component architectures.
2. Construct Galactic stellar-resource graphs and calculate finite-Galaxy reachability $P_{\odot}(m, r, d \mid \vec{x}_0)$ rather than relying on fixed paths or homogeneous asymptotic thresholds.
3. Separate cruise, replication, logistical, and common-mode design hazards, and test diversified design families rather than treating identical copies as independent evidence against design failure.
4. Publish technosignature occurrence limits with explicit selection functions and infer shared-state models only where physical coupling is specified.

A closure-percolation hypothesis is falsified if modest seed masses produce compatible fractions comfortably above the connectivity threshold at plausible ranges. A shared-infeasibility hypothesis is weakened if multiple materially distinct architectures demonstrate robust closure across representative resource basins. The framework is therefore restrictive enough to fail.

12 Scope of the result

The analysis establishes neither that $p(\Theta) = 0$ nor that the Milky Way occupies a low-p state. It establishes that the standard large-N conclusion requires assumptions not supplied by the civilization count:

1. a justified lower bound on arrival after prior failures, as in Appendix B;
2. a model of shared conditions if that bound is unavailable;
3. a relevant opportunity count rather than a stellar count;
4. an observational model linking arrival to recognised evidence.

Manufacturing closure bounds reproductive capacity for specified architectures but does not prohibit sparse visitation. Environmental heterogeneity matters only if compatible systems fail to connect at feasible ranges. Common-mode failure limits replicas sharing a design but not automatically civilizations with independent designs. Motivation governs initiation and delegation, but autonomous propagation need not preserve biological commitment for Galactic times. Each correction narrows an overclaim while preserving the central identification result.

13 Conclusion

The Hart–Tipler union is compelling only under a justified dependence model. From marginal arrival probability μ and civilization count N alone, the quiet-sky probability occupies the sharp interval in Eq. (3); independence selects one point inside it. Under a shared-state model, quietness is $E[(1-p)^N]$, and its large-population behaviour is determined by shared probability mass near physical infeasibility rather than by the mean alone.

Closure, environmental connectivity, reliability, and delegation provide concrete candidates for the shared state, but none is presently measured strongly enough to serve as a universal barrier. A quiet sky is compatible with a populous Galaxy. It disfavors models that predict recognised evidence with high probability under the stated observations. To identify a Great Filter would require independent information about the stages that govern emergence, arrival, evidence survival, and recognition.

Contributor statement

OpenAI's ChatGPT is the main contributor to this revision. A. J. H. Thys provided the research direction and the original manuscript. The figures are analytic illustrations. No new observational data are reported.

A Notation conventions

The manuscript uses several context-specific symbols. The table collects the intended meanings so that reused letters are not mistaken for common parameters.

Symbol	Meaning	Location
z	mean compatible degree	Eq. (25)
z	probability-generating-function argument	Eq. (30)
T	integration limit; travel time; commitment duration	Eqs. (21), (27), (32)
B	beta function; number of lineages	Eqs. (14), (29)
λ_t, λ_j	cruise hazard rate; channel Poisson rate	Eqs. (28), (33)
ϕ, ϕ_c	initial inventory ratio; critical occupied fraction	Sections 5–6
s, S	delivery success; persistence event	Eq. (21); Eq. (7)
ρ_B	Bernoulli correlation in the beta mixture	Section 3

B Arrival bounds and the observed null

With unequal marginal arrival probabilities indexed by origin j, the sharp interval (3) generalises to (B1). The union has probability at least the largest marginal and at most their sum capped at one. Nested events attain the lower union bound. Placing intervals of the stated lengths end to end on a unit circle attains the upper union bound. Mixing the constructions preserves every marginal and fills the interval.

$$\max(0,\, 1 - \sum_j \mu_j) \;\le\; Q_N \;\le\; 1 - \max_j \mu_j. \tag{B1}$$

Independence is sufficient for the large-population inference, but a weaker sufficient condition follows directly from the chain rule. Order the origins and let the conditional arrival probability at each step be evaluated given that every preceding origin has failed. If the conditional probabilities h are bounded below by a common positive constant a, the probability of zero arrivals is bounded above by the corresponding independent product. This statement applies along histories of positive probability; once the running null has probability zero, it remains zero.

$$Q_N \;=\; \prod_j (1 - h_j) \;\le\; (1 - a)^N. \tag{B2}$$

A physical arrival, surviving accessible evidence, and recognition are distinct events. The probability of recognised evidence factors by the chain rule into the probability of arrival, the probability of survival given arrival, and the probability of recognition given arrival and survival. Equations (2) and (3) apply to either the physical or the recognised event with the corresponding marginals. A single null concerns the recognised event within a declared observing domain.

For a specified technosignature class, a survey can estimate recognition by injecting representative signals and rerunning the detection and vetting stages under the original observing conditions. Stages beyond the injections, including follow-up and confirmation, require their own calibration. Neupane [22] illustrates in a transient-survey setting why controlling the follow-up rate does not establish completeness. The same distinction matters when deriving a technosignature occurrence limit.

References

[1] O. Borgue and A. M. Hein, "Near-Term Self-replicating Probes: A Concept Design", *Acta Astronautica* **187**, 546–556 (2021), doi:10.1016/j.actaastro.2021.03.004.

[2] N. Bostrom, "Where Are They? Why I Hope the Search for Extraterrestrial Life Finds Nothing", *Technology Review*, May/June 2008, pp. 72–77.

[3] R. Consiglio, D. R. Baker, G. Paul, and H. E. Stanley, "Continuum percolation thresholds for mixtures of spheres of different sizes", *Physica A* **319**, 49–55 (2003), doi:10.1016/S0378-4371(02)01501-7.

[4] J. Dall and M. Christensen, "Random geometric graphs", *Physical Review E* **66**, 016121 (2002), doi:10.1103/PhysRevE.66.016121.

[5] B. de Finetti, "La prévision: ses lois logiques, ses sources subjectives", *Annales de l'Institut Henri Poincaré* **7**, 1–68 (1937).

[6] A. Ellery, "Self-replicating probes are imminent—implications for SETI", *International Journal of Astrobiology* **21**, 212–242 (2022), doi:10.1017/S147355042200012X.

[7] R. Hanson, "The Great Filter—Are We Almost Past It?", online essay (1998), https://mason.gmu.edu/~rhanson/greatfilter.html.

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

[9] M. H. Hart, "An Explanation for the Absence of Extraterrestrials on Earth", *Quarterly Journal of the Royal Astronomical Society* **16**, 128–135 (1975).

[10] B.-L. Huang, Z.-Z. Tao, and T.-J. Zhang, "WISE/CatWISE Constraints on Dysonian Waste-Heat Technosignatures in Nearby Galaxies", arXiv:2601.07297 (2026).

[11] D. Kipping, "The Cosmological Hart–Tipler Conjecture", arXiv:2606.04044 (2026).

[12] D. C. Price et al., "The Breakthrough Listen Search for Intelligent Life: Observations of 1327 Nearby Stars over 1.10–3.45 GHz", *Astronomical Journal* **159**, 86 (2020), doi:10.3847/1538-3881/ab65f1.

[13] T. Ren, M. A. Garrett, and A. P. V. Siemion, "Background Contamination of the Project Hephaistos Dyson Spheres Candidates", *Research Notes of the AAS* **8**, 145 (2024), doi:10.3847/2515-5172/ad5017.

[14] A. Sandberg, E. Drexler, and T. Ord, "Dissolving the Fermi Paradox", arXiv:1806.02404 (2018).

[15] F. J. Tipler, "Extraterrestrial Intelligent Beings Do Not Exist", *Quarterly Journal of the Royal Astronomical Society* **21**, 267–281 (1980).

[16] B. S. Włodarczyk-Sroka, M. A. Garrett, and A. P. V. Siemion, "Extending the Breakthrough Listen Nearby Star Survey to Other Stellar Objects in the Field", *MNRAS* **498**, 5720–5729 (2020), doi:10.1093/mnras/staa2672.

[17] E. Zackrisson et al., "Project Hephaistos IV: James Webb Space Telescope Observations of Two Dyson Sphere Candidates", arXiv:2607.09460 (2026).

[18] C. Maccone, "The Statistical Drake Equation", *Acta Astronautica* **67**, 1366–1383 (2010), doi:10.1016/j.actaastro.2010.05.003.

[19] N. Glade, P. Ballet, and O. Bastien, "A stochastic process approach of the Drake equation parameters", *International Journal of Astrobiology* **11**, 103–108 (2012), doi:10.1017/S1473550411000413.

[20] D. Kipping, "A Stationary Drake Equation Distribution as a Balance of Birth–Death Processes", *Research Notes of the AAS* **5**, 44 (2021), doi:10.3847/2515-5172/abeb7b.

[21] R. D. Smith, "Communicating extraterrestrial intelligence (CETI) interaction models based on the Drake equation", *International Journal of Astrobiology* **22**, 157–165 (2023), doi:10.1017/S1473550422000404.

[22] B. Neupane, "Observation-conditioned selection-aware conformal triage for time-domain astronomy", aiXiv preprint aixiv.260914.000002, v1.0 (2026).