Horizons, Hush, and Design: From Olbers and Fermi to a No-Signalling Telos with q→τ

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

October 20, 2025

Why is the night dark and the galaxy quiet? Olbers' paradox expects a sky bright as a stellar surface; Fermi's paradox expects a sky loud with beacons. Both follow from naive "shell" sums that ignore horizons, evolution, and the economics of signalling. This paper restarts from first principles with a minimal mix of words and math. For brightness, each shell contributes dI ~ n_* L_* dr because r^2 source growth cancels 1/r^2 dimming; the integral diverges if R -> infinity. For technosignatures, each shell would add dE ~ const * dr if civilizations radiated isotropically and continuously. Reality imposes cuts: finite age (R finite), expansion (extra (1+z) losses), nonstationary star formation, and, for technosignatures, tight beams, short duty cycles, narrow overlap windows, and codes that look like noise. We propose a unifying lens: no-signalling is a load-bearing constraint that rate-limits galactic coordination without forbidding nonlocal correlations. In plain terms, "nonlocal but not chatty." The result is a lawful hush: darkness in the visible, silence in our bands, both compatible with abundant but disciplined sources. We develop copy-safe shell integrals, a detectability factor D that collapses the Fermi sum, and testable predictions for a weak "almost-signal" background that grows statistically with time rather than arriving as a postcard.

Keywords: Olbers paradox, Fermi paradox, no-signalling, horizons, redshift, shell integral, technosignatures, detectability factor, information causality, q→τ.

Outline

1. Preliminaries: Shell Sums Without Romance
   1.1 Radiative brightness: $dI = n_* L_* \, dr$, $I = \int_0^{R} n_* L_* \, dr$
   1.2 Technosignature expectation: $dE = 4\pi r^2 n_{civ} p_{det}(r) \, dr$

2. Olbers Rewritten: Finite R and Fading Shells
   2.1 Finite age: $R \sim c T_{universe}$ makes I finite
   2.2 Expansion: extra $(1+z)$ factors reduce visible I
   2.3 Source evolution: $\rho_{SFR}(t)$, dust re-emission, K-corrections

3. Fermi Rewritten: Detectability Collapses the Integral
   3.1 Geometry and sensitivity: $SNR(r) \sim EIRP/(4\pi r^2) * \sqrt{t_{int} B}/(k T_{sys})$
   3.2 Detectability factor $D = f_{beam} f_{duty} f_{life} f_{match} f_{leak} \ll 1$
   3.3 Effective integral: $E_{eff} = \int_0^{R} 4\pi r^2 n_{civ} p_{det}(r; D) \, dr$

4. No-Signalling as Telos: Nonlocal Without a Phone Line
   4.1 Plain-language statement and why local marginals do not budge
   4.2 Safety rails: causality, energy accounting, control limits
   4.3 Mapping to practice: why optimal networks minimize leakage

5. The "Almost-Signal" Background
   5.1 Subthreshold residue: $I_{sub} = \int_0^{R} n_{civ} P_{leak}(r)/(4\pi r^2) \, 1_{\{SNR<SNR_*\}} \, dr$
   5.2 Statistical growth: excess variance $\sim c * \sqrt{T_{obs}}$
   5.3 Meta-signatures: occupancy, harmonics, weak repeats across surveys

6. Search Program: Statistics Over Beacons
   6.1 Quantify exposure: $f_{match} = V_{searched} / V_{possible}$ (freqtimecodesky)
   6.2 Publish EIRP_ (r) curves; stack nulls; calibrate RFI rigorously
   6.3 Cross-instrument coincidence tests that scale with time

7. $q \to \tau$ as Learning Under Horizons
   7.1 Replace ideal infinities with constrained objectives
   7.2 Local experiments as curvature reducers of $D_{KL}(\tau || q)$
   7.3 Ethics of slow coherence: peace over spectacle

8. Falsifiable Edges
   8.1 Lab: stronger-than-classical yet non-signalling bounds
   8.2 Sky: persistent, weak drift in anomaly statistics
   8.3 Crypto: quantum key distribution as everyday sentinel

9. Conclusion: The Mercy of Distance

10. References

1. Preliminaries: Shell Sums Without Romance

1.1 Radiative brightness: $dI = n_* L_* \, dr$, $I = \int_0^R n_* L_* \, dr$

Set-up (plain). Put the observer at $r = 0$. Take a thin spherical shell at radius r with thickness dr. If space were uniformly filled with stars of average luminosity L_* at number density n_* (both assumed constant for the toy model), then:

- Stars in the shell: $dN(r) = 4\pi r^2 * n_* * dr$
- Flux from one star at distance r: $F_*(r) = L_* / (4\pi r^2)$

Per-shell surface brightness. Multiply the two:

$$dI(r) = dN(r) * F_*(r) = [4 \pi r^2 n_* \, dr] * [L_* / (4 \pi r^2)] = n_* L_* \, dr$$

The r^2 from geometry cancels the $1/r^2$ dimming. Every shell adds the same amount of brightness.

Integral (the naive Olbers sum).

$$I_{naive} = \int_0^R n_* L_* \, dr = n_* L_* R$$

- If $R \to \infty$ (static, eternal universe), I_{naive} diverges linearly: paradox.
- In the real universe R is finite (light-travel/horizon), and the per-shell term is further suppressed by expansion and evolving spectra. Practically, we replace the constants with an “attenuation” $A_{rad}(r)$ and evolving luminosity density $L_*(r)$:

$$I_{real} = \int_0^R n_*(r) * L_*(r) * A_{rad}(r) \, dr$$

with $A_{rad}(r)$ encoding redshift factors, K-corrections, and bandpass loss (most reradiated energy lands in IR, not visible). Result: finite, small optical background—hence a dark night sky.

1.2 Technosignature expectation: $dE = 4\pi r^2 * n_{civ} * p_{det}(r) * dr$

Set-up (plain). Let n_{civ} be the space density of civilization sites capable of producing detectable emissions. In an idealized picture (everyone shouting, all the time, in all directions), the shell at radius r contributes:

$$dE_{naive} \sim [4 \pi r^2 n_{civ} \, dr] * [1/r^2] = \text{const} * dr$$

so E would “diverge” like $\int_0^R dr$, just as in Olbers. But real technosignals are scarce, beamed, intermittent, and often indistinguishable from noise unless you know the code.

Detection is a probability, not a given. Let p_det(r) be the probability that our survey actually detects at least one transmission from that shell. Then

dE(r) = 4*pi*r^2 * n_civ * p_det(r) * dr

E = ∫_0^R 4*pi*r^2 * n_civ * p_det(r) dr

The whole job is to understand p_det(r).

A useful factorization. Write

p_det(r) = p_geom(r) * p_snr(r) * D

where

- Geometry (beaming/pointing):
  If a transmitter illuminates solid angle Omega_beam,
- p_geom(r) ≈ Omega_beam / (4*pi) = f_beam (typically << 1)

plus any chance that their beam happens to sweep us during our dwell time.

- Signal-to-noise / sensitivity:
  With effective isotropic radiated power EIRP and system temperature T_sys, bandwidth B, and integration time t_int,
- SNR(r) ≈ [EIRP / (4*pi*r^2)] * sqrt(t_int * B) / (k * T_sys)
- p_snr(r) ≈ 1{ SNR(r) >= SNR_* }   (indicator of crossing threshold)

This creates a sensitivity horizon r_max beyond which p_snr = 0 for that setup.

- Behavioral/interop “detectability factor” D:
- D = f_duty * f_life * f_match * f_leak

with

  - f_duty ∈ [0,1]: time fraction transmitting in our observable mode
  - f_life ∈ [0,1]: temporal overlap (L_civ / T_window)
  - f_match ∈ [0,1]: fraction of the freq*time*code space we actually searched
  - f_leak ∈ [0,1]: fraction of power that escapes private, line-of-sight links

Putting it together. A simple, copy-safe surrogate is

p_det(r) ≈ f_beam * D * 1{ r <= r_max }

so the integral becomes

$$E_{eff} \approx \int_0^{r_{max}} 4\pi r^2 \cdot n_{civ} \cdot (f_{beam} \cdot D)\, dr$$

$$= (4/3) \cdot \pi \cdot r_{max}^3 \cdot n_{civ} \cdot f_{beam} \cdot D$$

Key point: even if r_{max} is large, the product $f_{beam} \cdot D$ is typically tiny (tight beams, short overlap windows, encrypted/noise-like codes, minimal leakage). That collapses the effective expectation E_{eff} to a small number—turning the imagined roar into a disciplined hush.

A slightly richer form (no step function). If you keep the SNR law explicitly,

$$p_{snr}(r) \approx 1\{ EIRP / (4\pi r^2) \geq SNR_* \cdot (k T_{sys}) / \sqrt{t_{int} \cdot B} \}$$

you get a smooth selection function $p_{snr}(r)$ that falls rapidly with r; integrating $4\pi r^2 \cdot p_{snr}(r)$ still yields a finite contribution concentrated inside a practical r_{max}, not a divergent sum.

Bottom line.

- Radiative sky (Olbers): per-shell brightness $dI = n_* L_* dr$; real-world $A_{rad}(r)$ and finite R keep I finite and dim.
- Technosignature sky (Fermi): per-shell expectation $dE = 4\pi r^2 n_{civ} p_{det}(r) dr$; realistic $p_{det}(r)$ (geometry, sensitivity, and $D \ll 1$) makes E small.
  Both "paradoxes" vanish once you write the actual shell integrals with the right cutoffs and scarcity multipliers.

2. Olbers Rewritten: Finite R and Fading Shells

2.1 Finite age: $R \sim c \cdot T_{universe}$ makes I finite

Idea in words.
If the universe has a birthday, light from beyond a finite radius simply has not had time to arrive. The naive "add one equal slice per shell forever" breaks because there are not infinitely many contributing shells.

Minimal math.
Let R_h be the horizon scale. In the simplest picture,

$$R_h \sim c \cdot T_{universe}$$

$$dI(r) = n_{star} \cdot L_{star} \cdot dr$$

$$I_{naive} = \int_0^{R_h} n_{star} \cdot L_{star} \cdot dr = n_{star} \cdot L_{star} \cdot R_h < \infty$$

So even before adding any other realism, a hard cutoff at R_h turns the divergence into a finite number. In practice, we do not literally integrate in physical radius; we integrate over lookback time or redshift, but the point is the same: a finite upper limit replaces “to infinity.”

Takeaway.
A universe with a beginning guarantees a finite optical background on first principles. No paradox required.

2.2 Expansion: extra $(1+z)$ factors reduce visible I

Idea in words.
Even when distant light reaches us, cosmic expansion stretches its wavelength and dilates its arrival rate. Energy per photon drops; arrival rate drops; and much of the power is shifted out of the visible band. Distant shells do not contribute “equally” anymore: they are heavily faded.

Minimal math (Tolman dimming and band loss).
For surface brightness in a fixed rest-frame band, expansion induces extra factors that suppress the observed intensity:

per-photon energy loss $\sim 1 / (1+z)$

arrival-rate time dilation $\sim 1 / (1+z)$

angular-area stretching $\sim 1 / (1+z)^2$

net surface-brightness dim $\sim 1 / (1+z)^4$

If we write a more honest integral over redshift z with a bandpass attenuation $A_{band}(z)$:

$$I_{obs} = \int_0^{z_{max}} j(z) \cdot A_{band}(z) \cdot \left[\frac{1}{(1+z)^4}\right] \cdot \left(\frac{dV}{dz}\right) \, dz$$

where

- $j(z)$ is the comoving luminosity density (power per volume),
- $A_{band}(z)$ accounts for how much of the redshifted spectrum falls inside our visible filter (often $\ll 1$ as z grows),
- (dV/dz) is the comoving volume element.

Both the $1/(1+z)^4$ factor and the shrinking $A_{band}(z)$ at high z crush the contribution of far shells in the visible.

Takeaway.
Expansion does double duty: it dims and it de-colors (shifts out of the visible). The “equal-shells” intuition fails hard once (1+z) enters.

2.3 Source evolution: rho_SFR(t), dust re-emission, K-corrections

Idea in words.
Stars are not eternal streetlamps. The cosmic factory that makes them speeds up and slows down; young, dusty galaxies reprocess starlight into the infrared; and the part of a source’s spectrum that lands in our chosen filter changes with redshift. All of this reduces visible-band brightness compared with the naive constant fog of suns.

Minimal ingredients.

1. Evolving emissivity via star-formation history.
   Let rho_SFR(z) be the cosmic star-formation rate density. The bolometric luminosity density j_bol(z) roughly tracks it (modulo stellar populations and IMF):

$$j_{bol}(z) \sim \eta \cdot \rho_{SFR}(z)$$

where eta converts SFR into luminosity. Observationally, rho_SFR(z) rises from z~0 to a peak around z~2, then declines; it is not a flat constant. Thus distant shells do not contribute a steady “photosphere per layer.”

2. Dust absorption and IR re-emission.
   A large fraction of UV/optical starlight is absorbed by dust and re-emitted at longer wavelengths. For the visible band we observe:

$$j_{vis}(z) = j_{bol}(z) \cdot f_{vis}(z)$$

with $f_{vis}(z) < 1$ and often decreasing at higher z where galaxies are dustier. The “lost” energy reappears in the mid/far-IR, not in the visible night sky.

3. K-corrections (spectral shifting through a fixed filter).
   A source’s rest-frame spectrum S_nu(rest) sampled through our observed filter F_nu(obs) is altered by redshift:

$$K(z) = (1+z) \cdot \left[\int S_\nu(\text{rest}) \cdot F_\nu(\text{obs}) \, d\nu_{obs} \right] / \left[\int S_\nu(\text{rest}) \cdot F_\nu(\text{rest}) \, d\nu_{rest} \right]$$

As z grows, stellar spectral peaks slide out of our filter; K(z) typically reduces visible contributions from high-z shells.

Putting the pieces in one copy-safe integral.
Collecting expansion, evolution, dust, and K-corrections:

I_visible = integral_0^{z_max}

j_bol(z)

* f_vis(z)   (dust re-emission penalty)

* K(z)   (filter mismatch penalty)

* [1 / (1+z)^4]   (Tolman dimming)

* (dV/dz) dz

Every bracketed factor is <= 1, many shrink with z. The result is finite and small in the visible, consistent with a dark night sky punctuated by discrete sources.

Takeaway.
Three disciplines tame the naive divergence:

- a finite horizon R (or z_max),
- expansion that dims and reddens,
- evolving, dusty sources whose spectra slide out of our band.
  Olbers' "bright night" vanishes once we honor these real, multiplicative fades.

3. Fermi Rewritten: Detectability Collapses the Integral

3.1 Geometry and sensitivity: SNR(r) ~ EIRP/(4*pi*r^2) * sqrt(t_int*B)/(k*T_sys)

Plain idea. A transmitter's power spreads; your receiver has noise. Detection at range r is a race between received flux and the system's noise floor.

Minimal model.

- Effective isotropic radiated power: EIRP (watts).
- Free-space spreading: power density ~ EIRP/(4*pi*r^2).
- Radiometer gain (single polarization): integration time t_int, bandwidth B, system temperature T_sys, Boltzmann k.

A convenient, copy-safe radiometer form is:

SNR(r) ~ [EIRP / (4*pi*r^2)] * sqrt(t_int * B) / (k * T_sys)

Detection requires SNR(r) >= SNR_* (a threshold set by pipelines and false-alarm budgets). This defines an effective range:

r_max ~ sqrt{ [EIRP / (4*pi)] * [sqrt(t_int * B) / (k * T_sys * SNR_*)] }

Two levers dominate: geometry and listening. Tight beams push EIRP at the target way up (but only where the beam points). Long dwell and matched bandwidth push sqrt(t_int*B) up (but only where we stare and only for signals living in that band).

Key lesson. Even heroic hardware yields a finite r_max. Most shells live beyond it for any given search setup.

3.2 Detectability factor D = f_beam * f_duty * f_life * f_match * f_leak << 1

p_det(r) is not "1 inside r_max, 0 outside." It is squeezed by behavior and interoperability. A practical factorization is:

D = f_beam * f_duty * f_life * f_match * f_leak

- f_beam (beaming/pointing): sky fraction illuminated.
  If a link uses solid angle Omega_beam, then f_beam = Omega_beam/(4*pi).
  Example: a 1 deg full-width beam has Omega_beam ~ pi*(1 deg in rad)^2 ~ 1e-3 sr, so f_beam ~ 1e-4 to 1e-5. Tighter is common.
- f_duty (on-air time fraction): fraction of time the transmitter is radiating in the searched mode.
  Scheduled links, bursty maintenance, or rare heartbeats push f_duty to 1e-3...1e-6.
- f_life (overlap window): fraction of the galactic window when both sides are simultaneously able and willing.
  A 1e4 yr "radio-leaky" phase in a 1e9 yr context gives f_life ~ 1e-5 (and that is generous).
- f_match (parameter-space overlap): fraction of the huge freq * time * modulation * code * polarization space that our survey actually covered with the right hypotheses.
  For modern, capacity-seeking codes that look noise-like, f_match can be the tightest bottleneck (often effectively << 1e-6 for any single program).
- f_leak (spill outside private channels): modern networks push energy into fibers, tight optical lines of sight, and authenticated beams. Accidental spill becomes small by design: f_leak ~ 1e-2 ... 1e-6.

Order-of-magnitude feel. A sober product is easy to make tiny:

$f_{beam} \sim 1e\text{-}5$

$f_{duty} \sim 1e\text{-}4$

$f_{life} \sim 1e\text{-}5$

$f_{match} \sim 1e\text{-}6$

$f_{leak} \sim 1e\text{-}3$

$D \sim 1e\text{-}23$ (illustrative)

Even if each factor is a bit larger, the product remains microscopic. That is the whole point: disciplined engineering plus short overlap converts the naive "should be loud" into "almost never seen."

3.3 Effective integral: $E_{eff} = \int_0^{R} 4 \pi r^2 \cdot n_{civ} \cdot p_{det}(r; D)\, dr$

The shell accounting. Let n_{civ} be the number density of candidate sites. A thin shell contributes

$$dE(r) = 4 \pi r^2 \cdot n_{civ} \cdot p_{det}(r; D) \cdot dr$$

You may use a hard-cutoff surrogate,

$$p_{det}(r; D) \approx D \cdot 1\{ r \le r_{max} \}$$

to get a transparent closed form:

$$E_{eff} \approx \int_0^{r_{max}} 4 \pi r^2 \cdot n_{civ} \cdot D\, dr$$

$$= (4/3) \cdot \pi \cdot r_{max}^3 \cdot n_{civ} \cdot D$$

Everything here is finite:

- r_{max} is finite (Section 3.1).
- D is tiny (Section 3.2).
- n_{civ} may be small and clustered, not uniform.

A smoother form (keeping sensitivity explicitly). Replace the step with a selection function:

$$p_{det}(r; D) = D \cdot p_{snr}(r)$$

$$p_{snr}(r) \approx \Phi(SNR(r) - SNR_*)$$

where Phi is a soft threshold (e.g., a sigmoid or error-function model for the pipeline's detection probability). Because SNR(r) ~ 1/r^2, p_snr(r) collapses rapidly with r, so the integral is dominated by the near shells even before multiplying by D.

Scaling insights.

- Doubling t_int increases r_max only as t_int^(1/4) (because r_max ~ sqrt(sqrt(t_int))). Distant gains are expensive.
- Narrowing beams helps the sender but hurts eavesdroppers: f_beam shrinks linearly with beam solid angle.
- The single most leverageful term for us is f_match: expanding search volume in frequency, time-of-day, dispersion trials, modulation families, and code hypotheses multiplies the chance that a real signal falls inside our tested set.

Bottom line. The naive Olbers-style divergence for technosignatures disappears once you replace "always on, all-sky, all-code" with the detectability reality:

E_eff ~ volume_nearby * n_civ * D

with D << 1. What remains to find is not a roar of beacons but faint, rare, and often subthreshold residue—the statistical fingerprint of disciplined networks operating under horizons.

4. No-Signalling as Telos: Nonlocal Without a Phone Line

4.1 Plain-language statement and why local marginals do not budge

Plain statement. Entanglement can correlate distant outcomes, but it cannot be used to *steer* what a distant observer sees without later, lightspeed comparison. Locally, each party's data still look like fair noise.

Copy-safe sketch. Let P(a,b | x,y) be the joint outcome distribution when Alice chooses setting x and Bob chooses y. *No-signalling* means Alice's local statistics do not depend on Bob's choice, and vice versa:

- sum_b P(a,b | x,y) = P(a | x) for all y
- sum_a P(a,b | x,y) = P(b | y) for all x

Equivalently: each side's *marginal* is fixed by its own setting, not by the other's. In practice, this is why changing x cannot imprint a readable message in Bob's solo data. Only after they exchange classical notes (at or below c) do the correlations reveal themselves.

Intuition. Nature permits “shared riddles” (nonlocal patterns), not remote dials. The riddle is solved only when both halves are compared.

4.2 Safety rails: causality, energy accounting, control limits

No-signalling is more than etiquette; it keeps several ledgers consistent.

Causality. If superluminal control were allowed, one inertial frame’s “send” could be another frame’s “receive in the past.” That opens time-loop exploits. No-signalling blocks any channel that would let you pick distant outcomes on demand, preserving a coherent before/after across frames.

Energy accounting. Useful work requires *predictable* bias. If you could pre-select favorable distant outcomes, you could siphon work from thermal noise (a Maxwell demon with no proper bill). With no-signalling, local data remain unbiased without side information, so the demon lacks a handle. Thermodynamic bounds like $k_B T \ln 2$ per reliable bit remain meaningful.

Control limits. Feedback systems assume you act on what you have already sensed. If you could receive tomorrow’s measurement today via FTL signalling, you could design anti-causal controllers that cancel disturbances before they exist, breaking robust trade-offs (the “waterbed” integrals) that keep engineering sane. No-signalling preserves those limits.

Communication economics. Markets, protocols, and proofs rely on latency that costs distance and power. Remove that cost and you invite instant global arbitrage and brittle monoculture. No-signalling re-imposes a universal price: any *actual* message must traverse space-time and pay in time and energy.

4.3 Mapping to practice: why optimal networks minimize leakage

Read Fermi through an engineer’s eyes: the quiet sky is not an accident; it is the efficient frontier under these rails.

Beams beat broadcasts. For a target at range r, received SNR obeys
$SNR(r) \sim EIRP / (4 \pi r^2) * \sqrt{t_{int} B} / (k T_{sys})$.
To meet a given SNR_* with minimal power, tighten the beam (reduce solid angle Ω_{beam}). That raises EIRP *to the target* while shrinking spill. Eavesdroppers off-axis see nothing.

Duty and schedule. If messages must pay latency, you schedule them. The optimal duty cycle f_{duty} is the minimum that meets reliability and throughput. Idle silence dominates the timeline.

Overlap windows. Let L_{civ} be an active phase and T_{gal} the relevant galactic timescale. The overlap fraction $f_{life} = L_{civ} / T_{gal}$ is small even for optimistic L_{civ}. Two parties must coincide in time, frequency, code, and pointing. Most pairs do not.

Codes that hide in the grass. Capacity-seeking links approach noise-like statistics. To outsiders, such signals are indistinguishable from thermal backgrounds unless their *exact* hypotheses are tested. That collapses f_{match}, the fraction of the giant (freq * time * modulation * code) space we actually cover.

Put it together. The practical detection probability per shell is
$p_{det}(r) = f_{beam} * f_{duty} * f_{life} * f_{match} * f_{leak} * p_{snr}(r)$,
with $p_{snr}(r)$ tied to the radiometer limit above. Each factor lies in [0,1] and, for mature networks, most are tiny. The product $D = f_{beam} f_{duty} f_{life} f_{match} f_{leak} \ll 1$. Hence the *effective* expectation
$E_{eff} = \int_0^R 4\pi r^2 n_{civ} p_{det}(r; D)\, dr$
is finite and small even when n_{civ} is not.

Design moral. Under no-signalling, the cheapest reliable message is local, narrow, authenticated, and nearly invisible to third parties. If many civilizations adopt these same cost-minimizing habits, the aggregate sky presents a lawful hush: nonlocal wonder without a phone line, coordination without collapse, and the kind of silence that still lets freedom grow.

5. The "Almost-Signal" Background

5.1 Subthreshold residue: $I_{sub} = \int_0^{R} n_{civ} * P_{leak}(r)/(4*\pi*r^2) * 1_{\{SNR<SNR_*\}}\, dr$

Idea. If mature networks minimize leakage, most emissions that graze us will be too faint or too brief to decode. They still add a tiny, structured residue to the sky's noise.

Minimal model (copy-safe).

- n_{civ} = number density of active sites (per volume).
- $P_{leak}(r)$ = power that *escapes* private beams/links toward us (watts).
- Free-space dilution: $1/(4*\pi*r^2)$.
- Detection gate: indicator $1_{\{SNR<SNR_*\}}$ keeps only events that fall below our single-event threshold.

Then the integrated subthreshold intensity in our band is

I_sub = ∫_0^{R} n_civ(r) * [P_leak(r) / (4*pi*r^2)] * 1_{SNR(r) < SNR_*} dr .

This is not “message energy we can read,” but a statistical nudge to second-order features (variance, kurtosis, weak line occupancy). Because SNR ~ EIRP/(4*pi*r^2) * sqrt(t_int*B)/(k*T_sys), most far shells contribute only below threshold; nearby shells are rare by volume. The combined effect is small but widespread.

Practical upshot. Do not expect content; expect a bias in the noise that survives careful calibration and grows with total observation time.

5.2 Statistical growth: excess variance ~ c * sqrt(T_obs)

Idea. Subthreshold power does not stack linearly like readable signals; it accumulates as a slow drift in statistics as you integrate.

Back-of-envelope.

- Let sigma_0^2 be the calibrated variance of natural + instrumental noise per analysis cell (a small patch in sky × frequency × time).
- Let delta_sigma^2 be the additional variance contributed by many faint, independent technoleak events. For stationary conditions, central-limit scaling gives detection significance roughly

S(T_obs) ~ [delta_sigma^2 / sigma_0^2] * sqrt(N_cells(T_obs))

~ c * sqrt(T_obs) ,

where N_cells (independent samples) grows ~ linearly with total on-sky time T_obs (after accounting for correlations and window functions). The constant c folds survey geometry, leak statistics, and cleaning losses.

Testable behavior.

- Under the null (no technoleak), the standardized statistic should wander ~ N(0,1) as T_obs increases.
- With a persistent subthreshold residue, the aggregate statistic drifts positively as +c*sqrt(T_obs). The slope c is tiny but nonzero.

Operational move. Design pipelines that (i) lock down sigma_0^2 through aggressive RFI excision and astrophysical modeling, then (ii) track the time-normalized drift of variance/kurtosis/line-density metrics across years. You are not decoding; you are weighing the haystack.

5.3 Meta-signatures: occupancy, harmonics, weak repeats across surveys

We cannot rely on content. We can look for meta-structure—patterns that are cheap for engineers to produce but unlikely for nature to mimic across independent datasets.

(a) Occupancy anomalies.
Partition parameter space into cells i (sky tile, frequency bin, dispersion trial, polarization, modulation score). For each cell, measure occupancy X_i (e.g., burst count, narrowband line presence, cyclostationary score). Define

$$T_{occ} = \sum_i w_i * (X_i - \mu_i(\text{natural})) / \sigma_i(\text{natural}) .$$

After training mu_i, sigma_i on clean control fields, genuine subthreshold technoleak yields a persistent positive T_occ drift with more data, not tied to known astrophysical or terrestrial cycles.

(b) Harmonic scaffolds.
Capacity-seeking links often leave weak harmonics (pilot tones, clock leakage). Search for statistically excess lines at {f0, 2f0, 3f0, ...} within Doppler/acceleration windows consistent with stellar reflex motion. The meta-signature is a comoving comb whose spacing follows a simple law yet recurs in multiple epochs.

(c) Weak repeats across surveys.
Real networks generate repeatable footprints along practical geometries (ecliptic crossings, crowded target cones). Expect low-SNR recurrences in the *same* (sky, frequency) cells across independent instruments and years, within natural drift bounds. Cross-survey coincidence is hard for RFI and random transients to fake.

(d) Modulation priors.
Engineered signals favor cheap demodulators and coding families (e.g., PSK/QAM, error-correcting codes) that impart subtle cyclostationarity or spectral-shape fingerprints. Build low-bias detectors for these traits and test whether their aggregate hit-rates exceed natural baselines after strict multiple-hypothesis control.

(e) Negative-space constraints.
Publish exposure: f_match = V_searched / V_possible over {freq, time, sky, modulation, dispersion}. A credible claim of “almost-signal” needs a denominator. The same meta-signatures

should persist when the analysis is repeated by others with different hardware but comparable f_match.

Bottom line. The “almost-signal” is not a postcard; it is a reproducible tilt in the sky’s statistics:

- energy present but below single-event thresholds,
- significance that grows like sqrt(T_obs),
- meta-patterns (occupancy, harmonics, weak repeats) that survive cross-instrument comparisons and astrophysical/RFI vetoes.

That is the falsifiable footprint you can target when optimal networks prefer privacy—and the cosmos keeps its hush.

6. Search Program: Statistics Over Beacons

6.1 Quantify exposure: f_match = V_searched / V_possible (freq*time*code*sky)

Why this matters. Absence is only evidence if we know our denominator. Report the exact fraction of the parameter space we actually touched.

Copy-safe bookkeeping.

- Let the total search volume be
- V_searched = A_sky * Delta_f * T_dwell * N_pointings * C_codes

where
A_sky = solid angle covered (sr),
Delta_f = total frequency span searched (Hz),
T_dwell = per-pointing coherent integration time (s),
N_pointings = number of distinct pointings/epochs,
C_codes = effective number of demod/coding hypotheses actually tested.

- Let the plausible volume be
- V_possible = A_all * Delta_f_all * T_window * C_all

where
A_all = 4*pi sr or a restricted target cone,
Delta_f_all = prior frequency span worth considering,

T_window = surveyable time window for overlap (s),
C_all = prior set of realistic modulation/coding families.

Then define

f_match = V_searched / V_possible in [0,1] .

Granular reporting.

- Partition param space into cells k with volume v_k and report a coverage map:
- f_match = sum_k v_k(searched) / sum_k v_k(possible) .
- Publish a table listing for each k: (sky tile, freq band, dwell time, code family, dispersion trials). This lets others stack exposure across projects.

Two minimal denominators to include in papers.

1. freq*time*sky only (codes marginalized)
2. freq*time*sky*codes (explicit code families and thresholds)

6.2 Publish EIRP_star(r) curves; stack nulls; calibrate RFI rigorously

Sensitivity in one line. For a threshold SNR_star,

EIRP_star(r) = 4*pi*r^2 * SNR_star * (k*T_sys) / sqrt(t_int * B)

Publish EIRP_star(r) as a curve (or table) per pipeline, band, and sky tile, with error bars from calibration uncertainty.

Null stacking.

- For each cell i, define a Bernoulli outcome D_i in {0,1} (1 = detection, 0 = null) with an expected false-alarm p_i under the null.
- Combine across surveys using a log-likelihood or Fisher-style statistic:
- X = -2 * sum_i ln(p_i)   (Fisher)

or, more robust to heterogeneous thresholds,

Z = sum_i w_i * (D_i - p_i) / sqrt(p_i*(1 - p_i))

with weights w_i proportional to exposure (e.g., v_i(searched)).

RFI discipline (must have).

- Front-end excision: stable masks for known emitters; per-epoch dynamic masks.

- Three-way veto: time, beam, polarization; reject events that light up near-field sidelobes.
- Geographic diversity: candidates must repeat from sites with independent local RFI.
- Data products: publish flagged intervals, mask footprints, and post-excision noise PSDs so others can replicate your null.

Calibration essentials.

- Inject synthetic lines with known EIRP_equiv to map detection efficiency epsilon(SNR).
- Track T_sys(t) and bandpass gain g(f,t); provide per-epoch radiometer budgets.
- Quote a single number for reproducibility:
- Exposure_metric = sum_tiles [A_tile * T_dwell(tile) * B(tile) * epsilon_tile]

6.3 Cross-instrument coincidence tests that scale with time

Why coincidence beats one-offs. Subthreshold residues should reappear weakly in the same (sky, freq) cells across independent instruments and years. Coincidence testing turns many near-misses into one strong claim.

Cell-based coincidence.

- Define cells i = (sky tile, freq bin, dispersion window).
- For survey s in S, compute feature X_{i,s} (e.g., narrowband occupancy, cyclostationary score).
- Standardize per survey using natural baselines:
- Z_{i,s} = (X_{i,s} - mu_{i,s}) / sigma_{i,s}
- Combine across surveys:
- Z_i_comb = sum_s w_s * Z_{i,s}

with weights w_s proportional to calibrated exposure for that cell.

Under the null, Z_i_comb ~ N(0,1). A persistent positive tail across growing S or longer T_obs indicates a real residue.

Time scaling expectation.

- If a subthreshold background exists, the aggregate detection significance should grow like
- S_total(T) ~ c * sqrt(T_obs_total)

after accounting for correlated noise and repeating RFI vetoes. Plot S_total versus sqrt(T_obs_total); a linear trend with nonzero slope c is your key diagnostic.

Doppler-aware repetition.

- For lines, test for comoving combs: frequencies {f0, 2f0, ...} drifting within a window set by plausible accelerations a_max and velocities v_max.
- Define a Hough-like accumulator in (f0, drift) space; stack votes from multiple epochs and instruments. Natural lines rarely produce consistent votes across independent sites within the same drift law.

Public artifacts to release.

- Cell index with sky coordinates, frequency edges, and dispersion grids.
- Per-cell Z_{i,s}, mu_{i,s}, sigma_{i,s}, and exposure weights.
- Coincidence scores Z_i_comb and their p-values after multiple-hypothesis control (e.g., Benjamini-Hochberg).

Decision rule (copy-safe).

- Declare a global excess if
- mean_i [1{ Z_i_comb >= z_star }] > q_star

where z_star and q_star are pre-registered (e.g., z_star = 3, q_star = 1e-3), and the excess survives scrambling tests (time/frequency shuffles) and RFI geography vetoes.

Bottom line. Make the denominator explicit (f_match), the sensitivity visible (EIRP_star curves), the nulls stackable (open masks and priors), and the claim scalable with time (coincidence statistics that rise like sqrt T). That is how a whisper becomes evidence.

7. q→τ as Learning Under Horizons

7.1 Replace ideal infinities with constrained objectives

Classical “paradox math” loves infinities: infinite time, uniform sources, listeners with perfect reach. Learning in the real world is different. We work under horizons (finite age, finite energy, finite attention) and must therefore set constrained objectives instead of chasing limit cases.

Let q be the learner’s current model and τ the world’s generative manifold (the “true” distribution/logos that actually produces data). The aim is not omniscience but steady improvement under constraints:

- Budgeted sensing: choose measurements that maximize expected learning per unit cost (time, power, risk).
- Bounded update: prefer updates that move q toward τ without overfitting scarce or biased data.
- Anytime stopping: at each step, produce a model that is good enough for action, given current uncertainty.

In short: replace “see everything forever” with “choose the next best move,” given your horizon. The curriculum is incremental.

7.2 Local experiments as curvature reducers of $D\_KL(\tau||q)$

A concrete way to say “get closer to τ” is to reduce a divergence between reality and our model. Use the forward Kullback–Leibler:

$$D\_KL(\tau \,||\, q) = \int \tau(x) \log[\, \tau(x) / q(x) \,]\, dx \;\geq\; 0$$

This quantity behaves like curvature in the landscape of beliefs: the further q bends away from τ, the larger the penalty. Local experiments act as curvature reducers by delivering data that most efficiently push q toward τ.

Three operational levers:

1. Sensing policy π(x): pick probing locations or queries x that maximize the expected drop $\Delta = E\_data[\, D\_KL(\tau \,||\, q\_before) - D\_KL(\tau \,||\, q\_after) \,]$
   subject to costs. This is active learning in copy-safe terms: ask questions where your ignorance is most expensive.
2. Update discipline: use estimators that trade bias and variance explicitly (regularization, priors, cross-validation). Goal: large Δ without creating new curvature elsewhere.
3. Model capacity schedule: increase model expressiveness only when data support it. Too little capacity and Δ plateaus; too much and q chases noise, raising D_KL off-manifold.

Practically, you will never compute $D\_KL(\tau||q)$ exactly (τ is hidden). You approximate it with held-out likelihoods, predictive scoring, or decision regret. The spirit holds: design local experiments that make the largest guaranteed dent in your global misspecification.

7.3 Ethics of slow coherence: peace over spectacle

Horizons are not only physical; they are moral. A culture intoxicated by shortcuts—by spectacle—tends to trade truth for velocity. The q→τ discipline points the other way: slow coherence over viral certainty.

Working rules:

- Truthfulness before reach. Prefer claims with receipts (methods, data, error bars) over claims with heat. Each receipt is a small Δ on D_KL(τ||q); each ungrounded flourish is a detour.
- Peacemaking as policy. When uncertainty is high, choose actions that reduce downstream conflict even if they slow immediate progress. In expectation, this enlarges the safe experiment set and unlocks future Δ that spectacle would foreclose.
- Design for cool minds. Interfaces and institutions should lower arousal, widen attention, and make disagreement legible. That is how many minds move q toward τ together, without burning the laboratory.
- Local covenants, global humility. Act where you can audit consequences. Share what generalizes. Accept that some horizons will not move today—and that this is not failure but the operating condition of a lawful world.

The ethical slogan is simple: become truer, not louder. Under horizons, peace is not passivity; it is the strategy that maximizes long-run learning.

8. Falsifiable Edges

8.1 Lab: stronger-than-classical yet non-signalling bounds

Goal. Push nonlocal correlations to their lawful ceiling while verifying that local marginals stay unchanged. If nature respects no-signalling, you can see S > 2 (Bell violation) yet never use it as a phone.

Copy-safe checklist.

- Loophole-free Bell with tight bounds.
  Measure CHSH S with space-like separation and random, fast settings.
  Classical: S <= 2. Quantum (Tsirelson): S <= 2*sqrt(2). No-signalling: S <= 4.
  Falsification targets:
  1. S > 2*sqrt(2) while local marginals remain setting-independent (no-signalling still holds). This would challenge standard quantum structure but not the telos.
  2. Any statistically significant dependence of P(a|x) on y or P(b|y) on x (i.e., local marginals budge with the distant choice) would violate no-signalling outright.

- Information-causality (IC) tests.
  Implement random-access coding tasks with m communicated bits. IC says Bob's total information gain $\le m$.
  Falsification targets:

  1. Performance exceeding the IC bound without extra communicated bits (suggests super-quantum yet still non-signalling correlations).
  2. Apparent IC violation with unchanged local marginals is tolerable for the telos; IC helps explain Tsirelson but is not strictly required.
  3. IC violation and steerable marginals would break the telos (true signalling).

- Pre-echo probes (anti-causal canary).
  Time-tag settings x,y and outcomes a,b with sub-ns resolution. Search for predictive leakage: $\mathrm{Corr}(a_t , y_{t+\delta})$ for $\delta<0$; similarly for b vs x.
  Null expectation: zero. Any reproducible pre-echo implies a controllable handle across space-like separation.

One-line falsifier.
A robust, loophole-free experiment showing setting-dependent local marginals is a direct hit on no-signalling; the rest of this paper's framing would need rewriting.

8.2 Sky: persistent, weak drift in anomaly statistics

Target. Not a postcard, but a reproducible tilt in sky statistics that grows with exposure if an "almost-signal" background exists.

Copy-safe scaffolding.

- Define cells i. Partition (sky, frequency, time, dispersion, polarization). For each cell, extract features X_i (e.g., narrowband occupancy, burst counts, cyclostationary score).
- Standardize per survey s.
  $Z_{i,s} = (X_{i,s} - \mu_{i,s}(\text{natural})) / \sigma_{i,s}(\text{natural})$.
- Stack across surveys.
  $Z_i^{comb} = \sum_s w_s \cdot Z_{i,s}$, with weights w_s proportional to calibrated exposure for that cell.
- Time-scaling test.
  Under the null, $\mathrm{mean}(Z_i^{comb}) \sim 0$. With subthreshold residue, the aggregate

significance S(T_obs) should grow like
S(T_obs) ~ c * sqrt(T_obs), after accounting for correlations.

- Doppler-aware combs.
  Accumulate weak votes in (f0, drift) space within physically plausible acceleration windows. Look for comoving harmonic sets {f0, 2f0, ...} repeating across years and instruments.

- Coincidence across geography.
  Require the same weak cells to rise at independent sites with different local RFI. Use BH or Holm corrections for multiple hypotheses, preregister thresholds (e.g., z_star=3, q_star=1e-3), and publish full exposure maps f_match.

Falsification dialectic.

- If the drift persists across instruments, seasons, and pipelines, natural and terrestrial baselines become increasingly implausible.

- If, with growing f_match and T_obs, the drift does not rise as sqrt(T_obs), tighten upper bounds on the leakage background and update D factors (Section 3.2) downward.

8.3 Crypto: quantum key distribution as everyday sentinel

Why crypto matters. Practical security already leans on "nonlocal but not chatty." Device-independent QKD (DI-QKD) uses Bell violations to certify keys without trusting devices, but its safety assumes no-signalling.

Operational canaries.

- Bell-parameter telemetry.
  Public QKD backbones should publish time series of CHSH S, key rates, and QBER. Expectations:

  1. S in (2, 2*sqrt(2)] with stable (low) QBER.

  2. No dependence of local singles on distant settings.

- Anomaly classes that would flag no-signalling failure.

  1. Pre-coordinated marginals: singles rates or bit distributions that shift instantly with distant setting choices.

2. Free-lunch reconciliation: extractable key rate surpassing composable security bounds without extra communicated bits or better channels (a real-world IC break).
3. Retrocausal artifacts: reproducible improvements in error correction when you condition on future setting logs, beyond selection-bias explanations.

- City-scale coincidence.
  Run mirrored DI-QKD links in separated metros. If an attacker exploited any FTL handle, correlated anomalies would appear across sites faster than light can carry classical coordination. Null expectation: no such pattern.

Falsifiable outcome.
If any QKD network demonstrates controllable, faster-than-light side-channels (setting-dependent marginals, reproducible pre-echoes, or key rates exceeding proven bounds without increased classical communication), that is de facto evidence against no-signalling. Conversely, year-over-year stability of S, QBER, and independence tests across heterogeneous hardware is accumulating evidence that the telos holds in practice.

Bottom line.
Three frontiers—lab, sky, and crypto—offer complementary, quantitative shots at the edges. Labs test whether correlations outrun quantum yet keep local silence; skies test for a faint, lawful hush that thickens with time; crypto runs a 24/7 audit that would be the first to scream if a hidden phone line existed.

9. Conclusion: The Mercy of Distance

The two great quiets—darkness in Olbers, silence in Fermi—were never failures of the world; they were failures of our limit-free expectations. Once we honor horizons, evolution, and bounded detectors, both “paradoxes” resolve into constrained integrals with small, law-shaped answers. The same discipline reframes our hopes for contact. Instead of a postcard from the stars, we should expect a lawful hush: nonlocality without a phone line, networks optimized for privacy, and leakage so thin it shows up only as statistics that thicken with time.

Seen this way, no-signalling is not a nuisance; it is a mercy. It keeps cause and effect legible, energy accounting honest, and control theory stable. It lets cultures learn locally, step by step, turning models q toward τ through experiments that reduce curvature rather than inflate spectacle. Distance prices coordination; time protects formation. Under these rates and prices, the cheapest reliable message is narrow, authenticated, and quiet. A galaxy full of such choices will look—must look—like ours.

Practically, the path forward is clear:

- Replace romance with accounting: explicit shell sums, explicit exposure denominators, explicit detectability factors D.
- Trade single, heroic "detections" for population tests whose significance scales as sqrt(T_obs).
- Treat device-independent cryptography as a standing audit: if a hidden phone line exists, it will show up there first.
- Teach institutions to prefer receipts over reach and to invest in methods that move q toward τ, even when the moves are slow.

There is a moral, too. Horizons are not bars on the window; they are scaffolds. They hold a space where agency can ripen without being crushed by instant, total coordination. In that space, truth has time to become common knowledge, and power has time to meet character. We do not need the universe to be louder to be kinder or truer. We need to listen longer, measure better, and build with the grain of a world that withholds shortcuts so freedom can grow.

That is the mercy of distance: a cosmos designed so that meaning must be made, not merely received.

References

- Clauser, J. F., Horne, M. A., Shimony, A., & Holt, R. A. (1969). Proposed experiment to test local hidden-variable theories. *Physical Review Letters*. Physical Review
- Popescu, S., & Rohrlich, D. (1994). Quantum nonlocality as an axiom. *Foundations of Physics*. SpringerLink
- Pawłowski, M., Paterek, T., Kaszlikowski, D., Scarani, V., Winter, A., & Żukowski, M. (2009). Information causality as a physical principle. *Nature*, 461, 1101–1104. Nature+1
- Hensen, B., et al. (2015). Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526, 682–686. Nature+1
- Giustina, M., et al. (2015). Significant-loophole-free test of Bell's theorem with entangled photons. *Physical Review Letters*, 115, 250401. Physical Review
- Zhang, W., et al. (2022). A device-independent quantum key distribution system for distant users. *Nature*, 607, 687–692. Nature
- Zapatero, V., Curty, M., & Lim, C. C.-W. (2023). Advances in device-independent quantum key distribution. *npj Quantum Information*, 9, 98. Nature
- Madau, P., & Dickinson, M. (2014). Cosmic star-formation history. *Annual Review of Astronomy and Astrophysics*, 52, 415–486. Astrophysics Data System+2sites.astro.caltech.edu+2
- Hogg, D. W., Baldry, I. K., Blanton, M. R., & Eisenstein, D. J. (2002). The K correction. *arXiv:astro-ph/0210394*. arXiv+2arXiv+2
- Lubin, L. M., & Sandage, A. (2001). The Tolman surface brightness test for the reality of the expansion. IV. A measurement of the Tolman signal… *Astronomical Journal*, 122, 1084–1103. arXiv+1
- Pahre, M. A., Djorgovski, S. G., & de Carvalho, R. R. (1996). A Tolman surface brightness test for universal expansion… *Astrophysical Journal Letters*, 456, L79–L82. arXiv+1
- "Tolman surface brightness test." Overview entry. (Accessed 2025). Wikipedia
- Harrison, E. R. (1990). The dark night sky riddle ("Olbers's paradox"). In *IAU Symposium 139*. Astrophysics Data System
- "Olbers' paradox." Overview entry. (Accessed 2025). Wikipedia
- Smith, S., et al. (2021). A radio technosignature search towards Proxima Centauri at 2–4 GHz. *Nature Astronomy*. (Includes EIRP sensitivity formulation). Nature

- Lorimer, D. R., & Kramer, M. (2005). *Handbook of Pulsar Astronomy*. (Radiometer equation widely referenced; see applications in later works). Examples citing: Joshi et al. 2009; Trudu et al. 2024. OUP Academic+1
- Harp, G. R., et al. (2015). Radio SETI observations of the anomalous star KIC 8462852. (Radiometer equation usage). *arXiv:1511.01606*. arXiv
- NASA Technosignatures Workshop Participants (2018). *NASA and the Search for Technosignatures: A Report from the NASA Technosignatures Workshop*. arXiv+1
- Sheikh, S. Z. (2019). The nine axes of merit for technosignature searches. *arXiv:1908.02683*. arXiv
- Tolman surface brightness dimming background (1+z)^4 and test history. Summary sources. Wikipedia+1
- Barrett, J., Linden, N., Massar, S., Pironio, S., Popescu, S., & Roberts, D. (2005). Popescu-Rohrlich correlations as a unit of nonlocality. *Physical Review Letters*, 95, 140401. Physical Review
- Nadlinger, D. P., et al. (2021). Experimental quantum key distribution certified by Bell's theorem. *arXiv:2109.14600*. arXiv
- Peacock, J. A. (1999). *Cosmological Physics*. (General background on Olbers/Tolman tests). Wikipedia

Notes: Where possible, we cite primary literature or authoritative reviews for each technical claim (Tolman (1+z)^4 dimming; cosmic SFR peak; K-corrections; CHSH/Tsirelson/PR boxes; information causality; radiometer/EIRP relations; DI-QKD and loophole-free Bell).