


A combination-space search for embedded technosignatures in solar and heliospheric archives

Robert Griffin · Dextra Inc., Massapequa, New York · draft of 13 September 2026

We call the approach set out here **behavioral SETI**: it looks for the behavior of a source rather than for a transmitter. Existing technosignature searches presuppose the channel. We consider the case in which *no dedicated radiating apparatus is assumed*: structure is carried by emission a star already produces, imposed by redistributing that flux rather than generating it, and neither carrier nor encoding is known in advance. Two constraints make the problem finite: a sender sharing no units with the receiver can use only **dimensionless** quantities — ratios and normalized combinations of measured observables — and a sender minimizing transmitted energy will pose a **proof-of-work gate**, cheap to set, expensive to solve, trivial to verify, so that no energy is spent on receivers that could not act on the message. We formalize the resulting **combination space** and report 29 searches of public solar and heliospheric archives against it. Twenty-five are null, three void, and one recovers a known signal as a positive control. Seventeen carry stated sensitivities, **thirteen of them verified by injection**. The deepest verified limit is 1.4×10^{-6} in fractional Lyman- α irradiance at two minutes, at 95% recovery; the remaining one, the sidereal fold, resisted a faithful injection (§4.2). Sensitivity is not the binding constraint: the same instrument, days and pipeline are $245 \times$ more sensitive at two minutes than at one day, and the limiting quantity is instead the fraction of the combination space examined. Acquiring seventeen further observables closed the pair space from 18% to **88%** — 382 of 435 pairs, the remainder having no overlapping record — and returned **no survivor in 510 tests spanning the Sun–heliosphere boundary**. Of 4,060 triples, 2,934 are reachable and **2,749 completed**, returning **void** for want of a matched three-body null. A first **cross-viewpoint** test, Earth line against Mars line, is null for a line-of-sight modulator to 8.9×10^{-4} at 3 d and 4.2×10^{-3} at 307 s, the fast-band figure set by the Mars instrument rather than by the geometry. Finally, a designer-side ordering of channels places the most plausible first-contact carriers in archives that already exist and are resolved at the required level. **Two of those rows are searched here for the first time** — SOHO/VIRGO SPM photometry at one minute over 27.3 yr, and SOHO/GOLF Doppler velocity at 20 s over 25.9 yr — both null, the latter to an injection-verified 1.46×10^{-5} in fractional mode frequency, which is **fifteen times short of the 10^{-6} a designer would set**. The remaining rows are specified.

Keywords: technosignatures · SETI · solar irradiance · search strategy · time-series analysis



The paper's honest headline: the depth a cost-minimizing designer would choose for each channel (diamonds) against the depth this program reaches (circles). Red rows have been searched and returned null; grey rows sit below our current capability. Section 3.6 explains the ordering.

1. Introduction

1.1 Every search so far has presupposed the channel

The modern search for technosignatures begins with a choice of channel. Cocconi & Morrison (1959) argued for the 21 cm line, and essentially every program since has followed the same logic: name a carrier both parties can identify, then search it as deeply as instruments allow. Wright, Kanodia & Lubar (2018) formalized SETI as a coverage problem, but their axes are sky position, frequency, sensitivity, polarization, modulation and repetition rate: **a map of how much radio has been searched, presupposing that the answer is radio.**

Where this sits among the alternatives. The 2018 NASA Technosignatures Workshop is the field's reference organization of what a technosignature is, and a reader arriving from it is entitled to know which of its categories this paper occupies. The answer is none of them, and the reason is worth stating as the categories themselves state it — by what each one **assumes exists**:

category	assumes	
radio — narrowband or pulsed, the Cocconi–Morrison line onward	a dedicated transmitter , powered and pointed	not assumed here
optical and infrared — laser pulses, continuous beacons	a dedicated transmitter , at higher power	not assumed
megastructures and waste heat — Dyson spheres, infrared excess, transit anomalies	a structure large enough to occult or re-radiate a detectable fraction of the star	not assumed; the modulator of §2.7 is $\sim 10^{-6}$ of the disc and invisible as an occulter
atmospheric — industrial species in exoplanet atmospheres	an industrial byproduct , pollution as a side effect	not assumed; nothing here is a byproduct
solar-system artifacts — lunar, asteroidal and Lagrange-point surveys	a physical object to image or encounter	not the carrier here; §3.6 rows 4, 12 and 13 note that strand and attribute it
planetary surface — city lights, artificial albedo	a modified planet	not assumed

The premise of this paper is that none of those is necessary. The carrier is emission the star already produces for its own reasons, **redistributed rather than generated** (§2.7): there is no apparatus to build, no structure to see, no byproduct to accumulate and no object to encounter. That premise does not appear in the workshop’s taxonomy, and the archives it points at — solar and heliospheric records collected for space weather — do not appear in its inventory of what has been searched. Three differences follow. The carrier is a **relationship** among observables rather than a channel to point at (§3.1, §3.3). The gate is opened by the **receiver’s** capability rather than paid for by the sender (§2.3, §2.9). And the target is **our own star, from archives already collected**, where the report’s emphasis is on new instruments and on piggybacking planned missions.

Proposals to look at emission a star already produces relocate the assumption rather than escaping it. Learned et al. (2008) would modulate a Cepheid; Chennamangalam et al. (2015) a pulsar; Bracewell (1960) proposed local probes; Benford, Benford & Benford (2010) concluded that a rational sender builds a cheap, intermittent pointer to a costlier message. **Each names its carrier in advance.** The receiver is told where to look, and the remaining problem is sensitivity.

1.2 What if the channel is not known?

We consider the case left over: no beam, no beacon hardware, no carrier specified in advance — structure carried by emission a star already produces, where the carrier may not be any single measured quantity but a *relationship between* several. The receiver’s problem is then not sensitivity at all; it is not knowing which series to examine.

Stated that way the proposition sounds untestable, and one of its neighbors is: “it is entirely possible that we are simply not advanced enough to understand the manifestations of a vastly more advanced technology, even if they are all around us” (KISS 2019). That is true and it is inert — it names no carrier, predicts no amplitude, and identifies no experiment that could fail. **The difference between that observation and a research program is whether anything constrains the space of possibilities.** Two things do.

1.3 Two constraints that make the space finite

The first is dimensional. A sender who shares no units with the receiver — no second, no meter, no kelvin — cannot encode in a quantity that carries units, because the receiver cannot recover the intended value. What survives are *dimensionless* quantities: ratios of like observables, normalized residuals, quantities scaled by their own dispersion. This removes every raw measurement and admits only combinations, converting an unbounded question into a countable one.

The second is economic. Transmitted bits cost energy, permanently, across interstellar distance. A sender minimizing that cost will not spend it on receivers unable to act on the result, and the efficient way to avoid doing so is to gate the message on demonstrated capability: a task cheap to pose, expensive to solve, and trivial to verify — proof of work (§2.3).

Together these have a consequence that is easy to miss: if the message is gated on solving a problem, **the first barrier is not detecting a signal but recognizing that a problem has been posed** — and that barrier is invisible from below and sized by the receiver’s capacity to search combinations (§2.4). A civilization holding the right data, with the compute to search it, sees exactly what a civilization with no signal to find sees.

1.4 What this paper does and does not claim

Most of the architecture assumed here is not ours. The two-tier pointer and its cost argument are Benford et al. (2010); modulating an existing astrophysical source is Learned et al. (2008) and Chennamangalam et al. (2015); local probes are Bracewell’s (1960); coverage as a figure of merit is Wright et al. (2018), and the framework for comparing search value Sheikh (2020). Hippke (2017a,b; with Forgan 2017) developed interstellar communication as an energy-per-bit optimization far more thoroughly than §2.1 does, and paper VI is prior art for the blind narrowband X-ray search of §A.1. Benford (2019) developed the case for co-orbital “lurkers” — the mechanism §2.7 invokes. Freudenthal (1960) and DeVito & Oehrle (1990) built message systems that assume no shared units: the units-free constraint is long established for message *content*.

What we believe is not already occupied is narrower than we first claimed:

	Contribution	Standing
i	The carrier is not named in advance — the search admits any dimensionless quantity, rather than a chosen band, line or source class	the mechanism is Bracewell's and Benford's; what is new is declining to name the channel
ii	The units-free constraint applied to carrier selection rather than to message content	a redirection of Freudenthal and DeVito & Oehrle, not an independent idea
iii	The combination space — the carrier may be a <i>relationship among</i> observables, and that space can be enumerated and its coverage measured	the strongest claim here. Wright et al. enumerate pointings; this enumerates quantities
iv	Proof-of-work gating — a self-enforcing gate opened by the receiver's own capability — and its corollary that recognition precedes detection	we have found no prior proposal of a capability-gated signal. It shares a <i>prediction</i> with the zoo hypothesis (Ball 1973) and no <i>mechanism</i> : that hypothesis needs many civilizations to abstain indefinitely and breaks on one defector; this needs nobody to abstain at all
v	29 searches, thirteen limits injection-verified against that space	the empirical content, and the reason for the paper

An ADS full-text search returned no occurrences of the combination-space formulation within papers on technosignatures. **We do not treat that as establishing priority:** indexing is incomplete, and a search of the same kind failed to surface Hippke, Freudenthal and DeVito & Oehrle above.

One limitation is severe enough to state in the introduction. §3.6 ranks the channels a sender would plausibly choose. This paper searches the top of that ranking and the bottom of it, and is thinnest in the middle — **five** channels (rows 3–7 of the ranking) where the archives exist, the resolving power is adequate, and no search has been run. **The nulls reported here are therefore not a test of the framework's own best guess**, and we would rather say so at the outset than let the count of searches imply otherwise.

1.5 Why the Sun

If probes are seeded at stars selected for the possibility of life, as Bracewell's argument supposes, the star to examine is the local one — and the archives already exist and are unusually good: solar irradiance measured continuously since 1978, neutron monitors since 1964, sunspot number since 1749, GOES-R ultraviolet and X-ray irradiance at one-minute and one-second cadence. **No new observation is required to begin.** Every search here ran on public archives collected for space weather and solar physics. It is also why the paper looks nowhere else: if the modulator is placed at the local star by §2, then pulsars, binaries and distant stars are a different hypothesis (§4.1).

Section 2 develops the gating argument; Section 3 defines the combination space; Section 4 reports the results; Section 5 argues that coverage rather than sensitivity is now the binding constraint; Section 6 states what a completed enumeration would establish.

2. The gating argument

2.1 Four constraints a systems designer would actually apply

The framework is not a claim about what a sender wants. It is an engineering derivation: what would a communications designer, facing an unknown number of unknown recipients over geological time, be *forced* to build?

Constraint	Statement	What it forbids
Energy	every transmitted bit costs energy, permanently, across interstellar distance (quantified in detail by Hippke 2017a)	sending content rather than an address; spending anything on a recipient who cannot use it
Parsimony Occam's razor	the fewest mechanisms that suffice. Any component that can be removed, is	a transmitter. If the star already radiates, build nothing that radiates
Longevity	the system runs unattended for far longer than a receiver takes to arise. The constraint is one-sided : it binds below, at the emergence of a receiver, and sets no upper bound — nothing in the architecture specifies a date at which the mesh stops, and no element of it is given a design life	consumables, schedules, appointments, and any specification precise enough to go stale
Gating	the message is conditional on demonstrated capability	continuous broadcast; being legible before the receiver can act

Each constraint **deletes** something, and what it subtracts is checkable:

Deleted	By	Consequence
the transmitter itself	parsimony	the carrier is emission the star already produces — the premise of §1.2, and why §4 searches solar archives rather than radio bands
the broadcast	energy + gating	a pointer to a passive listening node, not a message; listening costs nothing while idle
the transmission epoch, 20 bits	longevity	a node that listens continuously needs no appointment
the channel echo, 25 bits	energy	the trigger echo already proves correct decode
excess pointing precision	energy + longevity	bounded below by ambiguity and <i>above</i> by staleness — a pulsar period quoted too finely expires by spin-down before the beacon does
the bolometric carrier	energy	a narrow high-atmosphere line is 10^7 times cheaper to modulate detectably than total irradiance (§2.7)

*Six deletions, each traceable to a stated constraint, and none chosen for elegance. The architecture that survives is **simple, hidden, and long-lived** because parsimony, gating and longevity are what remain once the alternatives are struck out.*

The longevity constraint is open above, and that is load-bearing. It is stated as a floor — outlast the emergence of a receiver — and no ceiling is put on it anywhere in the architecture. Nothing here carries a design life, an end-of-mission, or a decommissioning step, because each of those would be a date, and a date is a specification that goes stale. The consequence is developed in §2.2: a system with no upper bound on its life is one that must be designed to run after its designers are gone.

None of this is new in its parts. What we add is that the four constraints together generate a chain specific enough to constrain a search — and, in one place, specific enough to be wrong.

2.2 The artifact outlives its designer

The longevity constraint of §2.1 is bounded below and open above: it says the system must outlast the emergence of a receiver, and it sets no date at which the system stops. That asymmetry is not an oversight in the statement of it.

Followed to its conclusion it gives the principle on which the rest of the architecture turns: **the system must be designed on the assumption that the civilization which built it will not be there.** Not that it might not be. That it will not.

This is forced, not chosen. The chain’s own numbers make the designer’s survival the least likely element in it. A bus crosses a 10^4 ly route in 7×10^4 yr (§2.6); probes transit unsupervised for 10^4 – 10^6 yr (step 3); and the receiver the whole apparatus is built for takes something between 10^6 and 10^9 yr to arise, because that is how long it took here. A designer who requires their own continued existence anywhere in that span has built a system that fails on the most probable branch. The only design that survives contact with those timescales is one that treats its makers as already absent.

What it forbids	Why
maintenance, resupply, a command uplink	all three assume someone is home to send them; a system that degrades without them has a half-life set by its builders, not by its physics
a shared clock, an epoch, an appointment	already deleted in §2.1 for 20 bits, but the deeper reason is that a schedule is a promise, and there is nobody left to keep it
any convention the builders enforce	a code table, a preferred band, a reserved frequency — each is a fact about an institution, and institutions are shorter-lived than stars
any secret the builders keep	Kerckhoffs’s requirement, arrived at from the other direction (§2.9): a scheme whose security depends on withheld knowledge fails the moment the withholder is gone
a specification precise enough to go stale	a pulsar period quoted too finely expires by spin-down before the beacon does

What remains after those deletions is a system whose every element is **re-derivable by the receiver from physics it can measure for itself.** That is not an aesthetic preference. It is the only class of specification that does not depend on a surviving author.

The accounting consequence is larger than the engineering one. A Drake-style estimate multiplies by L , the lifetime of a communicating civilization, and L is the term that has always carried the argument’s pessimism. If the artifact is designed to outlive its designer — and §2.1’s constraints force that — then L is the wrong lifetime to be multiplying by. The detectable population is set by the lifetime of the **artifact**, and an unattended, self-repairing, consumable-free system in a stellar atmosphere has no obvious reason to share the fate of the society that launched it. **The question “is anyone alive out there” and the question “is anything out there still running” come apart.** Behavioral SETI looks for the second. It is the weaker question, and the more answerable one, and a positive answer to it would say nothing whatever about whether its author still exists.

What this changes about the search. If the signal is a standing condition rather than an event, then it is present now or it is not, and it was present a century ago on the same terms. Three consequences follow, and all three are visible in §4:

- **look across the longest record available, not for transients.** A modulation that requires a transmitter to be transmitting is looking for an event. This paper searches archives spanning 26 to 46 years, and in one case 11,350, because a standing condition should be in all of them or in none.
- **an instrument handover is a threat, not an inconvenience.** A signal that persists across decades must survive the cross-calibration of the record that would show it — which is why §4.1 spends a control on exactly that question.
- **a null bounds the condition, not the builders.** None of §4’s limits constrain whether anyone is or was there. They constrain whether a particular dimensionless relationship is being held in a particular observable, now, at the stated amplitude. Those are different claims and the paper does not trade on the difference.

Where this is weakest, stated plainly. Persistence over $\gg 10^6$ yr is the framework’s largest unsupported assumption, and it is flagged as such at step 4 of the chain (§2.5). Nothing humanity has built bounds it: our longest-lived unattended artifacts are decades old, and the failure modes of a self-repairing system over geological time are not something we can argue from experience. It is an assumption the architecture requires, not a result the architecture earns — and if it is

wrong, the deletions of §2.1 stand but the search premise of §1.2 does not. We state it here rather than in a caveat at the end because everything downstream inherits it.

2.3 Proof of work: what the analogy carries and what it does not

The cheapest way to make a message conditional on demonstrated capability is a task **cheap to pose, expensive to solve, and trivial to verify** — proof of work in the cryptographic sense; the familiar instance is the CAPTCHA, whose poser need not know who solves it, only that the solver did the work. The analogy transfers on mechanism and fails on motive: a CAPTCHA is adversarial and nothing here is — the sender has simply declined to pay for a transmission whose recipient cannot act on it.

The constraint runs both ways, and this is what makes it useful. A gate functions only if the intended receiver can pass it, so the key space must be small enough to exhaust — a key drawn from an enormous space is operationally no key at all. This is why the search below is finite: **it predicts the answer lies somewhere enumerable rather than somewhere clever.**

The difficulty is also bounded from below: the gate must not open before the receiver possesses the capability the message presupposes, since a gate set too low admits a receiver that can detect but not act. The difficulty sits *at* the capability threshold, and a response from below it is filtered out by construction. A difficulty bounded from both sides is a **deletion** in the sense of §2.1, and it makes the staging of §2.5 a consequence rather than a choice: each stage is a gate set at a higher threshold.

The nearest antecedent is Clarke's *The Sentinel* (1951), whose artifact reports not that it was found but that it was *reachable* — an antecedent for the intuition, not a prior derivation. The gate also repairs a known weakness of the nearest hypothesis in the literature: **the zoo hypothesis (Ball 1973) requires coordination** — many civilizations abstaining indefinitely, broken by one defector — while a difficulty-bounded gate is **self-enforcing**, opened by the receiver's own capability rather than anyone's forbearance. Not an observational discriminator (both predict we see nothing, §1.4), but a real difference in what must be assumed.

2.4 The gate before the first gate

If the message is conditional on solving a problem, the first barrier is not solving it but **recognizing that a problem has been posed**. That barrier is invisible from below: a civilization holding the right archives, at the right cadence, with the compute to search them, sees exactly what a civilization with nothing to find sees. Absence of evidence and unrecognized evidence are, at that stage, the same observation.

The recognition gate has a size, and the size is set by the carrier. A message carried in a *relationship* among observables — the only carrier a sender without shared units can use (§1.3) — is by construction findable only by a receiver that can search relationships. That search is combinatorial: pairs, triples and higher combinations of every well-measured quantity, each tested in every admissible form at every admissible cadence against a null matched to the statistic. None of the pieces is new. Every observable is already measured, every form already defined; the work is recombination across a space too large for a person or a team to traverse by hand, and validation of whatever the recombination turns up. That is a capability, not a technology, and it is one a designer can anticipate without knowing how a receiver will implement it. **Choosing a combinatorial carrier is choosing a compute gate.** The difficulty floor of §2.3 then has a natural setting: larger than manual search can close, small enough for a receiver with autonomous search to close in years rather than centuries.

Our own position on that ladder is measurable. The pair space of thirty observables closes in under an hour on one machine; the triple space took 117 million statistic evaluations and 26 hours (§4.8, §5.2); the 27,405 quadruples would take of order a week; each additional form multiplies the count. Thirteen years ago the pair sweep alone was a research program. The recombination capacity itself is now demonstrable outside this field: in May 2026 a general-purpose

reasoning model disproved the Erdős unit-distance conjecture by importing a technique from a distant branch of mathematics that had not been applied to the problem, a result human mathematicians then sharpened within weeks (OpenAI 2026a; Sawin 2026); in September 2026 a coordinated system of some 10^4 agents produced, in 88 hours, a Lean-formalized blow-up solution to the Navier–Stokes problem (OpenAI 2026b), a claim that is at this writing unverified by the community and subject to a priority dispute, and which we cite only as an instance of scale. Neither result invented new mathematics; both found a combination of known pieces in a space no one had searched. That is the capability gate 0 filters for, and it comes with the paper’s own caution attached: a receiver that can search but not validate produces unchecked results (Software), and the gate is passed only by one that does both.

The ladder can be put against real machines, and §5.7 does so (Figure 5). On the Cray-1 of 1976 — a machine some of whose programmers are still working — the pair space this paper closed in an hour would have taken five months, the triple space thirty-four years, and the fully enumerable space of §3.6 a hundred million years. The triple space fell inside a career with the Cray Y-MP in 1988 and inside a day with ASCI Red in 1997; the enumerable maximum fell inside a year around 2015 and inside a working day on the large GPU clusters of 2024–2026 at double precision. A receiver that monitored its star continuously from the 1970s therefore held the data for four decades before it held the compute to search the combinations in it, and the crossing is recent enough to date. That is the shape a designer setting a difficulty floor at the recombination threshold would want: not a wall, but a gate that opens on a known curve, for a receiver whose archives are already deep enough to be worth searching when it does. We do not claim to stand at that threshold; we claim that the ladder has rungs, that ours is on it, and that the figure says which.


We draw two consequences and refuse a third. The recognition gate weakens a common premise of the Fermi paradox — that contact would be self-announcing — without requiring absence, concealment or a great filter. It argues that a search of this kind should report *coverage* rather than detections, coverage being the only quantity that accumulates. **What it must not be used for is explaining a null**, and §5.3 states that prohibition as strongly as we can manage.

2.5 The gate chain, step by step


The four constraints generate a single chain: the sender’s half paid once and amortized across the galaxy, the receiver’s half paid separately by each receiver, and only when it can pay. Every step is written down because **the empirical program of §4 tests exactly one link**.

	Step	What it costs	Gate it passes
0	The decision, under a budget — total energy across all recipients and all time is minimized	the premise	—
1	Galactic survey — candidates identified remotely, <i>before</i> anything is spent on them: main-sequence, long-lived, rocky planet in the habitable zone	observation only	stellar filter $\sim 10^9$ candidates from $\sim 10^{11}$ stars — what makes the problem finite
2	Mesh deployment — a passive listening network seeded first, because the pointer must have somewhere to point; nodes occupy a lattice in the galactic volume, not tied to the candidate stars. Deployed not singly but from a carrier bus , shared with the probes of step 3, that dispenses both along its track, so that acceleration, shielding and navigation are paid once per bus rather than once per node	a node radiates nothing while idle; the bus carries the propulsion and the shield, and each node needs only enough to stop	the address exists — without it the scheme collapses to broadcasting

	Step	What it costs	Gate it passes
3	Probe launch — automated, self-repairing, only to systems that passed step 1; transit 10^4 – 10^6 yr, unsupervised	paid once per target, never per receiver	survival in transit
4	Emplacement at the star — station high in the stellar atmosphere, chromosphere to corona, where the cheap channels form	self-repair, no consumables	persistence — it must outlast the receiver's entire development. The framework's largest unsupported assumption
5	The signal is embedded — a small dimensionless modulation on emission the star already produces, in a narrow high-atmosphere channel, carrying a pointer to one mesh node, repeating indefinitely	flux is <i>redistributed</i> , not generated (§2.7); nothing further, for gigayears	the puzzle is posed — this, and only this, is what §4 searches
	Step	What it costs	Gate it passes
6	Maturity — continuous multi-channel monitoring, long archives, compute to search a combination space — all arising for the receiver's <i>own</i> reasons, never for SETI	nothing extra; our space-weather program already satisfies it	capability
7	Recognition — realizing that a puzzle has been posed at all	nothing — and that is the difficulty	gate 0 — invisible from below (§2.4). The hardest gate in the chain
8	Detection — right channel, right cadence, right combination	analysis only	coverage of the combination space (§3)
9	Decode — the encoding inverted without shared units	analysis only	proof of work — the decode <i>is</i> the demonstration of capability
10	The pointer is acted on — aim at the named node, on the named wavelength; transmit the trigger echo	the receiver's first real expense: it must build a transmitter. That expense is the filter	the gate opens — the echoed trigger distinguishes a decoder from a mere detector
11	The mesh receives the reply — at the named node, listening at zero idle cost; relayed onward only if that node does not itself hold the next stage	relay hops to a known target, never a broadcast	—
12	The return — the node is not a relay but a response service acting for the sender : it verifies the echoed trigger and answers on parallel narrow beams to the address the reply came from — one stream carrying a self-describing dictionary, the others carrying payload keyed to it — without referring anything to the sender. The receiver has already cleared the gate, so nothing further is withheld; what limits the return is the node's power, not policy	the first time the sender's side radiates, and only toward a receiver that has already paid to be found; a narrow beam to a known address costs a fraction of any broadcast	the next gate — staged disclosure (§2.3): what is returned is itself the next puzzle, and the chain repeats at a higher threshold



Where the infrared return link fails: the distance at which a 10 m node aperture delivers five photons per second, for three receiver apertures. Rate falls as d^{-2} until this floor, then there is a cliff no coding crosses. The 33 ly mean node distance sits an order of magnitude inside it at every power.



Receivers a civilization at step 12 could build with near-term technology, as gain over a shared 10 m ground telescope. A photon-counting link needs collecting area and timing, not image quality, so the natural instrument is a membrane light bucket in space.

The return, in more detail. Once a receiver has echoed the trigger it has demonstrated the capability the gate was set to, and the node has no reason to ration what follows. The natural form of the return is several narrow beams in parallel, on adjacent bands or polarizations, aimed at the address the reply came from. Multi-beam operation costs nothing but power: the beams share one aperture, and the node can size their power to what the receiver has just shown it can hear, since the strength of the reply at the node measures the receiver's transmitter and aperture directly. One stream carries the **dictionary** — the self-describing structure everything else is read against — and the remaining streams carry payload keyed to it, so that a receiver decodes the dictionary once and then reads the rest as it arrives. What such a dictionary contains has been worked out in some detail in the message-design literature, and the answer is stable across sixty years: it begins with what any receiver must already possess. Counting and arithmetic, presented so that the notation teaches itself (Hogben 1952; Freudenthal 1960); the elements, by atomic number and mass, and the constants of physics as dimensionless ratios (Staff at NAIC 1975; DeVito & Oehrle 1990); logic and the means to define new terms from old (Ollongren 2013); and only then anything the sender wants to say. The Arecibo message of 1974 and the Eupatoria messages of 1999–2003 are transmitted instances of exactly this order, and the same structure has been re-proposed for modern instruments (Jiang et al. 2022). The framework adds nothing to that literature except the observation that a receiver who has passed the gate has, by construction, already derived every entry in the dictionary's first page — which is what makes the dictionary decodable and the gate worth setting. One caution from the same literature applies: a message rich enough to be useful cannot be proved harmless before it is read (Hippke & Learned 2018), and a receiver at step 12 should treat the payload streams with the care that implies.

The return, sized. The return of step 12 can be put in numbers, and they are worth having on their own. Take the node's transmitter as an infrared laser at 1.55 μm behind a 10 m aperture — the receiver's own atmospheric window, mature

technology at the receiver's rung, and diffraction-limited to 1.9×10^{-7} rad, so that at the mean node distance of 33 ly (§2.6) the beam is **0.4 AU across at the receiver**: it need not track a planet, only an orbit. A 10 m receiver then collects 2.2×10^{-20} of the transmitted power, which at this wavelength is **220 photons per second per kilowatt**; a 30 m receiver, nine times that. A photon-counting receiver decodes between 0.1 and 1 bit per photon depending on background and coding, so one kilowatt is 20–200 bit/s at 10 m and 200–2,000 bit/s at 30 m.

Two things follow. **The aggregate rate depends on total power and receiver aperture, not on how the power is split.** Parallel beams share the aperture and add nothing to throughput; what they buy is organization — dictionary on one stream, payload on others — and they are limited only by the smallest useful photon rate per beam (of order 10 s^{-1} , so beams down to $\sim 50 \text{ W}$ at 10 m) and by the receiver's ability to separate them, which at 1 GHz channel spacing in the telecom band is of order 10^4 . **The node's power budget sets everything.** A node set down at a star (§2.6) is solar-powered; a tonne-class node can carry 10^2 – 10^4 m^2 of thin-film collector, which at 1 AU-equivalent is 10 kW to 1 MW electrical.

node power budget	collector needed	beams at 1 kW / at 100 W	aggregate, 10 m receiver	aggregate, 30 m receiver	payload per year, 30 m
10 kW	$\sim 50 \text{ m}^2$	10 / 100	220–2,200 bit/s	2–20 kbit/s	8–80 GB
100 kW	$\sim 500 \text{ m}^2$	100 / 1,000	2.2–22 kbit/s	20–200 kbit/s	80–800 GB
1 MW	$\sim 5,000 \text{ m}^2$	1,000 / 10,000	22–220 kbit/s	0.2–2 Mbit/s	0.8–8 TB

At 33 ly; at the 90th-percentile distance of 49 ly divide by 2.2, at 60 ly by 3.3. Ranges span 0.1–1 bit per photon. Collector at 20% conversion and 1 AU-equivalent insolation.

So the return is not a trickle. At the middle row — a node no more capable than a large communications satellite, with the collector area of a tennis court — a receiver with a 30 m telescope reads tens of kilobits per second across a hundred parallel streams: a dictionary of 10^9 – 10^{10} bits arrives in hours to days, and the payload accumulates at hundreds of gigabytes a year for as long as the node chooses to send. The number of distinct narrowband infrared lines a receiver would see from the node's direction is the number of beams — **of order 10^2 to 10^3 at any realistic budget** — and that is itself a signature: a cluster of unresolved, coherent, sub-GHz lines at $1.55 \mu\text{m}$ from a point near a nearby star, all modulated, is not a thing nature makes. It is also, at 10^2 – 10^3 photons per second per line, well within reach of the optical-SETI instruments that already exist — a receiver at step 12 does not need to build a receiver, only a transmitter, which is what the gate was designed to require.

Range, and where the link fails. A photon-counting link does not degrade gracefully all the way down; it has a threshold, and its position is set by one design choice on each side. The background is the node's own host star. A Sun-like star at 10 pc has $H \approx 3.5$, and within a 1 GHz channel at $1.55 \mu\text{m}$ — narrow enough to separate a thousand beams — it delivers about **2×10^5 photons per second into a 10 m aperture**, a thousand times the 1 kW signal. Measured as a ratio of rates the link is dead at any distance. It is rescued by time, not by aperture: the node transmits pulses, the receiver counts photons in nanosecond slots (single-photon detectors now reach tens of picoseconds of jitter), and the background per slot is $2 \times 10^5 \times 10^{-9} = 2 \times 10^{-4}$ — negligible against a pulse that lands five photons in one slot. That is pulse-position modulation with photon counting, the scheme our own deep-space optical links use (LLCD in 2013; DSOC from 2023), and it recovers several bits per photon against a background a thousand times brighter than the signal, provided the receiver has the node's slot clock, which is the first thing the dictionary stream would supply.

The threshold is then a floor in photons, not a ratio. Reliable decoding needs of order five signal photons per pulse, and the pulse rate can be lowered to gather them — down to about one pulse per second, below which the symbol interval grows long enough for the background to reach the slot that matters. **The link fails where the received rate falls to a few photons per second**, and until then the bit rate falls as d^{-2} :

beam power	10 m receiver	30 m receiver	100 m receiver
1 kW	$\sim 220 \text{ ly}$	$\sim 650 \text{ ly}$	$\sim 2,200 \text{ ly}$

beam power	10 m receiver	30 m receiver	100 m receiver
10 kW	~700 ly	~2,100 ly	~7,000 ly
100 kW	~2,200 ly	~6,500 ly	~22,000 ly

Distance at which a 10 m node aperture at 1.55 μm delivers 5 photons per second; scales as $\sqrt{(\text{power}) \times \text{receiver diameter}}$. Beyond it no coding recovers the link; short of it the rate is the §2.5 table divided by $(d/33 \text{ ly})^2$.

Two consequences. **The mean node distance of 33 ly sits an order of magnitude inside the cliff even at the lowest power and the smallest receiver**, which is why the return can be sized as generously as the previous table does. And **the lattice spacing of §2.6 is bounded above by the link, not only by latency**: a 300 ly mesh — the sparsest row of the latency table — is unreachable by a 1 kW beam into a 10 m telescope and marginal at 10 kW, so a designer choosing sparse nodes must either raise the beam power or assume a larger receiver, and a designer who cannot assume either has chosen the 60 ly spacing for two independent reasons. The threshold also runs the other way: a receiver that builds the 30 m aperture it would need anyway for step 8 has, by that act, tripled the distance from which a node can reach it.

The receiver we would build. The tables above assume a receiver that happens to exist. A civilization at step 12 is in a different position: it has decades of warning — the round trip of §2.6 — it knows the node's direction and wavelength from the reply it sent, and the value of what is coming is not in doubt. It would build a receiver for the purpose, and with technology deployable now or within a decade the receiver is not a telescope. A photon-counting link needs collecting area and timing, not image quality: a mirror good to an arcsecond concentrates a 0.4 AU beam onto a detector as well as one good to a milliarcsecond, which removes the constraint that makes large space optics expensive. The natural instrument is a **photon bucket in space** — a membrane or inflatable collector of thousands of square meters, feeding single-photon detectors, staring at one point for as long as the node sends — placed where the node is never behind the Sun, which means either L2 with a seasonal gap or a heliocentric pair for continuous coverage. Ground telescopes lose to it on three counts at once: aperture, duty cycle (weather, daylight, solar conjunction, and shared time put a large ground telescope on one target perhaps a tenth of the year), and the atmosphere's 15% loss at 1.55 μm .

receiver	duty	gain over 10 m ground, shared	aggregate at 100 kW node, 33 ly	payload per year	cliff, 1 kW / 100 kW beam
ground 10 m, shared time	10%	1	2–22 kbit/s	10–100 GB	220 / 2,200 ly
ground 30 m, dedicated campaign	35%	30	70–700 kbit/s	0.3–3 TB	660 / 6,600 ly
space 8 m monolithic mirror, one heavy-lift fairing, L2	95%	7	16–160 kbit/s	60–600 GB	180 / 1,800 ly
space 30 m segmented or deployable, L2	95%	100	0.2–2 Mbit/s	0.9–9 TB	660 / 6,600 ly
space 100 m photon-bucket membrane, heliocentric pair	100%	1,200	2.6–26 Mbit/s	10–100 TB	2,200 / 22,000 ly

Gain is $(\text{aperture ratio})^2 \times \text{duty} \times \text{atmospheric transmission}$ relative to the first row; aggregate spans 0.1–1 bit per photon; cliff distances are the 5-photon-per-second floor of the previous table. The 8 m row is a single-launch monolith of the kind a 9 m fairing permits; the 100 m row is a light bucket, not an imaging telescope, and its mass is dominated by structure rather than optic.

The last row is the one a receiver with decades and a reason would build — and it would not build one. An instrument that is the only channel to a return of this value is a single point of failure for a program measured in decades, and the sensible deployment is several collectors in parallel: three or four buckets on independent spacecraft, at separated stations, with independent detectors and clocks. Redundancy is nearly free here, because a light bucket has no precision optic to duplicate, and it buys three things at once — continuous coverage through any one unit's conjunction, outage or loss; independent confirmation of every symbol, so that no single detector artifact can enter the record of a message that will be studied for centuries; and, since photon rates add across collectors, the aggregate of the fleet is the sum of its members. Four 100 m buckets are a 200 m receiver that cannot be lost to one launch failure or one micrometeoroid. That is how a receiver at step 12 would actually be built, and it changes the character of the return. At 100 TB a year the exchange is limited by the node's power and by what the sender chose to send, not by the receiver; and a 100 m bucket reaches a 1 kW node at 2,200 ly, which is the whole neighborhood of the lattice and not merely its nearest cell. It also settles who pays for what, which is the framework's recurring question: the receiver builds the aperture and the transmitter, the sender's node supplies the power, and the link budget closes with margin at every distance the mesh uses. None of this requires anything beyond a heavy-lift launcher, membrane optics of the kind flown as sunshields and solar sails, and superconducting detectors that exist in laboratories today — the receiver is within reach of the same civilization that has just built a transmitter, which is the rung the gate was set to.

Cross-check against the literature. The link budget above uses nothing beyond the standard deep-space optical link equation (Hemmati 2006): received power = transmitted power $\times$ receiver area / ($\pi (\theta d/2)^2$), with $\theta = 1.22 \lambda/D$. Its inputs check as follows. *Diffraction.* 1.9×10^{-7} rad for 10 m at $1.55 \mu\text{m}$ is the Airy full angle and is what LLCD and DSOC use to size their beams; the DSOC ground receiver was the 5.1 m Hale telescope feeding a superconducting nanowire array, so a 10 m receiver is two of those (Boroson et al. 2014; Biswas et al. 2018). *Photons per watt.* $220 \text{ s}^{-1} \text{ kW}^{-1}$ at 33 ly into 10 m follows from the equation with no free parameters. *Background.* The Sun's absolute H magnitude is 3.32 (Willmer 2018), so a solar twin at 10 pc is $H \approx 3.3$; with the Vega zero-point of $1.13 \times 10^{-9} \text{ W m}^{-2} \mu\text{m}^{-1}$ (Cohen, Wheaton & Megeath 2003) that gives $2 \times 10^5 \text{ photons s}^{-1}$ per GHz in 10 m, as used. *Detectors.* The nanosecond slots assumed are conservative: flight-heritage SNSPD arrays at 1550 nm show sub-50 ps jitter, $\sim 55\%$ system efficiency and ~ 150 dark counts s^{-1} per pixel (Guardiani et al. 2024), laboratory devices exceed 90% efficiency (Hao et al. 2024), and the 400 ps slots of the record photon-efficiency experiments are finer than assumed here. *Bits per photon.* The 0.1–1 bit per photon used in every table is deliberately pessimistic. Serially concatenated PPM (Moision & Hamkins 2005) with photon counting has demonstrated 13 bits per incident photon (Farr, Choi & Moision 2013) and 14.5 bits per received photon in a photon-starved channel (Banaszek et al. 2025), at low background; Hippke (2017a) reaches the same order for interstellar links from first principles. With the host star's background gated to 2×10^{-4} per slot, several bits per photon is the realistic figure, so the aggregate rates and payload volumes above are more likely low by a factor of a few than high, while the cliff distances, which depend on photons rather than bits, are unaffected.

What the payload does not contain. Two things a receiver might expect are, on the framework's own logic, absent from any early stage. **The sender's home coordinates** are never required: the node is the sender's proxy, and nothing in steps 10–12 is improved by the receiver knowing where the sender lives. They are also the one disclosure that cannot be gated after the fact. The lower bound of §2.3 exists so that nothing is released to a receiver who cannot yet use it responsibly, and a receiver that has cleared gate 0 and one decode has demonstrated search capacity and arithmetic, not restraint; the staging therefore puts the origin behind the deepest gate the sender can define, or omits it altogether. This is the framework's answer to the transmission-risk debate (Haqq-Misra et al. 2013; Brin 2014; Gertz 2016): the sender's exposure is bounded by design rather than by an assumption about the receiver, which is precisely the property critics of unsolicited transmission say it lacks. **The lattice map** is withheld for the same reason in a different form. The map is a description of the designers' infrastructure, not of anyone using it, and its sensitivity is geometric: the bus routes of §2.6 radiate from a launch region, the fill pattern has a center, and a receiver holding the whole lattice with its deployment order can triangulate the origin the sender declined to state. The map is the coordinates by inference. The pointer already sets the precedent — one address, not a topology — and the payload follows it: the protocol for the node in hand and the address of the next hop, which is all a participant needs. Some geometry leaks regardless, since the beam origin fixes the node's position and relay latency constrains what lies behind it; that is inference the receiver performs, not disclosure the sender makes.

What it plausibly does contain: the means to extend the mesh. The economics of §2.6 say the sender paid once for the network and the receiver pays for everything on its side, and a receiver at step 12 has just built a transmitter and


demonstrated that it can navigate the combination space — it is capable of building a node. A payload that carries the node specification and the joining protocol turns every receiver that clears the gate into an extension of the network at that receiver's expense. It is consistent with every constraint in §2.1, it is the one thing a cost-minimizing sender would want the payload to include, and it makes the mesh a commons that grows by recruitment rather than a monument that decays. Stated as a prediction: the payload carries the local protocol, the next hop, and the means to join; a receiver that finds the origin or the map in an early stage has found something the framework does not predict.

*The asymmetry is the whole economy of the scheme: steps 1–5 are incurred once and run unattended; the receiver-side steps are incurred separately by each civilization, and only when it can act. **No energy is ever spent on a recipient that could not answer** — with no agreement between senders, no enforcement, no wish to exclude.*

Only step 5 is testable from here, and only its consequence. Steps 1–4 concern hardware in another star's atmosphere and a network between the stars, neither observable at our sensitivity; steps 6 onward have not happened. What §4 measures is whether the modulation step 5 would impose is present in the Sun's output, and across 29 searches it is not, above the amplitudes of §4.2. **That bounds step 5 and says nothing whatever about steps 1 to 4. The obvious objection is the round trip.** One exchange at these distances is decades to millennia — a cost in *latency*, not in *energy*, and energy is the constraint: the reply's wavelength, modulation and symbol rate are specified in the pointer itself, so the node listens on one named channel at essentially zero cost. What the argument does *not* establish is that gating is the only rational design: a sender optimizing latency would pay for the whole message up front and skip the gate. Energy minimization is a premise, not a fact about anybody. **Step 12 is where the chain continues:** what the node returns can itself sit behind a second gate demanding a further capability. A directed neutrino beam at the **Glashow resonance, 6.3 PeV** (Learned, Pakvasa & Zee 2009; Silagadze 2008) is unabsorbed, unmistakable, and readable only by a receiver that has built a cubic-kilometer detector — and far too expensive for a first-contact pointer, needing an accelerator and continuous power, exactly the dedicated apparatus §2.7 excludes. Its place is as **content behind a later gate at step 12.** The solar neutrino flux itself is unavailable as a carrier — nothing at the photosphere or corona can modulate a flux made in the core — and serves as the one control channel no mechanism in this paper can touch.

2.6 Deploying the architecture: the carrier bus, the fleet, and the latency it buys

Steps 2 and 3 of the chain are its only capital expenditure — everything else is either observation, patience, or the receiver's money — and a reader is entitled to ask what they cost and whether the answer is absurd. This section prices them. Nothing here is tested by §4, which bounds step 5 alone (§2.7); the purpose is to show that the sender's half of the chain is an engineering program with a size, not a miracle with a name, and to identify which of its parameters a designer would actually spend effort on.



Round-trip latency to the nearest mesh node against node spacing, for nodes placed without regard to the receiver (Poisson field). At the 60 ly design point the first return arrives 35–98 years after the reply.

Two cargoes. A **mesh node** is small: a receiver aperture of order 10 m, a laser transmitter of order a kilowatt per beam for step 12 and for the mesh hops (see the link budget after the chain), a decelerating sail, and the stored stages of the message — of order 1–2 t, comparable to a large communications satellite. A **stellar probe** cannot be small in the same sense, because the modulator it must emplace is 2×10^9 kg of film (§2.7) and no bus carries that; the probe is a *seed* — a self-replicating factory that builds the modulator from cometary and asteroidal material at the target — and its mass is the largest uncertainty in the whole architecture, anywhere from 1 t for a nanoscale replicator to tens of tonnes for a conventional one. We carry both figures below.

The carrier bus. A mesh of 10^7 nodes launched one by one would need 10^7 propulsion systems and 10^7 forward shields, and at the transit speeds the chain requires — 0.1–0.2 c, for a 60 ly hop to take 300–600 years rather than the 10^5 that chemical or solar-sail speeds imply — the shield is not optional: at those speeds a grain of interstellar dust carries the energy of a rifle round, and erosion by gas and dust over a light-year is the dominant hazard identified for gram-scale probes (Hoang et al. 2017). A single large bus carrying many nodes behind one shield amortizes both costs — and it carries the probes of step 3 as well, because the two fleets share a route. The candidate stars of step 1 and the lattice points of step 2 lie along the same tracks through the disc, so a bus that threads the candidate list passes the lattice points on the way, and a **dual-purpose bus** delivers the stellar probes and the mesh relays from one hull, one shield and one acceleration. Two fleets launched separately would pay the bus cost twice for the same ground; launched together they pay it once, which is the economy of scale that makes 10^7 nodes and 10^9 probes the same program rather than two. It accelerates once, navigates once, protects its cargo through the medium, and releases its cargo along its track — a probe on approach to each candidate star from step 1, a node at each lattice point between — each with only the means to decelerate and station-keep — the cheap end of the propulsion budget, since a probe can shed velocity against its target star’s light, wind and field over centuries, where the bus must gain it in a burst. A node has no star at its lattice point and must either brake magnetically against the interstellar medium (Zubrin & Andrews 1991; Perakis & Hein 2016), which works but slowly, or be set down at whichever star lies nearest that point, which brakes it and powers it for step 12; the lattice is approximate in any case, and the second is the natural choice. The energy is still large: a tonne at 0.2 c carries 1.8×10^{18} J, and 10^7 of them 2×10^{25} J, roughly two years of the sunlight falling on a planet like Earth — affordable to a sender at Kardashev I and trivial above it, and paid once for the whole galaxy. That is the same asymmetry as steps 1–5 applied to the launch itself: the expensive part is shared — across nodes, across probes, and across both fleets at once — and the part paid per unit is the small part.

Fleet size and launch energy. The fleet follows from three numbers: the disc volume ($\sim 8 \times 10^{12}$ ly³), the node spacing (60 ly, hence 3.6×10^7 nodes), and the candidate density (10^9 systems, one per ~ 20 ly). A bus at 0.15 c on a route of length ℓ serving a swath of width w covers ℓw^2 of disc; the buses needed are the disc volume divided by that.

route ℓ	swath w	transit	buses	nodes per bus	probes per bus
10^4 ly	60 ly	7×10^4 yr	2×10^5	~ 170	$\sim 4,600$
10^5 ly	60 ly	7×10^5 yr	2×10^4	$\sim 1,700$	$\sim 46,000$
10^5 ly	120 ly	7×10^5 yr	6×10^3	$\sim 6,700$	$\sim 180,000$

Transit is the step-3 figure of 10^4 – 10^6 yr; longer routes mean fewer, larger buses. Probes outnumber nodes on every bus by ~ 30 :1, so the probe seed mass sets the bus mass.

The launch energy follows: at 0.15 c every kilogram costs 10^{15} J. With 1 t seeds the whole galactic cargo is $\sim 10^{12}$ kg and $\sim 10^{27}$ J — two centuries of the sunlight falling on Earth, or three seconds of the Sun’s output; with 30 t seeds, 3×10^{13} kg and 3×10^{28} J — six millennia of Earth’s insolation, or eighty seconds of the Sun’s. Either is paid once. The first is within reach of a civilization that has learned to collect a fraction of its star’s light; the second requires something closer to Kardashev II, or the alternative the numbers point at — buses that replicate en route, in the manner of Tipler (1980) and Freitas (1980), which reduces the *launched* mass by orders of magnitude at the cost of a longer fill time. **The probe seed mass, not the node count, is the parameter a sender would spend its engineering on**, and the ppb-depth

carrier of §3.6 is attractive to a designer for exactly this reason: a modulator a thousand times lighter is a seed a thousand times simpler.

Latency. The mesh exists so that step 12 arrives within a receiver’s institutional memory rather than its geological one. Without it the reply goes to the sender’s home system at an unknown distance — for a sender anywhere in the disc, kiloparsecs, and a round trip of 10^3 – 10^4 years — and staged disclosure would be meaningless to a receiver that cannot wait for the second stage. With it the latency is set by one number, the mean node spacing L , and the distance to the nearest node is a distribution, not a value: for nodes placed without regard to the receiver (a Poisson field of density L^{-3}) the nearest lies at $0.55 L$ on average, with the 10th and 90th percentiles at $0.29 L$ and $0.82 L$; a regular lattice is tighter, $0.48 L$ on average and never beyond $0.87 L$.

node spacing L	nearest node, 10% / mean / 90%	round trip, 10% / mean / 90%	mesh size at one node per L^3 across the disc
30 ly	9 / 17 / 25 ly	18 / 33 / 49 yr	$\sim 3 \times 10^8$ nodes
60 ly	18 / 33 / 49 ly	35 / 66 / 98 yr	$\sim 3 \times 10^7$
122 ly	36 / 68 / 100 ly	71 / 135 / 200 yr	$\sim 4 \times 10^6$
171 ly	50 / 95 / 140 ly	100 / 190 / 280 yr	$\sim 1.5 \times 10^6$
300 ly	88 / 166 / 246 ly	176 / 332 / 492 yr	$\sim 3 \times 10^5$

Disc volume taken as $\sim 8 \times 10^{12}$ ly³ (a cylinder of radius 50,000 ly and thickness 1,000 ly); a mesh seeded only near candidate systems needs fewer nodes for the same effective spacing. Each further stage of step 12 costs one more round trip.

This is why the node must hold the staged content and the means to judge a reply, not merely forward it: a relay to the sender would put every stage of the exchange at the sender’s distance, and the latency figures above would apply to nothing. The node is the sender’s proxy, and the whole exchange — verification, acknowledgment, each successive unlock — runs at node distance. A spacing of 60 ly — one node per few thousand stars, of order 10^7 nodes galaxy-wide — puts the first return between 35 and 100 years after the reply is sent, and a three-stage disclosure inside two or three centuries: long for a person, short for a civilization that has just built a transmitter, and three orders of magnitude shorter than a reply to the sender itself. The mesh’s node count is therefore not an architectural flourish; it is the parameter that decides whether the gate chain is a conversation or an archive.

Scale, for reference. The mesh at 60 ly spacing is of order 10^7 nodes with optical links across a galactic disc. On 30 January 2026 a single company on a single planet filed with its regulator for a constellation of up to 10^6 actively powered satellites with optical inter-satellite links in one planet’s low orbit, for compute rather than communication, and the filing was accepted for review five days later (FCC 2026; SpaceNews 2026); the launch vehicle intended to deploy it carried its first test payloads of the associated satellite generation in May 2026, each rated at roughly $10 \times$ the downlink of its predecessor. Nothing about that system resembles the mesh — its nodes are powered, short-lived and a few hundred kilometers apart — and it is cited for one fact only: node counts within an order of magnitude of what the architecture needs are within the reach of a civilization at our own rung, on our own timescale. The mesh is not an extrapolation past anything we can see being built.

What this section claims. That the sender’s half of the chain is a fleet of order 10^4 – 10^5 carrier buses, each delivering thousands of probe seeds and hundreds to thousands of mesh nodes along a route of 10^4 – 10^5 ly, at a one-time energy cost between a few centuries and a few millennia of a planet’s insolation depending on the seed mass; that the node count is what makes the exchange a conversation rather than an archive; and that node counts of this order are being proposed by a civilization at our own rung. It does not claim that any of this has been built. §4 does not test it and cannot.

2.7 What must sit at the star, and what it must cost

Step 4 of §2.5 is the load-bearing step. A signal in stellar output requires a persistent agent in the stellar environment: surviving on the timescale of the beacon, self-repairing, able to impose a small modulation on emission the star produces anyway. **This is the largest unsupported assumption in the framework.** Nothing in §4 tests whether such a thing exists; the searches test only whether the modulation it would produce is present.

Where it would have to be. Not the photosphere. The channels that are cheap to modulate — Lyman- α , the EUV lines, the soft X-ray background — form in the chromosphere, transition region and corona, which are optically thin and mechanically tenuous. **The cheapest emission to modulate is formed exactly where a persistent structure is most plausible** — and where the radiative flux limits our own probes to brief perihelion passes.

What it would cost to run. The limits of §4.2, read backwards, give the power a sender must command to sit at the threshold of our detection:

Channel	Power in channel	Our limit	Power to sit at it	× world energy use
total solar irradiance	3.8×10^{26} W	62 ppm	2.4×10^{22} W	1.2×10^9
Lyman- α 121.6 nm	1.7×10^{21} W	1.4×10^{-6}	2.4×10^{15} W	118
soft X-ray 0.1–0.8 nm	2.8×10^{17} W	6×10^{-4}	1.7×10^{14} W	8.4

Taking world primary energy use as 2×10^{13} W. The channel powers follow from each band’s share of the solar luminosity.

This is the strongest argument in §2 that survives contact with numbers. Modulating the bolometric output at our sensitivity would take a billion times humanity’s energy budget; a narrow line costs **10⁷ times less** for the same detectability, and the soft X-ray background at our published limit about eight times what our civilization already generates. **Bit-energy conservation therefore chooses the carrier, not just the message length,** and §4.2 reports limits in all three channels.

Read naively, this refutes the premise. If a 1.4×10^{-6} modulation of Lyman- α means *generating* 2.4×10^{15} W, the sender has built a transmitter more expensive than any radio beacon. **The escape is that a modulation need not be generated.** Flux the star already radiates can be *redistributed* — occulted, scattered, or shifted between channels — and the cost is area, not power. To occult a fraction 1.4×10^{-6} of the disc requires **2.1 × 10¹² m²**, a square some 1,460 km on a side, or 2.1×10^9 kg of micron film. To produce a *two-minute* modulation that structure must also change what it intercepts on that timescale — rotate, translate, or vary its opacity. **We do not know which, and the searches do not distinguish them:** the area is what the energy argument bounds, and the mechanism is what it does not.

Quantity	Value
projected solar disc	1.52×10^{18} m ²
area intercepting 1.4×10^{-6}	2.1×10^{12} m ²
as a square	1,460 km on a side
mass at 1 g m ⁻² (micron film)	2.1×10^9 kg
radiated power to sustain it	zero

That is the difference between 10^{15} W and nothing at all. It is also, plainly, **Bracewell’s mechanism and not ours** — a local artifact in the line of sight, developed for co-orbital positions by Benford (2019); what this paper claims as new is

the search strategy (§1.4).

Two consequences follow. **The fast-band limits of §4.2 bound any modulator at the cadences examined, gray or selective alike:** a gray occulter moves Lyman- α along with everything else, so at two minutes it is bounded directly by the 1.4×10^{-6} Lyman- α limit, not by the daily achromatic figure. The achromatic and line-ratio searches (§A.1) add not a separate bound but a **discriminator between mechanisms**, and only at the cadences at which they were run: a spectrally selective modulator is the more demanding object, and distinguishing it from a gray one at two-minute cadence has not been done here.

An honest reframing. “No transmitter exists” is stronger than we can defend. The defensible claim is narrower: **no dedicated radiating apparatus is assumed, and the carrier is not specified in advance.** That leaves the work a systematic anomaly search over a principled combination space — where its value lies in any case, since the limits of §4.2 stand however the modulation is imagined to arise, or whether it arises at all.

Why not look for the thing gravitationally? Lunar laser ranging reaches millimeters; the search is not worth trying, by eleven orders of magnitude.

$$dr \sim 2GMr/(D^3 n^2)$$

modulator of this section, 2.1e9 kg, at 1 AU
short of a 1 mm normal point by
mass at 1 AU that WOULD give 1 mm
or: how near a 2.1e9 kg probe must be

r = Earth-Moon separation, n = lunar mean motion

4.5 femtometres
2.2e11
4.6e20 kg = half of Ceres
24,800 km, inside geostationary

The radiative route is worse: the Moon’s response to radiation pressure is 8.3×10^{-5} m per unit fractional irradiance, and Lyman- α is 4.4×10^{-6} of the total, so the 1.4×10^{-6} limit of §4.2 moves the Moon by 5×10^{-16} m. **Ranging would have to improve by 8.6×10^6 to compete with pointing a photometer at the Sun;** planetary ephemerides do better, but by order 10^2 , not 10^{11} . **No dynamical method reaches an object this light** — a modulator that redistributes rather than generates has almost no mass to find.

2.8 What the argument has to predict

An argument of this shape earns its place only by making the searches it motivates sharper than the nulls it might be used to excuse. Four predictions follow:

	Prediction	From	Tested in
i	the carrier is a relationship among observables, not an observable	no shared units (§3.1)	§4.2, the sweep
ii	the key space is enumerable , not clever	the gate must be payable (§2.3)	§3.3
iii	the message is short , so the band is fast	a derived pointer is of order 10^2 bits, not 10^6	§4.2–4.3
iv	it is in the loudest data, not the quietest	a pointer must be found, not hidden	not tested here

2.9 The gate as a cryptographic object

The constraints of §2.1 and the proof-of-work argument of §2.3 describe the gate functionally. It is worth naming the cryptographic primitive it actually corresponds to, because the correspondence is exact, because it repairs an imprecision in §2.3, and because it makes a prediction about **search order** that the bottom-up program would not otherwise have made.

The gate is a symmetric cipher under brute-force key search. The key is the tuple (*which observables, which functional form*); the ciphertext is the public archive; the plaintext is the imposed modulation; the key space is $C(N,k) \times |F|$, which is the combination space of §3.3 under another name. Everything is public except the key — the archives, the physics, the method, this paper — which is **Kerckhoffs's principle** in its purest available form, and not by accident: a sender wanting concealment would not modulate a star.

The property that makes the scheme work at all is **self-identifying plaintext**. A receiver knows when the right key has been tried because the statistic fires; no crib and no known-plaintext pair is needed. That is the same property that makes brute force viable against a block cipher — one can recognize English without knowing in advance *which* English — and it is what the detection criteria of §3.5 formalize.

It is not proof of work in the hashcash sense, and the difference matters. In hashcash the *sender* performs the work, to price a message or prove commitment. Here the *receiver* performs it. The direction of payment is inverted, and with it the design goal: this is a **proof-of-capability challenge** — a CAPTCHA turned inside out, a puzzle a computer is meant to *pass* rather than fail. The analogy in §2.3 is right about the mechanism and wrong about the polarity.

The nearest technical antecedent is not in the SETI literature. Rivest, Shamir & Wagner (1996) pose a *time-lock puzzle* whose solution requires a chosen amount of computation, and set its difficulty from *projected future hardware speed* so that it opens at a chosen future date. That is the two-sided difficulty bound of §2.3, arrived at independently in a different field.

The construction differs in one respect, and the difference is favorable to the argument here. A time-lock puzzle, and its descendant the verifiable delay function, is **inherently sequential**: the delay cannot be bought off with more machines. A key search is embarrassingly parallel.

resource	time to exhaust a 2.8×10^{23} FLOP tier
El Capitan, dedicated	6.2 d
Frontier, dedicated	9.8 d
a volunteer network at Folding@home's 2020 peak	1.3 d

A designer wanting a *calendar* clock therefore chose the wrong construction. A designer wanting to gate on **total civilization compute** chose exactly the right one — and that is the more sensible quantity to threshold, because it is what the lower bound of §2.3 is trying to measure in the first place. Calendar dates are parochial; one cannot reach a capability level by waiting.

Key-space entropy, and why the data requirement is mild. The required record length is set by the need for the true key's statistic to stand above the maximum of K draws from the null. That maximum grows as $\sqrt{(2 \ln K)}$ while detectable amplitude falls as $1/\sqrt{N}$, so $N \propto \ln K$.

tier	combinations	forms	H(K) bits	record length, relative
pairs / 30 channels	4.35×10^2	4	10.8	1.00 ×
triples / 30	4.06×10^3	4	14.0	1.30 ×
quadruples / 30	2.74×10^4	4	16.7	1.56 ×

tier	combinations	forms	H(K) bits	record length, relative
triples / 60	3.42×10^4	10	18.4	$1.71 \times$
quadruples / 60	4.88×10^5	10	22.2	$2.06 \times$
combinatorial maximum, $k = 30$ of 60	1.18×10^{17}	10	60.0	$5.58 \times$

The key space spans **fifteen orders of magnitude**; the required record length spans $5.6 \times$. Compute scales with K and data scales with $\log K$. **The gate is compute-bound, not data-bound** — which is how one would design a lock intended to be opened eventually, and which is an independent route to the empirical finding of §4.3 that sensitivity is not the binding constraint.

Unicity distance. Shannon’s unicity distance is the ciphertext length below which more than one key yields a plausible decryption. That is the same quantity as the multiple-testing threshold of §3.5: below it, wrong keys fire. The $_N^2/\alpha$ scaling of §3.4 is therefore an instance of a known information-theoretic bound rather than a bespoke statistical fix — and a **void** result is, in this language, a search conducted below the unicity distance for its key space.

Weak keys, and the search order they imply. A cipher is strong when its keys are drawn uniformly from the key space. These are not. The sender is constrained: the combination must be one the receiver plausibly *measures*, on records long enough to hold a message, and that concentrates the key distribution onto the best-instrumented, longest and most obvious observables. In cryptographic terms these are **weak keys**, and brute force is the wrong attack against them. No password cracker begins with brute force; it begins with a wordlist.

attack	keys	family-wise threshold
exhaustive, quadruples over 30 channels	27,405	1.8×10^{-6}
wordlist: pairs, triples and quadruples over the six most plausible channels	50	1.0×10^{-3}

The wordlist is $548 \times$ **cheaper and, because the multiple-testing penalty scales with the number of keys tried, roughly $550 \times$ more sensitive** — a signal that an exhaustive sweep would bury under its own trials correction can survive the wordlist’s. That gain is earned only by committing to the wordlist in advance; choosing it after seeing exhaustive results would be selection.

The wordlist that follows from the designer’s own constraints — the channels any civilization measures early, precisely, and for a long time — is the §3.6 ordering reached by a different route. **Two independent arguments converging on the same search order is itself evidence the ordering is right**, and the consequence for this program is concrete: the unsearched rows of §3.6 should be run before any exhaustive tier, not because exhaustion is expensive but because it answers a less likely question at greater cost.

One caveat, stated because it bounds what the argument licenses. The weak-key claim rests on a prior over what a sender would choose, and that prior is our reasoning about their design rather than anything testable. It justifies search **order**, which costs nothing if it is wrong. It would not justify restricting the search **space**, and nothing here should be read as license to drop the exhaustive tiers.

3. The combination space

3.1 Admissible carriers, and a constraint that had to be sharpened

A sender sharing no units with the receiver cannot encode in a quantity carrying units. Admissible carriers are therefore **dimensionless**: ratios of like observables, normalized residuals, quantities scaled by their own dispersion. This removes

every raw measurement and admits only combinations, which is what converts an unbounded question into a countable one.

A fractional modulation depth of a single observable is admissible. The deepest limits in §4.2 are single-channel; a depth is a quantity divided by its own baseline — a normalized residual, dimensionless in exactly the required sense — so **the single-channel searches lie inside the combination space**, in its $k = 1$ corner.

Dimensionless is necessary and not sufficient, and we learned this the expensive way. A first sweep included plasma beta and the Alfvén Mach number *because* they are dimensionless, and returned sixteen Bonferroni survivors — every one a pair involving those two quantities. Neither is a measurement: both are *functions* of density, temperature, speed and field strength, all already in the channel set, so testing $|\mathbf{B}|$ against β tests $|\mathbf{B}|$ against a formula containing $|\mathbf{B}|$. The geomagnetic indices failed the same way, as a downstream response to the solar wind. Requiring that a candidate be dimensionless *and* independently measured took sixteen survivors to four, all textbook heliophysics.

Extending the set to thirty (§3.2) showed that two clauses are still not enough: nine of the twenty-two survivors are pairs each dimensionless and independently measured yet not independent of *each other* — the four Ca II K indices come from a single spectrum, the abundance ratios Fe/O, C/O and He/O share a denominator, CME rate and speed come from one catalog. **A pair sharing an instrument, a spectrum or a denominator will produce a survivor whether or not anything is modulating it.**

The regularizer therefore has three clauses, not two. An admissible carrier is a pair of quantities that are (i) dimensionless, (ii) independently measured, and (iii) not derived from a shared instrument, spectrum or normalizing quantity. Clause (iii) was added *after* seeing the thirty-channel result and is a post-hoc correction — which is why §4.8 reports the uncorrected survivor list in full, and why a pre-sweep correlation audit at a **declared threshold of $|r| > 0.90$** (the value used in §4.8) is now part of the procedure, to be applied *before* any future sweep rather than offered as an explanation after it.

3.2 The channel set

Thirteen channels were used in the first sweep: sunspot number, F10.7, cosmic-ray flux, solar wind speed, density and proton temperature, interplanetary field magnitude, the alpha-to-proton ratio, total solar irradiance, Mg II core-to-wing, the GOES X-ray background, the > 10 MeV integral proton flux, and Lyman- α ; the geomagnetic indices are excluded by §3.1. An earlier version used ten and was written up as covering the pair space — wrongly: it covered the pairs of ten channels chosen by the author, **the same opportunism the sweep exists to escape, moved up one level.**

Seventeen further channels were then acquired to close the pair space, all public:

Added channel	Source	Days	Span
O^{7+}/O^{6+} , C^{6+}/C^{5+} , $\langle q \rangle(Fe)$, $\langle q \rangle(Si)$, $\langle q \rangle(O)$, Fe/O, He/O, C/O	ACE SWICS 1.1, quality-flagged	$\sim 4,783$	13.5 yr
sunspot area; hemispheric asymmetry $(N - S)/(N + S)$	RGO + USAF via MSFC	51,918 / 44,240	142 yr
mean line-of-sight solar magnetic field	Wilcox Solar Observatory	14,564	48.8 yr
Ca II K emission index, K2V/K3, K3, $\Delta K1$	Sac Peak via LASP LISIRD	4,112	38.9 yr
CME rate; mean CME speed	SOHO/LASCO CDAW catalog	$\sim 10,100$	30.4 yr

The charge-state and abundance ratios are the strongest additions on the criteria of §3.1: dimensionless by construction; the hemispheric asymmetry likewise, over 142 years.

That gives **thirty observables**, the denominator of every coverage fraction in this paper; §5.4 records that a different but defensible list moves it by a factor of order unity.

3.3 Forms, and the size of the space

Each pair is tested under four dimensionless functional forms — ratio, log-ratio, and the corresponding normalized residuals — over the bands the records support. A pair is dropped only when its two channels share fewer than 2,000 days.

Enumeration	Combinations	Examined	Tests	Coverage
13 channels, pairs (first sweep)	78	78	229	18% of the 30-observable pair space
30 observables, pairs	435	382	1,074	88% ; 53 pairs lack 2,000 d of overlap
30 observables, triples	4,060	2,934	11,736	72% reachable, but void — see §4.9

The size of this table is the recognition gate of §2.4 made explicit: nothing in it is new data, and closing it is a matter of recombination capacity. The pair space is now essentially closed at thirty observables. What remains unexamined is an archival fact, not a choice: 53 pairs and 1,126 triples have no 2,000-day window in common. Coverage is now limited by record overlap rather than by effort, a different and more tractable problem.

The sweep is stratified by physical category, and that stratification is what makes its null informative: 510 of the 1,074 pair tests cross the Sun–heliosphere boundary, exactly where a beacon coupling something at the Sun to something measured at 1 AU would appear.

3.4 Nulls, and matching the surrogate to the statistic

Significance is assessed against surrogates rather than a parametric model, since solar series are red, non-Gaussian and gapped — and the choice of surrogate is not a matter of taste. Two controls in this work were **no-ops**: they could not have failed, and were caught only after returning reassuring answers.

Failure	Case	Mechanism
Preserved by its own surrogate	phase randomization against a mode-comb statistic reproduced the comb, at 135.3 μHz in the <i>surrogate</i> , with the same significance as the data	phase randomization preserves $ \text{FFT} ^2$ exactly and the statistic is computed from the power spectrum. The surrogate preserved the very thing it was meant to destroy
Invariant by construction	sorted-gap dispersion against IAAFT returned $p = 1.00000$ with a null standard deviation of exactly zero	IAAFT preserves the marginal distribution, and a function of the sorted values alone cannot move under it

The rule that falls out: match the null to what the statistic is made of. Nonlinear structure needs IAAFT; level structure needs phase randomization. **Anything determined by the power spectrum alone cannot be tested this way at all** — lag- k autocorrelation is the Fourier transform of the power spectrum, so no spectrum-preserving surrogate can serve as its null, while permutation destroys all autocorrelation and is beaten trivially by any red series.

3.5 What counts as a detection


A candidate had to clear four hurdles, fixed before the data was examined:

	Requirement	In practice
1	Bonferroni significance across the full test count of its sweep	for the 13-channel sweep, $p < 2.18 \times 10^{-4}$ across 229 tests
2	a null matched to the statistic, per §3.4	an unmatched null certifies nothing
3	replication on an independent instrument viewing the same sky	the twin gate removed 9 of 9 survivors in one search and 163 of 163 in another
4	survival of a control the signal cannot occupy, supplied by physics where possible	acoustic cutoff, ecliptic pole, anti-sidereal frequency (§4.4). A physical control cannot be tuned after the fact

Two caveats belong with the sweep itself. Several pairs sit at the p -floor set by the shift count, $1/6001 = 1.67 \times 10^{-4}$, so a reported $p = 0.00017$ means only *as significant as six thousand shifts can show*; that floor is adequate here and **not** for a larger enumeration — §5.2 gives the scaling and a run of ours that violated it. And the exact survivor count moved between four and five across two runs, because the shift generator was seeded from a string hash Python randomizes per process. **The category structure is stable across runs; the exact survivor count is not**, and no claim in this paper rests on it.

3.6 Where a designer would put it, and where we have looked

The sender's side gives a sharper question — **given a choice of channel, which would a designer pick?** — answerable before any data is examined, and set out ahead of the results so that the ordering it produces is a **prediction** rather than a description of what we happened to find.



Where a designer would put the pointer, and where this paper has looked. Diamonds: the depth a cost-minimizing designer would set for each channel. Circles: the depth this program reaches at 95% recovery. Rows 1–3 and 5–9 have been searched; rows 10–15 sit below our threshold.

The detection threshold for a periodic fractional modulation of depth d , per-sample fractional noise σ_1 , coherent integration over N samples, family-wise false alarm α and empirical noise scale μ is the expression behind §4.2:

$$d_{\min} \approx 2 \sigma_1 \sqrt{(2 \mu \ln(N/\alpha) / N)}$$

For $\sigma_1 \approx 10^{-3}$ per minute and $N \approx 2.6 \times 10^6$ (five years) this gives $d_{\min} \approx 5 \times 10^{-6}$, consistent with the injected Lyman- α limit. Sender cost is linear in d (§2.7): 1 ppm needs about 1.5×10^{12} m² of interceptor, 1 ppb about 1.5×10^9 m². **The designer therefore prefers the shallowest depth any receiver channel resolves** — and resolving power is highest in timing and frequency, lowest in composition and particle flux, with irradiance between.

Why the margin column is nearly all ones. A proof-of-work gate is set just above what a first-capable receiver can resolve (§2.3), so the level “required” in each channel is defined by our own capability. A margin near unity is what the gating argument predicts, not evidence for it; the informative column is the last one — **whether the channel has been searched at that level at all**. Rows where the margin is near unity and no search has been run are where the framework’s own logic says a first-contact pointer would sit.

#	channel, in the designer’s descending order	level a designer would set	our reach	margin	searched there?
1	narrow UV/EUV lines, minute cadence	1–10 ppm	1.4 ppm	~1–7	yes — this paper, §4.2
2	radio spectral index, second cadence (RSTN)	10^{-4}	$\sim 10^{-4}$	~1–10	yes — §4.14, added after this table was drawn
3	sub-minute EUV (PROBA2/LYRA, SDO/EVE ESP)	1–10 ppm	$\sim 10^{-6}$	~1–10	yes — §4.18, added after this table was drawn
4	static artifact, lunar surface	~1 m/pixel	LRO NAC 0.5 m	~1–2	partially — no systematic survey
5	passive occulter, information in transit <i>timing</i>	10 ppm dips	31–108 ppm at 6 h, measured	unreachable	yes — §4.17, void
6	broadband irradiance, coherent (VIRGO, TIM)	0.1 ppm	2.0×10^{-7} measured	0.5	yes — §4.15, added after this table was drawn
7	p-mode frequency structure	10^{-6}	1.46×10^{-5} measured	0.069	yes — §4.16, added after this table was drawn
8	neutron monitor / particle flux	10^{-4}	7.6×10^{-5}	~1	yes
9	solar-wind composition	10^{-3}	10^{-3}	~1	yes (§4.8); implausible as a carrier
10	line-profile ratios, disc-integrated	10^{-7}	10^{-6}	~0.1	no
11	polarization, disc-integrated	10^{-7}	10^{-5} – 10^{-6}	~0.01–0.1	no
12	static artifact, Lagrange points and co-orbitals	~10 m	few hundred m	$\sim 10^{-2}$	no
13	static artifact, asteroid surfaces	1 m/pixel	a dozen bodies imaged	$\sim 10^{-3}$	no
14	artificial neutrino line, MeV	$\sim 10^{-2}$ of solar rate	background-dominated	$\sim 10^{-2}$	no
15	artificial neutrino line, Glashow resonance 6.3 PeV	~1 event/yr	~1 per decade, poor direction	$\ll 1$	no

Order-of-magnitude throughout, except our own reach in rows 1, 8 and 9, which is the injection-verified figure from §4.2. **The margin column is the level a designer would set divided by our reach**, so a margin above 1 means the channel is resolved finely enough to carry the signal and a margin below 1 means it is not; where either column is a range, the margin spans the corresponding ends. Solar neutrino flux appears nowhere as a carrier: nothing at the atmosphere can modulate a flux produced in the core, which is precisely what makes it the one control channel no mechanism in this paper can touch.

The margin column is an ESTIMATE, and every estimate in it that has since been measured came in worse. Rows 5, 6 and 7 were searched after this table was drawn (§4.15–§4.17), and in all three cases the measured reach fell short of the estimated one:

row	margin as estimated here	margin as measured
5 — occulter transit timing	~1	unreachable ; the six-hour noise floor is 31–108 ppm against a 76 ppm Venus transit (§4.17)
6 — broadband irradiance	~1	0.5 (§4.15)
7 — p-mode frequency	~0.1–1	0.069 (§4.16)

Three of three, all optimistic, and row 5 by enough that the row is not searchable at the designer level at all. The estimates were made by reasoning from instrument specifications rather than by injection, and reasoning of that kind has been systematically generous: a specification describes what an instrument resolves, not what survives a detrending, a continuum estimate and a trials correction. **The unmeasured entries in this column should therefore be read as upper bounds on capability rather than as capability**, and the rows still marked ‘no’ in the last column may be further from reach than they appear. What the table continues to establish is the ORDERING — which channels a designer would prefer — and that ordering is unaffected.

What the ordering says about this paper. The program has been most thorough exactly where the designer’s probability is *highest* (row 1) and where it is *lowest* (rows 8–9), and thinnest in between — not a designed allocation, but where the archives happened to be easiest. **Rows 1, 2, 8 and 9 are capable and searched**: row 1 is the strongest single statement this paper can make — null to 1.4 ppm at the level a designer would plausibly choose — and row 2 was run after this table was first drawn (§4.14). **Rows 3, 4 and 5 are capable and not searched**, and rows 6 and 7 were searched after this table was drawn (§4.15, §4.16) and are null, row 7 at a measured margin of 0.05 rather than the 0.1–1 estimated here: the archives exist, the resolving power is there, and no search has been run. If the pointer is meant for a first-capable receiver, the framework’s own logic puts it there at least as plausibly as in row 1, and **the honest reading is that this paper has not yet tested the framework’s best guess**. **Rows 10–15 are below threshold**: if the designer aimed at a more advanced receiver, the message is there and cannot yet be seen — a gap that closes on a decade timescale (Rubin/LSST and small-body missions, rows 12–13; Hyper-Kamiokande, JUNO, DUNE, KM3NeT, IceCube-Gen2, rows 14–15).

The program this implies is to run the remaining rows in the order 5, 6, 3, 2, 7, then 4 — cheapest and best cost-per-bit first — and **that order is not the one the searches so far have followed**. Rows 5 and 6 are the same archive twice: SOHO/VIRGO SPM and TIM at minute cadence, first as a transit-timing-sequence search after Arnold (2005), then coherently at 10^{-7} . Row 4 belongs to the archaeological strand of the field (Davies & Wagner 2013; Benford 2019).

3.7 Channels above our threshold, and why a designer would not put the pointer there

Rows 14–15 of the §3.6 table sit far below unity, and there is a class of physics behind them that deserves its own statement, because a reader will ask why the search does not look there and because the answer is one of the framework’s more useful predictions.

What we cannot yet modulate or detect. For the electromagnetic channels of §3.2 our position is well inside the capable regime: we detect, we resolve, and we could in principle modulate. For two channels of known physics we are at or below the floor of capability, and for a third there is no floor to be below.

channel	our detection capability	what modulation would require of a sender	status
neutrino beam	IceCube, 1 km ³ : about one Glashow-resonance event per decade, cascade direction to a few degrees (IceCube Collaboration 2021); Super-Kamiokande, 50 kt: ~15 solar ⁸ B events per day, no useful spectral line sensitivity; Hyper-Kamiokande, JUNO and DUNE raise the mass by ~10× in the 2030s (Hyper-Kamiokande Proto-Collaboration 2018)	a directed beam from an accelerator — a dedicated, powered, radiating apparatus (Learned, Pakvasa & Zee 2009; Silagadze 2008), which is what §2.1 excludes for the pointer	primitive detection; no sender-side economy
gravitational waves	LIGO–Virgo–KAGRA: strain ~10 ⁻²³ Hz ^{-1/2} at 100 Hz, sources at 10–1,000 Hz (Aasi et al. 2015; Abbott et al. 2016); LISA in the mHz band from the mid-2030s (Amaro-Seoane et al. 2017); pulsar timing arrays at nHz	strain scales as (G/c ⁴) × (mass × velocity ²)/distance ≈ 8 × 10 ⁻⁴⁵ in SI; a 10 ⁶ kg mass at 1 km s ⁻¹ at 10 pc yields h ~ 10 ⁻⁵⁰ , twenty-seven orders below detectability. Only astrophysical masses radiate detectably	experimental detection; physically closed as a carrier
physics not yet known	none by definition	a sender cannot encode for a receiver in a channel the receiver has no instrument for and no theory to build one from	below gate 0 by construction

Neutrinos are excluded from the pointer by cost, gravitational waves by physics, and unknown physics by the recognition argument of §2.4: a puzzle posed in a channel the receiver cannot conceive of is not a gate but a wall, and a wall filters for nothing.

The framework’s prediction, and the design reason behind it. A designer *could* set the pointer in one of these channels, and would thereby select for a receiver centuries past the point of being able to reply. §2.3 says why a designer would not. The gate is bounded from above: it must be passable by the receiver it is meant for, or it is operationally no gate at all. The receiver the architecture is meant for is one that can build a transmitter and search a combination space — capable, but not necessarily mature — and a pointer that only a megaton neutrino observatory or a space-based gravitational-wave interferometer could read would choke off contact with exactly the receivers the safeguard exists to reach: species obviously capable of the exchange, whose remaining immaturity is the thing a staged disclosure is designed to accommodate. **The lower bound on difficulty filters out receivers that cannot act; the upper bound is there so that the filter does not also remove receivers that can.** The framework therefore predicts that the *pointer* is photonic and combinatorial, at the level of §3.6 rows 1–7, and that channels above the receiver’s threshold appear, if at all, as content behind later gates at step 12 — where a receiver that has already demonstrated a transmitter can be told to build a detector, and where the sender’s neutrino accelerator is paid for only after somebody has answered. This is the same asymmetry as everywhere else in the chain: the expensive channel is reserved for the stage at which the receiver has already paid to be found.

Two consequences for this paper. First, the absence of neutrino and gravitational-wave searches in §4 is not a gap in coverage of the pointer; the framework says the pointer is not there, and a designer who put it there would have defeated their own purpose. Second, the §5.5 recommendation of a monoenergetic-line search in existing neutrino

catalogs stands, but as a test for a *later-stage* channel or an unrelated beacon, not for the framework’s pointer — a result there would be interesting on its own terms and would not bear on the gating argument.

4. Results

4.1 Summary

Twenty-nine searches were run, every one against a public archive collected for unrelated purposes. Three sections at the end of these results are not searches and do not enter the tally: §4.10 puts the one result that came closest to a detection through the procedure of §3.5, §4.11 tests whether the fast-band searches lose signal to a geometric term the coherent model omits, and §4.13 applies the method to a dataset the architecture cannot reach. All are null.

Scope. Every search examines the Sun, the heliosphere, or the solar system out to 160 AU — forced by the architecture, since step 4 of §2.5 places the modulator *at the local star*. **A search of distant sources tests a different hypothesis.** A further 23 searches of pulsar, pulsar-timing-array, X-ray binary, eclipsing binary and astrometric archives were run during this program and are reported separately. **All were null or void; none is excluded here because of its outcome**, and their exclusion is stated so that the coverage claimed in §4.3 is not read as covering more than it does.

Outcome	N	Meaning
null, injection-verified limit	13	amplitude bounded and recovery measured (§4.2, Figure 2)
null, analytic sensitivity only	1	amplitude stated, recovery not measured
null, false-alarm and power pair	3	bounded by a measured false-alarm rate and power, not by a single amplitude (§4.2, Appendix A)
null, no amplitude limit	8	reported as constraining nothing
void	3	detector or null failed its own validation, or the channel proved unsearchable at the level required; §4.5, §4.9, §4.17
detection	1	a known signal, recovered as a positive control (§4.6)

Why each search ended where it did, which is not the same question. “Null” is one word covering four situations, and the difference decides where effort should go next.

why it ended	N	what it is a statement about
nothing above threshold	22	the sky. The detector worked, the control passed, the channel was empty at the stated level
underpowered	3	the analysis. The search ran and could not have found the effect if present — measured power 36–53% (§A.2 rows 6, 13, 14)
detector failed its own control	1	the analysis. A statistic that sat at 0.954 for data and surrogates alike, and another that fired on unmodified data — one search failing its own control twice over (§4.5)

why it ended	N	what it is a statement about
no defensible null	1	the analysis. The surrogate is correct and calibrated, and the Sun rejects it — “no signal” cannot be specified for a star with its own higher-order structure (§4.9)
the Sun is too loud	1	the data. A 76 ppm Venus transit against a 31–108 ppm noise floor at six hours. Not repairable by analysis (§4.17)
the archive removed the observable	—	the data. A level-2 product had already discarded the roll correction, and the spatial structure with it (§4.19). The LOI channel sits outside the twenty-nine: the observable was gone before a search could be scored
detection	1	a known signal, recovered as a positive control (§4.6)

Only the first line is astronomy. Three are statements about our own machinery and two about the archives we were handed. Twenty-nine searches and twenty-five nulls reads as broad coverage; the honest reading is **twenty-two channels searched competently, three searches too weak to count, two defeated by their own machinery, and one defeated by the data.**

4.2 Limits

All limits are fractional amplitudes of a modulation in the named observable, against the stated null. **Each is the amplitude recovered 95% of the time under injection**, against the threshold the search itself uses ($\mu = \text{median}(R)/\ln 2$, $\alpha = 0.05$).

Observable	Band	Limit, 95% recovery	Basis
Lyman- α irradiance	2 min	1.4×10^{-6}	GOES-16 EUVS 1-min, 1,945 d
Lyman- α irradiance	5 min	3.3×10^{-6}	as above; the 5-minute band is the p-mode band (§4.6)
cosmic-ray flux	5 min – 3 h	7.6×10^{-5}	Oulu 1-min, 26 yr, 100.0% coverage
cosmic-ray flux	15 min – 3 h	7.6×10^{-5}	same record, narrower band — the 62-yr 5-min product was not injected and is not tabulated
soft X-ray, 0.1–0.8 nm	2 min	6.0×10^{-4}	GOES XRS-B 1-min, 9.68 yr
soft X-ray, 0.1–0.8 nm	1 day	1.5×10^{-1}	GOES daily background, 13,289 d
total solar irradiance	30-day box	2.3×10^{-4}	LASP composite, 5,041 quiet days; matched box filter
spectrum, achromatic	daily	5.0×10^{-2}	7 UV lines + XRS-B, frozen filter; injected square wave
multi-scale construction	6 time bases	5.0×10^{-3}	Oulu 1-min, 26 yr; both repeated and spliced constructions detected, control null at $\epsilon = 0$

Observable	Band	Limit, 95% recovery	Basis
interplanetary field	self-keyed	0.10σ	ACE MAG + F10.7, 8 yr, 15.3M samples; 99.7% recovery, measured false-alarm 6.3% against a 5% target

Ten injection-verified limits in this table; the eleventh is the cross-viewpoint bound of §4.12, which is in different units and is stated there. Three further searches bound a modulation through a measured false-alarm and power pair rather than a single amplitude — cross-channel coherence (FP 3.5%, 100% power at 0.40σ), self-keyed spread spectrum (FP 6.3%, 99.7% at 0.10σ) and the Voyager radial coincidence (FP 3.0%, 100% at 50% injection); they are in Appendix A because they do not reduce to one number. One row remains unverified and is reported as a sensitivity, not a limit: the sidereal fold (1.8×10^{-4}). It resisted injection because the detector fires on unmodified data — the seasonal leakage of §4.4 is real — so a random-phase injection cancels as often as it adds; it needs an injection path built around its anti-sidereal gate rather than around the fold.


An earlier version of this table reported analytic threshold-crossing amplitudes as though they were injection-established confidence limits. Neither was true — running the injections showed those figures optimistic by factors of 1.2 to 4.8, and the analytic column was removed rather than corrected, because it could not be regenerated from the pipeline (§5.6). **Nothing downstream used it:** the energetics of §2.7 and the cadence ratio of §4.3 derive from the injected values. **Three rows needed no correction — the achromatic search, the multi-scale construction and the self-keyed test — and they are exactly the three whose injection curves were written into them from the start;** every figure that came instead from an analytic threshold was optimistic, by 1.2 to 4.8 times. **The correlation is perfect and the sample is eleven.** That is the argument for injection stated as compactly as this paper can state it.

Re-running the self-keyed search produced an alignment-specific excess. It is treated at length in §4.10, because the first thing we did with it was wrong.

On the shape of the null. Continuum-normalized power here has mean 1.44 rather than 1.0 and exceeds an $\text{Exp}(1)$ threshold some 4×10^3 times more often than that distribution predicts. That would matter if the searches had assumed $\text{Exp}(1)$; they do not — each estimates the scale empirically as $\mu = \text{median}(R)/\ln 2$, which returns 1.44 on this data, so the threshold actually used sits 44% above the nominal one and the clean candidate lists of §A.1 are consistent with the heavy tail.

4.3 Sensitivity is not the binding constraint

The limits above span five orders of magnitude: solar variability is red, so the noise a signal must exceed falls steeply with frequency, and every step toward shorter periods paid.



Coverage of the combination space. The pair space is 88% examined; the remaining 12% is blocked by retired records and cannot close by waiting. The triple space was attempted and returned void for want of a matched three-body null.

Comparison	Ratio	Interpretation
X-ray, 1 day → 2 min	245×	both rows injection-verified at 95% recovery ; the analytic figures gave 119×, so measuring it made the cadence effect larger, not smaller
full spread of the limits table	10 ⁵	different observables and statistics — not a cadence-only comparison
pair space examined	88%	382 of 435; the other 53 are blocked by retired records and need cross-calibration, not time (§5.7)
triple space examined	0%	2,934 reachable, 2,749 completed, all void for want of a matched null (§4.9)

A search limited by sensitivity is improved by better instruments or longer baselines. This one is not: **the limiting quantity is the fraction of the combination space examined**, which no instrument improves — only measuring more quantities, and combining them more ways, does. The pair space has been closed from 18% to 88% (§3.2, §4.8); the residual 12% cannot close by waiting (§5.7); the triple space was attempted and returned void (§4.9).




Figure 1: Limits against the period at which each was set. **Filled points are injection-verified** at 95% recovery; the single open point is the sidereal fold, whose detector resisted a faithful injection (§4.2). The two cross-viewpoint limits of §4.12 are not plotted — they bound a different quantity. The dashed line joins the two soft X-ray points — same instrument, same pipeline, differing only in period: a factor of 245, both ends measured. **The spread is a fact about cadence, not about instruments.**

Figure 2: Recovery against injected fractional amplitude for seven of the eleven injection-verified rows of §4.2; three are not plotted because they do not report recovery against a fractional amplitude (a p -value curve, a threshold in ϵ , a recovery curve in σ). The dashed line is 95% recovery. **The families do not behave alike:** the cosmic-ray and 5-minute Lyman- α rows sit close to their previously quoted amplitudes, while the 2-minute Lyman- α and 1-day X-ray rows needed corrections of $3.6\times$ and $4.8\times$. Injection is on-bin and therefore a best case. **Plot in the repository.**

4.4 Controls

Each search was gated on a control that could fail, and four did useful work; where possible the control was supplied by physics, since a physical control cannot be tuned.

Control	Basis	Result
acoustic cutoff	waves are trapped only below ~ 5.3 mHz	135 μ Hz comb present at 2.5–4.0 mHz ($A = 0.293, 0.348$), absent at 6–8 mHz (0.002, -0.004)
ecliptic pole	$r \cdot n = 0$ there all year, maximal in the plane	a topocentric line falls $81.6 \rightarrow 5.9$, a $14\times$ gradient from geometry alone
anti-sidereal frequency	364.25 cycles/yr; nothing physical lives there	sidereal 4.9σ but anti-sidereal 6.6σ — the apparent signal is seasonal leakage
twin instrument	two units, one sky	9 of 9 Bonferroni survivors removed; separately 163 of 163. Removes per-unit artifacts only — GOES-16 and 17 are identical designs in geostationary orbit, so eclipse seasons, thermal cycling and the 24 h period are common-mode and survive the gate

The twin and anti-sidereal controls each removed an apparent detection that had passed every amplitude test — the sharpest case a 4.9σ sidereal line beaten by its own control at a frequency where nothing physical can live.

4.5 A void result, reported

One search is reported as void rather than null, because its detector failed validation: **a null from an uncalibrated detector is a statement about the analysis and not about the sky**, and recording it as a null would silently inflate the

coverage this paper claims.

Aperiodic structure in solar output (LASP TSI composite, 46 yr, 16,801 daily samples). The detector, built to find non-periodic but non-random structure, failed its own positive control twice over: the statistic was computed on a channel dominated by the activity cycle, and the surrogate preserved the very quantity the statistic measured. Changing the null altered the false-positive rate and left injection recovery untouched, localizing the fault to the statistic. **No limit is claimed and none should be read into it.**

A second void arose in the triple sweep (§4.9); two more arose in parts of the program not reported here (§4.1). The diagnostic was the same in all of them.

4.6 The one detection, and its relation to prior work

Solar p-mode oscillations were recovered in GOES EXIS Mg II irradiance at one-minute cadence. The detection statistic was not oscillation power but the **spacing of the mode comb**: the autocorrelation of the normalized power spectrum across 2.5–4.0 mHz peaks at **135.1 μHz** on GOES-16 and **135.0 μHz** on GOES-17, against an accepted large frequency separation of 134.9 μHz — fixed by $\sqrt{(M/R_*)^3}$ and therefore predicted rather than fitted.

Solar oscillations in this instrument were reported first by Eden et al. (2024), from the 3-second product — an order of magnitude finer in cadence and precision than the one-minute product used here. **We claim no priority for the detection.** What differs is method, not discovery: Eden et al. measure oscillation *power* at characteristic periods, whereas the statistic here is the *comb spacing*; their Letter reports no frequency separation, autocorrelation, or mode comb. The modest addition is that EUVS-C resolves the radial orders well enough for $\Delta\nu$ to be recovered from the coarser product. $\Delta\nu$ **itself is among the best-determined quantities in solar physics and is not measured here to any useful precision**; its value in this work is that it was known in advance.

Figure 3: Autocorrelation of the continuum-normalized power spectrum, both spacecraft. The comb peaks at **135.1** and **135.0 μHz** against a large frequency separation of 134.9 μHz that was **predicted, not fitted**. The two mode-band traces are separate spacecraft whose records differ in length and therefore in frequency binning — a bin-keyed artifact would land at different lags — and they lie almost exactly on one another, including the side structure at ~67 and ~202 μHz a genuine comb must produce. The two flat traces are the identical pipeline at 6–8 mHz, where waves are not trapped and no comb can exist: the feature is ~30 × weaker there. **A method that manufactured combs would make one in both bands.** *Plot in the repository.*

4.7 Why a re-detection is the most useful result in this paper

A program returning nulls faces an ambiguity it cannot resolve from within: nothing present, or a detector that does not work. Injection addresses part of this, but tests a detector against a signal of assumed shape chosen by the detector’s own author. The p-mode recovery closes the gap from outside: **the same pipeline — same detrending, same continuum normalization, same thresholds — applied to the same archives, returns a predicted astrophysical quantity at 7σ.**

Property	Why it matters
the value was predicted, not fitted	134.9 μHz follows from stellar structure; the pipeline could not have been tuned toward it
two independent spacecraft	records of 1,945 and 1,121 days, hence different frequency binning; a bin-keyed artifact would land at different frequencies. It did not
confirmed by an independent group	Eden et al. (2024) establish the signal exists in this instrument — a check from outside this work

Independent prior publication therefore *strengthens* the control: the nulls in §4.2 are limits from an apparatus demonstrated to detect a real signal of comparable subtlety in the same data.

4.8 The thirty-channel pair sweep

All 435 pairs were attempted; 382 had the required 2,000-day overlap, giving **1,074 tests** against 60,000 circular shifts each, a Bonferroni threshold of 4.66×10^{-5} and an attainable *p*-floor of 1.67×10^{-5} .

Category	Tests	Survivors	Reading
heliospheric × heliospheric	315	14	solar-wind structure; see below
solar × solar	249	8	activity proxies and same-spectrum indices
heliospheric × solar	510	0	the boundary a beacon would have to cross

129 tests fell below $p = 0.05$ against 53.7 expected by chance — the signature of genuine physical coupling in the data rather than of a miscalibrated null.

The first row of interest is the empty one. Five hundred and ten tests span the Sun and 1 AU — the largest block in the sweep, and the one place a signal coupling a solar observable to a heliospheric one must appear. It was empty at thirteen channels and it is empty at thirty. Of the twenty-two survivors, **none is unexplained**:

N	Class	Examples
9	shared instrument or denominator — the failure mode of §3.1, clause (iii)	the four Ca II K indices against each other (one spectrum); Fe/O against C/O and He/O (shared denominator); O^{7+}/O^{6+} against C^{6+}/C^{5+} ; CME rate against CME speed (one catalog)
13	known physics	wind density against five SWICS charge-state and abundance ratios, the standard fast/slow wind discriminators; $ B $ against density (stream interaction regions); proton temperature against alpha/proton; sunspot number against X-ray background, which SWPC publishes as an activity index; Mg II against Ca II K, the two classic chromospheric proxies
0	unaccounted for	—

A correlation audit run before the sweep predicted part of this. Eight pairs exceeded $|r| = 0.90$ on their own overlap — among them Ca II K emission index against K3 at $r = 0.975$ and O^{7+}/O^{6+} against $\langle q \rangle_{(O)}$ at 0.931 — and those pairs duly produced survivors. Sunspot *area* against sunspot *number* did not flag, the audit doing useful work in the other direction: area carries information the count does not.

Null. The sweep recovers the strongest genuine couplings in the data unprompted, confines them to the physically related and instrumentally coupled subsets, and returns nothing across the Sun–heliosphere boundary in 510 tests.

Figure 4: Cumulative *p*-value distribution of the 1,074 pair tests, split by the physical category of the pair, against the uniform expectation (dashed). Both same-domain blocks run hard against the Bonferroni threshold and cross it. **The 510 tests spanning the Sun–heliosphere boundary do not reach it at all** — their steepest *p* is 1.6×10^{-3} , two orders above the line. That is the block in which a beacon coupling a solar observable to a heliospheric one would have to appear. [Plot in the repository](#).

Both sweeps were re-run with deterministic seeds, and the result is reproducible for the first time. The surrogate seed was previously derived from Python’s `hash()`, which is salted per process: three consecutive runs returned 1689220225, 56146563 and 1311111642 for the same key, so no p -value here could be regenerated by anyone, including us. That is not a correctness fault — an arbitrary seed is still a valid seed — but for a paper whose standard is that a null is worth the fraction of a space it excludes, **an unreproducible null is worth less than it looks**. The seeds are now BLAKE2b over a canonical key, fixed across processes, versions and platforms.

Re-running draws different surrogates from the same null, so the two sets should differ in detail and agree in distribution. They do: a two-sample Kolmogorov–Smirnov test gives $D = 0.008$, $p = 1.00$ on the 1,074 pairs and $D = 0.003$, $p = 1.00$ on the 10,996 triples. **The claim this section rests on is unchanged — zero survivors cross the Sun–heliosphere boundary in either run.**

One instability is worth reporting rather than absorbing. Survivor counts at the family-wise threshold moved from 10 to 13 among the pairs and 198 to 206 among the triples, and **182 of 10,996 triples — 1.7% — cross the threshold in one run and not the other**. Every one of those sits below the counting floor of $1/(M+1) = 10^{-4}$, where the p -value comes from the generalized-Pareto tail fit of §3.4 rather than from counted exceedances. **The tail extrapolation is therefore not stable across surrogate realizations at the level of an individual test**, which is a real limit on the device that §3.4 introduces to escape the N^2 surrogate cost. It does not affect the distributions, the boundary result, or the void verdict below; it does mean that a single extrapolated p -value near threshold should not be read as a property of the data.

4.9 The triple sweep, and why it is void

Of 4,060 triples of thirty observables, 2,934 have the required 2,000-day overlap, giving **11,736 tests** under two forms that do not factor into pair statistics: the third-order residual product $d(a)d_{-}(b)d_{-}(c)$, and the log-space curvature $\log r(a) - 2 \log r(b)_{-} + \log r(c)_{-}$, each channel taking a turn as the middle term. 10,996 tests completed.

Stage	Shifts	Survivors	Method
screen	10,000	198	tail-fitted null (§5.2)
confirm	500,000	93	direct exceedance counting
form-level control	10,000	—	synthetic triples, no signal

The screen-then-confirm design worked as intended: **105 of 198 screened survivors failed direct counting**, the tail fit behaving exactly as §5.2 warns. But the 93 that survived are not reported as detections, because a third control invalidates both forms.

Every survivor was the same form. All 198 were log-curvature; not one was a residual product. A perfect segregation by form is a property of the statistic, not plausibly of the sky, so both forms were run against synthetic triples with the red spectra and shared activity driver of the real channels and **containing no signal whatever**:

Form	$p < 0.05$	$p < 10^{-3}$	median z	Verdict
expected, sound form	5.0%	0.1%	0	—
log-curvature	6.6%	0.9%	−0.04	heavy-tailed — $9\times$ excess where survivors live
residual product	0.0%	0.0%	−4.49	cannot fire — a no-op

Both forms fail, in opposite directions. The log-curvature form is roughly calibrated in the bulk and has a ninefold excess in the far tail — precisely the regime that sets Bonferroni survivors. The residual product sits 4.5σ below its own null by construction and can essentially never produce a positive. The clean split by form does not show one statistic firing on signal; it shows the other **unable to fire at all**.

The mechanism is visible in the algebra. $\log a - 2 \log b + \log c$ cancels common-mode variation when the three series are aligned and not when two are circularly shifted, so for channels driven by a common solar cycle the statistic separates from its null *by construction*. Consistent with that, 61% of the confirmed triples are entirely heliospheric, and the Ca II K emission index, K2V/K3 and K3 — three indices from **one spectrum** — confirm together at the floor.

Void, and what would fix it. The circular shift is the wrong null for a three-body statistic: it destroys pairwise structure along with three-way structure. The right surrogate **preserves every pairwise cross-correlation and destroys only the three-way alignment** — the third-order analog of IAAFT — and building it is the prerequisite for any triple enumeration. Two honest limits on the control itself: its 320 synthetic triples measure the false-positive rate at 10^{-3} , *not* at the 4.5×10^{-6} Bonferroni threshold, and the synthetic channels approximate the real correlation structure rather than reproducing it. Neither weakens the verdict: **a ninefold tail excess and a form that cannot fire are sufficient grounds to void a result, though not to confirm one.**

Void. No limit on three-way dimensionless structure is claimed, and the 93 survivors of direct counting are reported as artifacts of an uncalibrated statistic. The triple space remains, in the sense that matters, **unexamined**.

4.10 A candidate, and what the detection procedure did to it

Re-running the self-keyed search (§A.2, row 3) returned a matched-filter correlation against the true F10.7 key of $\rho = -3.52$, against a shifted-key null whose 95th percentile is 2. The quantity ρ is **not a correlation coefficient**. It is the matched-filter output normalized to unit variance under the null — $\rho = \Sigma(a_i/\sigma_a)c_i/\sqrt{n}$, with a the mean-subtracted field residual and c the ± 1 chip sequence — so it is a **z-score**, in units of its own null standard deviation, and equals the Pearson correlation times $\sqrt{n}$. The observed -3.52 therefore corresponds to a Pearson r of -0.065 over 2,919 samples. A nominal Gaussian would read -3.52 as $p \approx 4 \times 10^{-4}$; the shifted-key null gives 7×10^{-3} , because that null is **not** Gaussian, and it is the empirical null that is used throughout.¹⁴ $p = 0.007$ — an alignment-specific excess. The first version of this section set it aside on the grounds that the correlation is **negative**, and a beacon keyed to solar activity would not anticorrelate with it.

That reasoning was invalid, and it is the exact error §5.3 warns against — using the framework to decide what the data mean. The sign is not a property the beacon hypothesis constrains: a modulation imposed by redistributing flux can present with either sign, and dismissing a result for an unattractive sign is procedurally indistinguishable from accepting one for an attractive sign. The detection procedure of §3.5 is applied instead.

1. Is the null matched? F10.7 and the interplanetary field both carry solar rotation, and the published null's arbitrary lags destroy rotational alignment and any beacon alignment *together*; a second null shifted by **integer Carrington rotations** preserves rotational phase and destroys only arbitrary alignment.

2. Does it replicate? The same reduction, key and statistic were run against **Wind/MFI** over the same span; 2,897 of 2,898 days reduced. **This is an instrumental control, not a physical one:** both spacecraft sit near L1 and sample the same plasma, so agreement rules out an ACE artifact, not a real coupling.

3. Both halves of the record?

series	n	ρ	p , arbitrary shift	p , Carrington shift	
ACE 2012–2019	2,919	− 3.52	0.0073	0.0063	excess
ACE first half	1,459	− 3.45	0.0100	0.0097	excess
ACE second half	1,460	− 1.40	0.144	0.146	none
Wind 2012–2019	2,896	− 3.07	0.0123	0.0096	excess
Wind first half	1,448	− 3.14	0.0100	0.0097	excess
Wind second half	1,448	− 1.07	0.285	0.272	none

*The excess survives the rotation-matched null and replicates on an independent magnetometer, with the same sign and comparable magnitude. It does **not** hold in the second half of either record.*

4. Is the second-half absence real, or has the test lost power there? This decides the result, and it has to be measured. The chip sequence is the sign of the day-to-day change in F10.7, which collapses from a median of 3.61 sfu in the first half to **0.83 sfu** in the second as the cycle declines from the 2014 maximum — so the code might simply have degenerated toward a coin flip. Running the injection curve separately in each half settles it:

half	median $ \Delta F10.7 $	power at 0.05σ	power at 0.10σ	false alarm
ACE first	3.61 sfu	31%	84%	3%
ACE second	0.83 sfu	56%	100%	6%
Wind first	3.61 sfu	40%	90%	6%
Wind second	0.83 sfu	42%	94%	5%

The second half is not the insensitive one — it is the more sensitive one: 100% and 94% power at 0.10σ , against 84% and 90% in the first half, because the field residual is quieter near solar minimum. Since power is quoted in units of the local standard deviation, a beacon of fixed absolute amplitude would be easier to see in the second half, not harder. The absence is real.

Verdict: a real anticorrelation, and not a beacon. It survives a rotation-matched null, so it is not solar rotation; it replicates on an independent magnetometer, so it is not an ACE artifact; but it is **confined to 2012–2015 in both instruments**, and the window in which it is absent has more power to detect it. A set-and-forget beacon does not switch off, and the interval carrying the excess is the maximum of cycle 24 — the temporal signature of an activity-driven physical coupling, which is the physics step of §3.5 doing its work. **And it would not clear the trials correction in any case.** The self-keyed family tried one configuration, so there is no within-search look-elsewhere factor; but this is one result among **29 searches**, giving a per-search Bonferroni threshold of $0.05/29 = 1.7 \times 10^{-3}$ (the denominator is 29, not 1,099, because the pair sweep carries its own within-search correction over its 1,074 tests). The observed p of 0.006 to 0.012 does not reach it. **We report it as a candidate that failed replication in time and does not clear the paper’s own trials threshold** — not as a result dismissed for having the wrong sign.

Two things are worth keeping. The rotation-matched null is a control of exactly the kind §4.4 argues for — supplied by physics, not tunable. And we would not have run any of these tests had the sign argument stood: **the sign was doing the work that four measurements should have done**, and it happened to reach a similar destination — the most dangerous way for a shortcut to fail. The anticorrelation itself is a modest heliophysical result we make no claim to have explained.

4.11 The one-year geometric term

Every fast-band search in §4.2 integrates coherently for 1,945 days against a *stationary* carrier, but a modulator fixed in inertial space and viewed from a moving Earth does not present one: the light-travel time and projected velocity acquire an annual term. Two cases: **phase** modulation, from the annual swing in light-travel time, settled analytically; and **amplitude** modulation, from the changing geometry to a fixed occulter, settled by injection.

Phase. The semi-amplitude of the Earth–Sun distance is $e = 0.0167 \text{ AU} = 8.3 \text{ light-seconds}$; a sinusoidal phase modulation of amplitude φ leaves $J_0(\varphi)$ of the carrier amplitude:

candidate period	phase amplitude	$J_0(\varphi)$	power lost from carrier
120 s (shortest searched)	0.437 rad	0.953	9.2%
300 s	0.175 rad	0.992	1.5%
1,200 s	0.044 rad	0.9995	0.1%
1 day	0.0006 rad	1.0000	< 0.01%

The term is negligible over the whole band searched — 9.2% of carrier power at the shortest period in the paper and under 1.5% above 300 s, because 8.3 light-seconds is small compared with every candidate period.

Is there an annual term in the data anyway? Over $T = 1,945 \text{ d}$ the frequency resolution is $5.95 \times 10^{-9} \text{ Hz}$ and the annual frequency $3.169 \times 10^{-8} \text{ Hz}$, so the sidebands sit **5.33 bins** from the carrier — resolvable, but split across bins and spread by the Blackman mainlobe — and the statistic must sum a block, $R_{\text{(geo)}}(i) = R(i) + R(i \pm 5) + R(i \pm 6)$. A first threshold for that sum assumed $\text{Gamma}(5)$ and over-fired grossly (§5.6). The matched null needs no distributional assumption: a sideband at **any other offset** has identical structure and no geometric meaning, so the exceedance count at the true offset is compared with eight **decoy offsets**. This is a control that can fail: had an annual term been present, the true offset would have stood above the decoys.

channel	above single-bin thr	true offset	decoy median	decoy max	p
25.6 nm	32,609	96,383	110,560	114,439	0.78
28.4 nm	30,944	94,174	104,492	109,101	0.78
30.4 nm	25,651	83,797	89,438	92,417	0.78
117.5 nm	26	141	134	141	0.22
Ly- α	16	105	108	124	0.78
133.5 nm	15	68	79	93	1.00
140.5 nm	24	106	112	121	0.78
Mg II	22	189	181	212	0.44

***Null in all eight channels** — the true-offset count sits at or below the decoy median almost everywhere. The three EUVS-A channels carry the known per-sensor comb (§4.4); the decoy comparison is insensitive to it, which is the point of using a count ratio rather than an absolute count.*

And the method arm. Injecting a carrier with a full-depth annual envelope, the geometric statistic gains 15–25% over the coherent one and changes no detection: both find the signal at $\alpha = 1 \times 10^{-5}$ and both miss at zero. Amplitude

modulation *preserves the carrier* — an envelope of $1 + \cos$ leaves half the power where the coherent search is already looking — and the available phase modulation is 0.44 rad at worst.

Conclusion: the coherent integration of §4.2 is sound, and its limits stand as quoted. The analytic term is 9% at the extreme and under 2% across most of the band; what the exercise buys is that the fast-band limits no longer rest on an unstated assumption of a stationary carrier.

4.12 A different line of sight

Every search above reads the Sun from one place, and §2.7 proposes a modulator that *redistributes* flux along the receiver’s line of sight — line-of-sight-specific by construction. An observer elsewhere sees either nothing or something different, whereas a modulation intrinsic to the Sun is seen by everybody, offset only by geometry. No search so far could tell those apart. This one can.

Earth line against Mars line. GOES-16 EUVS Lyman- α against MAVEN/EUVM diode C, also Lyman- α , over 2019-12-10 to 2025-04-06 — 1,946 days, MAVEN returning data on 94.3% of them. Both series are normalized to 1 AU and shifted to photon-emission time.

On product choice, because it decides the answer. MAVEN’s L3B product is FISM-M *model* output, partly driven by Earth-based inputs; using it would contaminate the non-Earth line with the very line it is compared against and return agreement at every viewpoint — a result manufactured by the choice of file. L2B is measurement, and is what is used. Its cost is cadence: orbit-averaged at 3.66 h, so this test reaches periods above about 7 hours.

The control, which had to come first. Solar rotation must be seen by both platforms at a lag fixed by geometry, $\tau = (\lambda_{\text{(M)}} - \lambda_{\text{(E)}})/360 \times 27.2753$ d, and over this span the Earth–Mars angle sweeps the full 360°.

quantity	value
windows used	86
median peak cross-correlation	0.826
RMS(measured – predicted lag)	1.97 d
shuffled-pairing null	7.65 d
p	5×10^{-4}
anti-geometric control (wrong-sign lag)	7.64 d

The geometry is recovered — the license for everything that follows: a viewpoint null means nothing unless the viewpoint machinery can be shown to find something.

The search. Nine peaks above threshold in the Earth line, seven in the Mars line, six in both. The three Earth-only peaks are reported as unresolved, not detections: MAVEN had only 66%, 20% and 93% power to see a signal of the amplitude observed at Earth, short of the 95% bar — **“absent at the other viewpoint” is worthless unless the other viewpoint could have seen it.**

period	95% recovery	fractional Ly- α	false alarm	status
3 d	0.02 σ	8.9×10^{-4}	0%	limit
7 d	0.08 σ	3.6×10^{-3}	0%	limit

period	95% recovery	fractional Ly- α	false alarm	status
13.5 d	—	—	60%	no limit — detector fires unmodified
27 d	—	—	99%	no limit — detector fires unmodified
60–180 d	—	—	$\leq 6\%$	no limit — red noise, no sensitivity

The limit holds over 2 to about 10 days and nowhere else, and the two failure modes above that are kept apart. At 13.5 and 27 days the circular-roll surrogate preserves the real rotational peak, so the detector fires on unmodified data — the same failure as the sidereal fold in §4.2. Read without the false-alarm column beside it, the 27-day row would have appeared as the strongest limit in this paper.

The result is a null, and it is the first in this paper that constrains **the mechanism §2.7 actually proposes** rather than modulation in general. (An earlier version reported ten Earth-only candidates; all were one uncorrected geometric factor — §5.6.)

The fast band. The slow-band test is limited by cadence, not principle, so it was run again on the MAVEN **L2 band** product at full diode cadence: 1,879 days reduced to one-minute means on the matched 1,945-day grid — coverage 72.6% at Earth, 75.9% at Mars, 54.7% in common (the 66 missing days are the MAVEN safe-mode outage of 2022).

The band has to be capped at 6 hours, and the reason is not conservatism. Solar rotation has a *different synodic period from each viewpoint* — 27.275 d from Earth, 26.354 d from Mars — so rotation-band structure is present in one line and absent in the other **by construction**. Run without the cap, this search duly returned a 26.6-day “Earth-only” candidate at 99.2% Mars power — exactly the signature it exists to find, and entirely a geometric artifact. Periods near and above a day belong to the slow-band test, which carries the rotational lag.

In the judged band, 125 s to 6 h, **the Earth line has no peak above threshold at all**. What the search bounds is the amplitude at which a viewpoint-specific modulation *would* have been seen, measured by injection through the identical pipeline:

period	Earth line, 95%	Mars line, 95%	note
307 s	7.2×10^{-6}	4.2×10^{-3}	both measured
911 s	1.8×10^{-5}	—	Mars has no sensitivity at $\leq 5 \times 10^{-3}\sigma$
3,671 s	—	—	neither line sensitive
3,600 s	—	—	blind by construction — an exact harmonic of the day

The threshold was challenged and it holds. A band-wide control was added on review — the count of exceedances over all 1,336,602 judged bins per surrogate, whose expectation is $\alpha = 0.05$ if the threshold is calibrated — because the previous check tested a *single fixed bin* against a threshold set for 1.3 million and returned $\sim 0\%$ whether the pipeline was sound or not. The first run of the new control returned **197 exceedances per surrogate**, which would have meant the limits here were not family-wise 0.05 and would have loosened them roughly fourfold.

That number was an artifact of the control, not a property of the search. The surrogate rolled the series and then re-applied the gap mask, which on a gapped record moves the existing zero-blocks to new positions and adds zeros back at the original gaps: at MAVEN’s coverage it carried half again the data’s zero fraction, and zero-blocks generate large spectral structure. Rolling only the observed samples, so the gap pattern stays exactly where it is, gives:

line	analytic	empirical	ratio
Earth	24.7	23.8	$1.00 \times$
Mars	24.8	25.7	$1.02 \times$

The analytic threshold $\mu \ln(N/\alpha)$ is calibrated on this data to within 2%, and the limits below stand as printed. §4.2’s argument — that estimating the scale empirically as $\mu = \text{median}(R)/\ln 2$ accommodates the heavy tail — survives a direct test it had not previously been given. The episode is recorded in §5.6 rather than removed: a control that had to be rebuilt twice before it measured the thing it was auditing is worth more as a record than as a silence.

Fractional amplitude at 95% recovery. The Earth-line figure at 307 s is $2.2 \times$ the §4.2 limit at the same period, which is the expected cost of the daily fold and the common-grid interpolation and is the check that this pipeline is calibrated against the paper’s main one.

The binding number is Mars, and it is 580 times worse than Earth. A viewpoint claim requires *both* lines to be able to see the signal, so the limit is set by the weaker: 4.2×10^{-3} at 307 s, against 1.4×10^{-6} for the Earth line alone in §4.2. The limitation is the instrument, not the geometry: EUVM is an orbiter’s monitor, and its one-minute means carry a relative scatter of 0.84 in log space against 3.6×10^{-3} for GOES. **One row is blind by construction and is listed to say so.** The pipeline folds out the mean daily profile, which removes every harmonic of 1/86400 s; 3,600 s is the 24th. Injections there recover nothing at any amplitude — a property of the filter, not of the data — and the quotable rows use periods deliberately incommensurate with the day. (An earlier version injected *after* the fold, exempting the injected signal from a filter a real one would have met; §5.6 records it.)

The next line-of-sight test is an imager summed to an irradiance, not another orbiter monitor. Disc-integrated sums of STEREO/EUVI images (304, 171, 195, 284 Å) are a standard EUV irradiance proxy; STEREO-A has run since 2007, sweeping the full 360° from Earth, at a per-image cadence that reaches the fast band MAVEN L2 cannot; Solar Orbiter EUI/FSI adds a heliolatitude lever. Neither was run here — both need an image-to-irradiance reduction with its own degradation model, and the passband mismatch with GOES 304 is best handled with *platform-internal* ratios such as 304/171, dimensionless per platform. A dedicated disc-integrated photometer off the Earth line remains the instrument this measurement eventually wants; it is neither the only improvement available nor the cheapest.

4.13 A dataset outside the Sun: lunar laser ranging

Lunar laser ranging is the most precisely monitored geometric quantity in the solar system — five retroreflectors, four decades, millimeter normal points — exactly the archive the coverage principle of §5.5 points at. It is treated here, but it is **not one of the 25 searches** and does not enter the tally:

It cannot constrain the modulator of §2.7, and we say so before reporting the nulls. §2.7 puts ranging 2.2×10^{11} short gravitationally and 2.0×10^{12} short radiatively. A null here bounds any unmodelled periodic structure in the Earth–Moon system, but not the architecture this paper is about. It is a method demonstration and an honest negative.

The construction is the interesting part. A single-reflector range is contaminated by station coordinates, atmosphere, Earth orientation and the lunar ephemeris. Ranging to *two* reflectors from the same station on the same night and taking the **ratio of the round-trip times** cancels all of those to first order — and the ratio is dimensionless by construction, satisfying §3.1 with no shared unit. Five reflectors give ten pairs. No dynamical model is used anywhere — published LLR residuals are post-fit against a full ephemeris, so a signal shaped like any fitted parameter is removed before anyone sees it; nothing here is fitted, so nothing is absorbed. The price is that a local detrend is a high-pass filter, so **the detrend window is the coverage claim.**

search	data	pass-band	result
inter-reflector ratio, night cadence	7,630 normal points, 5 reflectors, 4 stations, 2012–2023; 15,329 pairings forming 29 (pair, station) series, of which 24 carry the 60 points the 1.5–19 d band needs under a 40 d detrend and are tested	1.5–19 d 40 d detrend	null — 24 series, Bonferroni $p < 2.1 \times 10^{-3}$, lowest observed 2.5×10^{-3}
within-session residual, full rate	236 sessions, 17,062 returns, Matera and Wettzell	0.5–160 s per-session degree-6	null — global $p = 0.114$; 3 of 400 frequencies above the 99.9th percentile against 0.4 expected
inter-reflector ratio, minute cadence	—	—	not attempted — data-limited, see below

Residual scatter after the per-session polynomial is $244 \text{ ps} = 36.5 \text{ mm}$ one-way, which is ordinary LLR single-shot precision and is the check that the detrend is doing its job rather than eating the signal.

Every near-survivor is a lunar harmonic, and finding that out required fixing the controls. The three lowest p -values land on 2.782 d, 3.357 d and 1.511 d. Against a control list of the lunar months alone, all three look clean; carried to its harmonics they are the anomalistic month over 10 (2.7555 d), the draconic over 8 (3.4015 d) and the half-anomalistic over 9 (1.5308 d). **A 40-day high-pass suppresses the fundamentals and passes exactly these harmonics**, so the harmonics are the relevant controls for this band. None clears the trials threshold in any case. (A parser error earlier in this search forced the withdrawal of a claim that full-rate LLR is background-dominated — §5.6; split correctly, the data were always clean.)

One thing is genuinely non-random. The laser fire epochs are strongly quantized: phase concentration 0.52 against a chance level of 0.0077, grid occupancy $\chi^2 = 1.1 \times 10^5$ on 19 degrees of freedom, 60% of returns in one phase bin — a real periodic measurement artifact, and the expected one (timing hardware). **We have not bracketed its fundamental:** the comb scan rises monotonically to whatever ceiling it is given, 2 kHz and then 12 kHz, so the grid is finer than 83 μs ; a greatest-common-divisor of the epoch differences would pin it, and we report the failure rather than quoting the ceiling as a result.

Why the minute-cadence version was not run. Stations cycle between reflectors within a night, so the same construction should run at minutes. But the open-mirror full-rate holdings are Matera, Wettzell and Beijing; APOLLO, which does cycle reflectors, publishes to a separate archive; and of the nights available, **6 carry more than one reflector and none yields 60 paired points**. The code path would run unchanged on APOLLO data.

4.14 Radio spectral index at one second

§3.6 ranks this channel second and records it as unsearched. It is now searched. RSTN reports eight frequencies per site at one-second cadence, so the ratio of any two is dimensionless by construction — and a redistributor is exactly the thing that would move a spectral index while leaving the total alone. Three sites — Learmonth, Palehua, San Vito — 84 to 89 days each, all 28 pairs per site, band 2.2 to 300 s (the lower bound clear of the 2.000 s Nyquist, the upper set by a 601 s detrend).

The result is null for the carrier, and the way it got there is the interesting part. The pooled search returns 44, 64 and 43 peaks above threshold at the three sites, and the twin gate flags two periods near **5.9 s at more than one site** — precisely the signature the search exists to find.

Learmonth, 5.900 s	R	threshold
channel 610 MHz alone	5.81	2.87

Learmonth, 5.900 s	R	threshold
ratio 245/15400	0.99	2.89
ratio 410/8800	0.98	2.93

The power is in a channel, not in the ratio. A solar modulation of the spectral index moves the ratio; this does not. The ratio peaks the pooled search reported are inherited from the affected channel, and the same pattern holds at the other two sites, where nothing reaches threshold in any ratio at these periods.

And the twin gate is not valid here in any case. RSTN sites are *standardised* receivers — different sites, not independent designs — so a per-design artifact appears at all three exactly as a solar signal would: the failure that defeated the GOES-16/17 gate in §4.4, recurring. Agreement across RSTN sites cannot carry the weight the Oulu/Kiel agreement carries.

Two implementation errors produced a false null — **zero peaks at three sites, twice** — before this search produced a real one; §5.6 records both, and after repair an injected 0.05σ tone is recovered at $R = 4.35$ against a threshold of 2.80.

4.15 Row 6: broadband irradiance, coherently, on SOHO/VIRGO

§3.6 ranked this row sixth by designer preference and recorded it unsearched, on the stated grounds that the archive was not reachable. **That was wrong, and the way it was wrong is worth recording.** PMOD's `ftp.pmodwrc.ch` publishes only an IPv6 AAAA record and does fail from our networks — retried here, three ways, still failing. But NASA mirrors the entire SOHO mission over plain HTTPS. The primary source was dead; the dataset never was. A dead primary source is not an unavailable dataset, and this one blocked three rows of the designer table for four revisions.

VIRGO's three sun photometers (402 nm, 500 nm, 862 nm) give 14,342,400 samples each at 60 s from 1996-01-23 to 2023-04-30 — 27.3 years, 90.5% / 94.3% / 95.4% present. Because the product is already relative, in ppm, the **difference of two channels in ppm is their log color ratio**: the dimensionless carriers of §3.1 come free, with no constructed quantity and no shared denominator.

Gate 1 is unusually strong here and it failed twice before it passed. VIRGO SPM is an instrument solar p-modes are classically measured with, so the five-minute oscillation must come out of a blind search of this data or the search is not calibrated. The first run failed it on two channels — and returned eighteen “candidates”, every one at exactly 180.0000 s or 360.0000 s, which are $\frac{2}{3}$ and $\frac{1}{3}$ of the 60 s Nyquist. **A real signal has no reason to land on a round number in our units.** They are instrument lines reaching $R \approx 19,000$, and they were strong enough to drag the envelope estimator off the p-mode peak. They are excluded a priori, as exact multiples of the sample interval, rather than after inspection of which bins came out large. The second failure was ours: an envelope filter of 2,001 bins is 2.3 μHz against an envelope roughly 1,000 μHz wide (§5.6).

channel	band peak	envelope peak	gate 1
BLUE 402 nm	2.9638 mHz	3.0753 mHz	pass
GREEN	2.9638 mHz	3.0753 mHz	pass
RED 862 nm	2.9638 mHz	3.0753 mHz	pass

Null. With the instrument lines excluded, **no bin of 7,164,744 exceeds threshold in any of the three channels.** The color ratios return two marginal bins each, and all six are identifiable without appeal to anything new: exactly 720.0000 minutes (the 12-hour spacecraft thermal cycle), 4,780,800 minutes (9.1 years, the solar cycle), and 178 and 212 days, which sit against the edge of the filter described below. **No candidate appears in all three photometers**, which is the gate a solar signal would pass and a detector artifact would not.

The limit, and where the search is blind. Injection into the prepared series puts 95% recovery at **0.20 ppm at 307 s**, rising to 0.80 ppm at 1,009 s and 1.60 ppm at 10 h. §3.6 estimated this row's reach at $\sim 10^{-7}$ with margin ~ 1 ; the measured value is 2.0×10^{-7} , **margin 0.5** — the closest any row in this program comes to the level the framework's own logic places a signal at, and null there.

A second instrument, and a correction. VIRGO also flies a PMO6 radiometer, and its one-minute total-irradiance product carries **no highpass at all**, so it reaches a band SPM cannot. Searched identically it is likewise null: every one of twenty candidates above threshold reads $R \approx 1,000$ –2,000 in the radiometer and $R < 5.1$ in all three photometers, and is therefore specific to one detector chain. **This matters as much as the null does**, because the a priori frequency veto had failed on those same candidates — it excluded exactly 180.0000 s and 360.0000 s at a half-width of 0.03%, while the real line carries sidebands to $\pm 2\%$ and a 120 s companion the veto never named. The cross-product gate caught all twenty without needing any of that to be predicted correctly.

The long-period blindness is ours, not the archive's, and an earlier version of this section said otherwise. It was attributed to the L2 two-month highpass. It is not: the search detrends against a one-day running median, which removes everything slower than about a day before the archive's filter is ever approached. Lengthening the detrend to 45 days was tried, and **it did not open the band** — injection recovers 6.40 ppm at one day, 42% at three days, and **nothing at all at 10, 30, 60 or 120 days at any amplitude tested**. The reason is physical rather than procedural: solar variability between roughly three days and two months runs to hundreds of ppm, so a part-per-million signal is far beneath it. **That band is not unsearched; it is unsearchable at this level**, and the same is true of the six-hour timescale in §4.17.

Two restrictions on the null that remain. The L2 product also states correction for orbit, degradation, outliers and “attractors”, and an outlier step can remove exactly the impulsive or quantized structure this statistic is built to detect. Testing that needs an injection through the VDC's own pipeline, which requires L1 data the mission bundles do not carry, and **it has not been done**. Separately, the widened line veto now excludes **12% of the band**, reported here because a notch broad enough to remove every candidate is not a null unless its cost is stated.

4.16 Row 7: p-mode frequency structure, on SOHO/GOLF

GOLF is the better instrument for this row than the EXIS photometry of §4.6, and for a reason that matters to the argument: it measures **Doppler velocity**, a different observable of the same oscillation, and its two photomultipliers PM1 and PM2 put the twin gate **inside the instrument**, where there is no cross-calibration to dispute. 40,845,600 samples at 20 s over 25.9 years at **98.4% duty** — the most continuous record in this program.

What is searched is not whether p-modes exist. That is settled, and here it is the control. A sender able to modulate the Sun at all could shift **mode frequencies**, the most precisely measured quantity in solar physics; the observable is therefore the frequency shift over time, and the question is whether it carries structure beyond what solar activity explains.

gate	quantity	result	
1a	large separation $\Delta\nu$, predicted from $\sqrt{(M/R^3)}$ as 134.9 μHz	135.00 μHz , $r = 0.90$ on all three channels	pass
1b	known solar-cycle frequency shift, ~ 0.4 μHz p-p	0.584 μHz p-p, fractional 1.89×10^{-4} against a literature $\sim 1.3 \times 10^{-4}$	pass
2	PM1 against PM2, 104 segments	$r = 0.982$	pass

Gate 1a is the rare control with no free parameter: $\Delta\nu$ follows from stellar structure and is never fitted. Gate 1b matters for a different reason — **a search for a 10^{-6} modulation that cannot recover a known 10^{-4} one is not a search**, so

the known activity term is an asset rather than a nuisance.

Null. Regressing the shift on F10.7 removes 51.7% of the variance (2.071 nHz per sfu; residual rms 0.0760 μ Hz). The strongest remaining peak sits at 2,455 d with power 0.128 against a permutation threshold of 0.145 — $p = 0.126$ across 200 independent frequencies and 20,000 permutations of the residual. Nothing beyond the activity term.

And the limit corrects this paper's own estimate. Injecting a sinusoid into the shift series **before** the activity regression — so the signal meets the same detrending a real one would, rather than being exempted from it — puts 95% recovery at 0.045 μ Hz, which is 1.46×10^{-5} in fractional mode frequency at $v_{\max}$. §3.6 estimated our reach for this row at 10^{-5} – 10^{-6} with margin 0.1–1; that estimate was explicitly order-of-magnitude, and the measured value gives a margin of **0.069**. **Row 7 is searched, it is null, and it does not reach the level a designer would set.** It falls short by a factor of fifteen.

That figure is what a deliberate search for a better estimator produced, not a first attempt. Ten candidates were ranked — segment lengths of 45, 90 and 180 days, centroid windows of ± 2 , ± 4 and ± 6 μ Hz, and a two-photon multiplier combination — and the ranking was done **on the known solar-cycle term alone**, never on sensitivity, so that the choice could not be tuned toward a nicer limit. The selected estimator (PM1 + PM2, 180-day segments, ± 4 μ Hz) recovers the solar-cycle shift at $1.08 \times$ the literature value and cuts residual scatter by $2.15 \times$, from 0.0760 to 0.0353 μ Hz.

A $2.15 \times$ cut in residual scatter bought only $1.33 \times$ in sensitivity, and the reason is worth stating. Longer segments buy precision per point and pay for it in points: 180-day segments leave 51 where 90-day segments leave 104. A periodic search loses power with fewer samples even when each sample is better, so the two effects largely cancel. **Residual rms is not a proxy for detection sensitivity when the sample count moves**, and an estimator ranked on scatter alone would have been chosen wrongly here.

The shortfall is not mainly method, and this is a revision of what an earlier draft of this section claimed. It was argued here that fitting individual mode profiles rather than cross-correlating band spectra would reach ~ 0.01 μ Hz, six times better, putting the limit near 3×10^{-6} . **That was an extrapolation from published mode-fitting precision and it did not survive being tried.** A ten-candidate estimator search, ranked on the known activity term, delivered $1.33 \times$ — from 1.94×10^{-5} to 1.46×10^{-5} — and the best of those candidates gave the same injected limit as one of the simplest. The remaining factor of fifteen to 10^{-6} is not visibly recoverable from this archive by better estimation, and the honest statement is that **GOLF does not reach the designer level for this row.**

4.17 Row 5: occultation dips and transit timing, and why it is void

§3.6 ranks this row fifth at a level of 10 ppm dips, after Arnold (2005), and records it searched only as a daily single-dip test. VIRGO's one-minute total irradiance — 14,199,837 samples over 27.0 years, no highpass — is the right archive for it.

The control is the strongest available anywhere in this program, because the Sun is transited by known bodies at known times and the depth is not fitted. It is the ratio of two disk areas: Venus $(R/R_{\odot})^2 = 75.6$ ppm, Mercury **12.3 ppm**. Venus transited on 2004-06-08 and 2012-06-06, Mercury on four dates between 2003 and 2019.

Void: none of the six is recovered. A matched filter at the known duration and the known mid-transit time returns Venus 2004 at **–57 ppm** — the Sun was *brighter* across the transit window than the surrounding six days — and Venus 2012 at 164.6 ppm against a null scatter of 108.2. Across all six the SNR spans -1.3 to $+1.5$.

The reason is measured, and it is the Sun rather than the method. A 76 ppm dip over six hours against 62.4 ppm per minute of white noise would be $\text{SNR} \approx 23$, which is why two earlier versions of this gate were rewritten before the answer was believed. **But the noise at six-hour timescales is not white: the null scatter of the same filter on transit-free days is 31–108 ppm**, ten to thirty times the white-noise prediction, because supergranulation and active-region evolution live there. Venus sits at $\text{SNR} \approx 1$ against that.

Two earlier versions of this gate are recorded in §5.6 because the sequence is instructive. The first compared the deepest hour on the transit day against $0.4 \times$ the predicted depth and reported **six of six recovered** — all six numbers were noise, Mercury reading 50–70 ppm against a predicted 12.3. The second added a control-day null and returned **zero of six**, correctly. **An injection would not have caught the first error:** the daily detrend was absorbing the transit, and a signal injected after that step recovers perfectly while a real one does not. Only an anchor with a predicted amplitude could catch it.

No limit is claimed, and the row is reported as void rather than null: a detector that cannot see a real 76 ppm occultation constrains nothing about an artificial one. §3.6's estimate of ~ 1 for this margin was made against a per-minute noise figure, and a transit is not a per-minute event.

4.18 Row 3: sub-minute EUV, with a dark diode

SDO/EVE ESP gives four EUV bands at **0.25 s over 120 consecutive days** — 41,472,000 samples, 97.7% present, Nyquist 2 Hz. That is a band forty years of daily records cannot represent at all, and §4.3's finding that the same pipeline is $245 \times$ more sensitive at two minutes than at one day argues for pushing to the fast end.

The gate is built into the instrument. ESP carries CH_D, a dark diode: same detector chain, same electronics, same telemetry, no photons. Anything appearing there is instrumental by construction, with nothing to model. The five-minute oscillation is recovered in all four light channels (envelope 3.03–3.23 mHz) and **peaks elsewhere in the dark channel**, which is the pairing that makes the control informative rather than merely available.

Null. Of 20,736,000 bins, CH_18 and CH_36 return nothing above threshold; CH_26 and CH_30 return 31 and 30 candidates, **every one of them at exactly 2.000000 Hz — the Nyquist frequency**. Nothing physical sits precisely at Nyquist. The dark channel vetoes nine of the twelve strongest as instrumental; the three it passes are at Nyquist as well, which is a reminder that a control discriminates only against the failures it was built for.

CH_36 is excluded from any limit. Its values cross zero (median 2.6×10^{-4} , minimum -5.2×10^{-4}), so the relative residual $x/\text{trend} - 1$ diverges and its residual rms is 76%. The channel is reported as searched and returning nothing, not as constraining anything, and it is excluded from the ratios below — a ratio against a divergent denominator is a division by something near zero, not a dimensionless carrier.

The dimensionless carriers. §3.1 restricts a sender to dimensionless quantities, so the single channels above are the wrong object and the color ratios are the right one. Because each channel is prepared as a relative residual, the difference of two prepared channels is their log ratio to first order, and the three ratios among CH_18, CH_26 and CH_30 cost nothing extra to form. **All three are null, and every candidate is again at exactly 2.000000 Hz** — 11, 2 and 31 bins above threshold across the three, all at Nyquist. Three of those pass the dark-channel veto, which records only that the dark diode happened to carry less power in those particular bins: **a control discriminates against the failures it was built for, and this one was not built against a sampling edge**. With Nyquist excluded a priori, as it can be from the cadence alone, no candidate survives in any ratio.

4.19 The one spatially resolved channel, and why it is void

Every other search in this paper is disc-integrated. SOHO/VIRGO's Luminosity Oscillation Imager gives **twelve science pixels across the solar disc**, plus four guiding references, at 60 s over **29.0 years with a 97.19% duty cycle** — the best-sampled record used anywhere here.

It was acquired because a bound computed **before** acquisition said it should reach further than any disc-integrated channel for a *localized* source. A feature covering fraction f of the disc at depth d contributes $f \cdot d$ to a disc-integrated series and d to the pixel containing it: for a feature one pixel across that is a signal gain of 12 against a noise cost of $\sqrt{12}$, a net $3.46 \times$, putting the projected limit at 5.8×10^{-8} against the 2.0×10^{-7} measured on SPM. The same arithmetic makes LOI $3.46 \times$ **worse** than SPM for a global modulation, because the photons have been divided for nothing.

Gate 1 passes on every pixel. LOI is a helioseismology imager, so the five-minute oscillation must be present: all twelve return an envelope between 3.025 and 3.175 mHz. The instrument, the data and the pipeline are all behaving.

The rotation control fails, and it takes the test with it. A localized feature crosses the disc in about thirteen days, entering and leaving pixels in a sequence fixed by geometry — so every pixel pair has a predicted lag, and solar rotation supplies the positive control for free, because active regions do exactly this. **One of sixty-six pixel pairs correlates above $r = 0.3$.** The rotational coupling that must be there is not.

The probable cause is in the product, and it was named in the header before the data was downloaded. L2 states correction for orbit, outliers, attractors **and roll sensitivity changes**. SOHO rolls; LOI's pixels are fixed in the instrument frame; removing roll sensitivity plausibly removes the inter-pixel spatial structure with it, leaving twelve near-independent photometers rather than twelve pixels on a rotating disc. **Void:** not for want of sensitivity, but because the observable the test needs appears to have been processed out upstream. Settling it requires L1, which the mission bundles do not carry.

This is the third search in this paper stopped at the same boundary — the others being the outlier-correction question for VIRGO SPM (§4.15) and for EVE ESP (§4.18). In all three the limiting factor is not the instrument or the analysis but a level-2 product having already decided what to remove, and §5.7 should read that as a structural constraint on archival work rather than as three unrelated caveats.

5. Discussion

5.1 What twenty-two nulls establish, and what they do not

A null result bounds an amplitude; it does not settle existence. What §4.2 establishes is that **no dimensionless modulation exceeding the stated amplitude exists in the stated observable over the stated band, during the interval observed** — and every qualifier is load-bearing: a signal below the limit, in a channel not measured, in a band not examined, or absent during the epoch of the archive, is untouched. That is the logical position of every radio search ever published.

What the nulls achieve is to convert speculation into bounded speculation. Before this work, “the Sun’s output carries structure” had no attached number. It now has twelve, eleven injection-verified, and any future version of the claim must place the signal below those limits, outside those channels, or outside those bands — each a substantive and costly restriction.

5.2 Coverage is the binding constraint, and what it actually costs

The natural objection to enumerating the space — of order 1,300 pair tests and 12,000 triple tests under the admissible forms — is cost. The statistical and computational prices scale differently and must be quoted separately.

The statistical price is logarithmic. The Bonferroni threshold falls linearly in the number of tests, but enters the amplitude through the inverse normal:

Quantity	229 tests	12,000 tests	Scaling
Bonferroni threshold	2.2×10^{-4}	4.2×10^{-6}	linear in N
equivalent Gaussian z	3.70	4.61	as $\sqrt{(\log N)}$
amplitude sensitivity lost	—	$1.25 \times$	for a $52 \times$ larger search

Trials penalties are widely feared and, in this regime, mild: a fiftyfold larger search costs 25% in amplitude, against the $245 \times$ same-instrument gain available from cadence alone in §4.3.

The computational price is quadratic, and we established this the hard way. Significance is assessed by counting exceedances among M circular shifts, so the smallest attainable p is $1/(M+1)$; for that floor to reach a Bonferroni threshold of α/N the shift count must satisfy $M > N/\alpha$, and total work is $N \times M = **_N_2/\alpha**$.

We report this because a run of ours failed on it. The thirty-channel sweep’s 1,074 tests have a Bonferroni threshold of 4.66×10^{-5} , while the 6,000 shifts inherited from the 229-test design floor p at 1.67×10^{-4} — **$3.6 \times$ above the line**. The run returned “0 of 1,074 survive”, which reads as a clean null and is nothing of the kind: no test could have survived whatever the data contained. **It is the mirror of the no-op controls of §3.4 — those could not fail, this could not fire** — and it is invisible unless the floor is checked against the threshold explicitly.

Enumeration	Tests	Shifts required	Evaluations	Runtime
13 channels, 78 pairs (first sweep)	229	4,580	1.0M	0.6 min
30 observables, 435 pairs	1,074	21,480	23M	13 min
30 observables, full pair space	1,300	26,000	34M	19 min
30 observables, triples	12,000	240,000	2.9G	26 h

Runtimes at a measured 3.0×10^4 statistic evaluations per second on 80 cores of one machine. The quadratic term is invisible at 229 tests and dominates by 12,000.

The remedy is to stop counting exceedances: a generalized Pareto fit to the upper tail of the shift null gives p -values below the empirical floor from a fixed surrogate budget, returning the cost to linear in N — at 10,000 shifts the full pair space costs seven minutes and the triples about an hour (the triple sweep of §4.9 took 64 minutes for 117M evaluations). **We validated the fit against direct counting on the full pair space before using it on the triples:**

Counted- p band	Tests	Median bias	90% dev
$p > 0.1$ (bulk)	885	− 0.000	0.01
0.01 – 0.1	119	+ 0.007	0.05
0.001 – 0.01	34	+ 0.034	0.13
$10/M - 0.001$	11	+ 0.223	0.46
$p < 10/60001$ (the tail)	25	+ 0.404	2.28

Agreement in dex between the tail-fitted and counted p . Over the resolved range ($n = 1,049$) the median bias is − 0.000 dex, the interquartile spread 0.005, and the rank correlation 0.9999. The fit also does what counting cannot: $|B| / \text{Fe/O}$, pinned at the 1.67×10^{-5} floor by counting, resolves to 6×10^{-14} .

The validation passes on its pre-registered criteria and does not certify the regime it is needed for. Agreement is essentially perfect wherever counting also resolves p — precisely where the fit is unnecessary. In the extreme tail, where the two disagree by up to 2.3 dex, no comparison is possible *because counting has no answer there*; and the two methods disagree on the survivor count, 22 counted against 10 tail-fitted. **A twin-instrument check on the fit itself settles how far it can be trusted.** The identical fit was computed independently on a second machine — different host, core count, SciPy, and random shift draws. In the bulk the runs are indistinguishable: median difference + 0.0000 dex, interquartile spread 0.0064, 90% of deviations within 0.026 dex. **In the tail they are not:** the largest single disagreement is **10.3 dex**, and of ten Bonferroni survivors on each machine **only seven are the same tests**. The extrapolated p in the far tail

depends on the surrogate draw at a level that changes which tests are called significant. This is the control of §4.4 applied to a method rather than to data, and it fired.

The tail fit is therefore a screen and not a verdict, and we use it only that way. Anything it selects is re-run by direct counting at a shift budget sufficient for its own threshold — affordable precisely because so few tests survive the screen.

The conclusion survives the correction, with its margin reduced. **A complete enumeration of the pair space is not an instrumental, funding or telescope-time problem** — it is under an hour on one machine, against archives that already exist; the triple space is 26 h by direct counting and about an hour with a tail fit. And the case for finishing it does not depend on §2 being correct: the limits stand as limits regardless of why one went looking.

5.3 The recognition problem, and the discipline it requires

The gating argument has an obvious hazard: if a message is withheld until a receiver demonstrates capability, *any* null can be attributed to the gate rather than to absence, and the hypothesis becomes unfalsifiable in the way that matters — not by predicting nothing, but by accommodating everything.

§4.10 is what this failure looks like when it is not hypothetical. An alignment-specific excess appeared in the self-keyed search, and the first thing we did was set it aside because its sign was unattractive to the hypothesis — the same failure in the opposite direction: discarding a signal on a criterion the framework had no right to supply. It was caught in review and the result run through the four steps of §3.5. The discipline is only worth stating because we needed it.

We therefore do not invoke the gate to explain the nulls in this paper, and we recommend that no one else does either. Each null was run against a stated carrier at a stated sensitivity; the correct reading is that the carrier is absent above that amplitude, not that a gate intervened. A gating argument allowed to absorb negative results has stopped being a hypothesis and become an excuse.

What the gate is legitimately for is *generating* hypotheses: the dimensionless constraint makes the space enumerable, the relationship prediction directs the search toward combinations, the short pointer directs it toward the fast band — all of which can be looked for and failed to find. The recognition barrier is a real epistemic limit and not a license: it argues for breadth over depth and for reporting coverage, and it has no bearing on whether a given limit is correct.

5.4 Limitations

Limitation	Effect on the conclusions
Adaptive search design. Searches were devised sequentially, some after inspecting data from earlier ones	unrecorded degrees of freedom that no per-search trials correction captures. No search returned an unexplained positive, so no correction is owed — but had one, its significance could not have been assessed. Future work should be pre-registered
Six nulls carry no limit (§4.1)	they constrain nothing, are reported only for completeness, and are excluded from every claim above
Both void results are failures of the null, not of the data (§4.5, §4.9)	neither is fixed by more data, and both are likely to recur wherever a new statistic is introduced without a matched surrogate
Thirty observables is a choice	the denominator of the coverage fraction is set by which quantities are judged independently measured and dimensionless. A different but defensible list moves the 88% of §4.8 by a factor of order unity
One epoch, one star	all solar results describe the last few decades of one star. Nothing here bears on other epochs or other stars

5.5 What to measure next

The fast band, on channels already recorded. The $245\times$ gap between daily and two-minute sensitivity means most existing high-cadence solar data has never been searched at the cadence at which it is most sensitive. This costs nothing but analysis: **SOHO/VIRGO SPM** (three-channel photometry at one minute, unbroken from 1996, the obvious first target), **PROBA2/LYRA** (four bands, nominally 50 Hz, 2010–), **SDO/EVE ESP** (0.25 s, 2010–) and the **RSTN** radio network (**one second** at eight frequencies, 1966–, the longest fast record of any kind). All four were verified fetchable during this work. **RSTN has since been searched at full cadence and is reported in §4.14**; the other three have not.

A neutrino-line search that nobody has run. A search of the public IceCube and Super-Kamiokande catalogs for a monoenergetic line, framed as a test for a *directed artificial* source, is defined, cheap and untried. Present capability is far below the level a designer would set — roughly one Glashow-resonance event per decade at IceCube (IceCube Collaboration 2021), and solar and reactor backgrounds dominating the MeV window — and it closes with Hyper-Kamiokande, JUNO, DUNE, KM3NeT and IceCube-Gen2 on a decade timescale: precisely the shape §3.6 expects of a channel gated on a capability the receiver is still acquiring.

Continuity before novelty. All 53 unreachable pairs are blocked by records that have stopped (§5.7), so the cheapest gain in the pair space is the cross-calibration of an existing successor onto a retired record — Solar Orbiter SWA/HIS onto ACE SWICS above all.

Quantities not measured at all. Four candidate observables have no continuous record over the seventy-year span: core g-modes, the interplanetary electric field, continuous solar polarimetry, and high-latitude solar wind, which existed only while *Ulysses* flew (1990–2009). These are blank columns, not sparse ones; no reanalysis reaches them, and they are the only part of the program that would require new instrumentation.

Because the enumeration is cheap and its value lies in completeness, it is better conducted as a registered, incrementally reported survey — carriers and thresholds declared before the data is touched, coverage published whether or not anything is found — than as a series of individually motivated searches.

5.6 A note on method: physics-supplied controls, void results, and the error record

Three practices used here have nothing to do with technosignatures and may be the most portable result in the paper.

Physics-supplied controls. Wherever possible the control was a place the signal cannot be, fixed by physics rather than chosen by the analyst: above the acoustic cutoff, at the ecliptic pole, at the anti-sidereal frequency, on the twin instrument. Such a control cannot be tuned after the fact, and controls of that kind removed two apparent detections that had passed every amplitude test — including one at 4.9σ whose anti-sidereal counterpart stood at 6.6σ . A statistical null calibrated on the same data would have passed both.

Reporting void results as void. A null from an uncalibrated detector is a statement about the analysis, not about the sky, and recording it as a null silently inflates apparent coverage. Any survey that measures its own coverage — the figure of merit this paper argues for — needs the category.

The error record, collected. The program’s mistakes are collected here because they share one shape — **a statistic or a control not matched to what it was testing** — and because several are the reason a stated null is trustworthy. A reader is entitled to judge how often it happened. None was found by inspection; every one was found by a control that could fail.

Where	Error	Consequence, and how it was caught
§4.2	analytic threshold-crossing amplitudes reported as injection-verified limits	limits optimistic by $1.0\text{--}4.8\times$; caught by running the injections, and the analytic column removed rather than corrected
§4.5	surrogate preserved the quantity the statistic measured, on a cycle-dominated channel	a null that could not have failed; the detector failed its positive control — void

Where	Error	Consequence, and how it was caught
§4.9	circular-shift null applied to three-body statistics	93 apparent survivors, all one functional form; synthetic no-signal triples showed one form $9 \times$ heavy-tailed and the other unable to fire — void
§4.11	five-bin sum of correlated, heavy-tailed bins thresholded as $\Gamma(5)$	63,856 spurious candidates in one channel, firing $3\text{--}5 \times$ too often; replaced by the distribution-free decoy-offset null
§4.12	MAVEN normalized to 1 AU, GOES not	ten “Earth-only candidates” at 1.07, 0.89, 0.76, 0.67, 0.59 and 0.53 yr — Earth’s 6.69% peak-to-peak eccentricity term, smeared by the window, presenting as exactly the signature the search exists to find; resolved into one uncorrected geometric factor
§4.12	injected amplitude divided by $\sqrt{\Sigma_w^2}$ instead of Σ_w ; injection applied <i>after</i> the daily fold	amplitudes inflated $33.6 \times$; a quoted sensitivity the search did not have. Injection path audited, injections moved before the fold
§4.13	CRD headers occur in both cases (329 H4, 256 H4); the parser matched one	merged “sessions” spanning 13.9 h, residual RMS 768 km, full-rate LLR reported background-dominated — withdrawn : split correctly, a degree-5 polynomial already leaves 0.1 m with 99.7% of returns inside 1 ns, and the search uses degree 6
§4.13	lunar-month control list without harmonics	three “clean” near-survivors that are harmonics a 40-day high-pass passes; caught by extending the list
§4.14	accumulator guarded on equal spectrum length; days differ in length	all but the first day of each pair silently discarded: zero peaks at three sites, twice — a false null; caught by a minimum-days audit
§4.14	threshold assumed one periodogram where the statistic is the mean of ~ 85 — $\Gamma(n, \mu/n)$, far lighter-tailed	an unfirable threshold; after repair a 0.05σ tone is recovered at $R = 4.35$ against 2.80
§5.2	p -floor of an inherited 6,000-shift budget above the sweep’s Bonferroni threshold	“0 of 1,074 survive” — a clean-looking null in which no test could have survived; caught by checking floor against threshold
§3.4	two no-op surrogates: phase randomization against a spectrum-derived statistic, IAAFT against a sorted-value statistic	a surrogate comb at 135.3 μHz as significant as the data’s; $p = 1.00000$ with null σ exactly zero. The rule of §3.4 followed

The voids (§4.5, §4.9) are the failure that, caught later, would have been a false limit; the false nulls (§4.14, §5.2) are the same failure inverted. This table is why the paper trusts the limits it does state.

A fourth practice, learned the hard way: an audit must reproduce what it audits.

The three practices above concern controls — a control must be able to return “no”. A separate discipline governs the checks applied to a completed search, and this program learned it by violating it repeatedly in a single afternoon.

The fast-band threshold $\mu \ln(N/\alpha)$ was challenged on review. The challenge was legitimate: the control that had certified it tested a **single fixed bin** against a threshold set for 1.3 million bins, and returned $\approx 0\%$ whether the pipeline was sound or not. Replacing it with a band-wide exceedance count was right. Everything concluded from the replacement was wrong.

attempt	result	why it was wrong
1	threshold $11,287 \times$ too low	surrogate rolled the series and re-applied the gap mask, carrying $1.5 \times$ the data's zero fraction
2	$5,647 \times$	surrogate fixed; day harmonics still counted
3	$19.6 \times$	day-harmonic veto applied at ± 3 bins
4	$1.02 \times$ on MAVEN	correct surrogate, and a pipeline that folds the daily profile out
5	$20.9 \times$ on GOES	same veto — but GOES's eclipse season drifts, so the comb is broadened and the veto caught line centers, not shoulders
6	$35.0 \times$	daily profile folded out, which made it worse

No stable value was obtained, and the reason is structural rather than statistical. A gapped series carries a comb in its *window*: EUVS Lyman- α is 26.9% absent on a daily pattern, and that mask is multiplicative. Subtracting a mean daily profile is additive and cannot remove it; notching requires a veto wider than a comb whose lines drift with the spacecraft's eclipse season. **Every excess measured localized entirely to that comb**, which the search rejects before declaring any candidate.

The limits are therefore left as printed, on three independent grounds: the one channel whose pipeline genuinely removes the daily structure returns a threshold calibrated to 2%; every measured excess sits in bins the search vetoes; and the candidate lists of §A.1 are clean, which a badly miscalibrated threshold could not produce. **That is not a verification — it is the absence of a defensible measurement against.** The distinction is stated because the alternative was to change a printed number on the first of six attempts, which is what happened, and was reverted.

The generalisable rule. Each of the six failures was the same kind: the audit did not replicate the procedure it was auditing. The surrogate was not the same kind of object as the data; the bin set was not the one the search reports on; the veto was narrower than the structure it targeted. None was a statistical error, and twice a diagnosis of one instance was followed immediately by committing the next. **A control must be able to fail; an audit must reproduce what it audits — and the second is harder, because a broken audit produces a number rather than an error.**

5.6b What has not been done, and what it would cost

A coverage claim is only meaningful beside a statement of what was left. Three categories, and only the first is a matter of deciding to.

Runnable now, no new data and no significant compute.

test	why it has not been run	cost
Rows 10–11 — line-profile ratios and disc-integrated polarimetry	below threshold by the §3.6 estimate, and that column has since been measured wrong three times out of three	hours
The remaining 23 scripts migrated to the shared library	each produced a committed number and needs re-validating against it	~1 day
The 52 windows not passed through a transfer check	most are certainly fine; four of the same class cost a day of this program	~1 day

test	why it has not been run	cost
VIRGO TSI minute at the slow end, with a longer detrend	the 45-day pass returned nothing at 10 d and beyond; the band between 3 d and 2 months is measured as unsearchable at ppm level, not merely unsearched	hours

Needs a larger compute budget.

test	scale	what it would buy
T3 — quadruples over 30 channels	27,405 keys $\times$ 4 forms, $\sim$ 18 d on one 88-core machine	nothing, until a defensible four-body null exists. §4.9 and §4.15–19 are explicit that the space is void for want of a null, not for want of compute, and running it would return 27,405 keys' worth of void
T6 — quadruples over 60 channels, fast band	2.8×10^{23} FLOP — 6.2 d on a dedicated exascale machine, 1.3 d on a volunteer network at Folding@home's 2020 peak	the same objection applies, and more strongly
Peak-bagging with full Lorentzian profile fits, GOLF	days, not hours	row 7 sits at 1.46×10^{-5} against a designer level of 10^{-6} . A ten-candidate estimator search bought $1.33 \times$; proper mode fitting is the one untried route to the remaining factor of fifteen, and it may not deliver
N² surrogate counts in place of the tail fit	the reason §3.4 introduced the fit at all	the generalized-Pareto extrapolation is unstable across realizations at 1.7% of triples (§4.9). Counting exceedances instead removes that instability and costs N ²

The compute argument is not the binding one and this table should not be read as asking for a machine. Two of the four entries above are blocked on a null rather than on cycles, and §4.3's central finding is that coverage rather than sensitivity binds. An exascale allocation would buy a larger void.

What the compute actually costs, in three machines. Sustained rates: Cray-1 160 MFLOPS (1976), the 88-core workstation used here $\approx$ 176 GFLOPS, Frontier 1.353 EFLOPS (Rmax).

	FLOP	Cray-1	88 cores	Frontier
triple sweep	5.5×10^{14}	39.8 d	52 min	0.4 ms
quadruple sweep + matched null	2.9×10^{14}	21.0 d	27.5 min	0.2 ms
every search in this paper	1.1×10^{15}	79.6 d	1.7 h	0.8 ms
T3 — quadruples over 30 channels	2.7×10^{17}	53.5 yr	17.8 d	0.2 s
T6 — quadruples over 60, fast band	2.8×10^{23}	5.5×10^7 yr	50,409 yr	57.5 h

The whole program is eighty Cray-1 days and under a millisecond of Frontier. T6 — 55 million Cray-1 years in 1976 — is a long weekend on a flagship machine in 2026. Frontier is 8.5×10^9 times a Cray-1, thirty-three doublings in forty-eight years, almost exactly Moore cadence.

This table is an argument against an allocation, not for one. T6 is not expensive; it is 2.4 machine-days. It is that 2.4 days spent on a question whose null is undefined returns 487,635 keys' worth of void, and §4.9 establishes that the higher-order space is blocked on the null rather than on cycles. **The gate also opens on its own schedule:** T3 is already 0.2 s on a flagship and eighteen days on a workstation. Waiting is cheaper than asking.

5.6c A plan for T3, and the prerequisite that is a research project in its own right

T3 — all 27,405 quadruples of thirty channels, four forms — is eighteen days on one machine and 0.2 s on Frontier.

Neither number is the obstacle. The obstacle is that a four-body search needs a null, and §4.15–§4.19 show that the one correct null available is rejected by the Sun itself at median $z = 45$, because solar output has intrinsic higher-order structure. Running T3 today produces void at scale.

Phase 0 — the null. This is the whole problem. Two routes, and only the second is a compute problem:

- *Targeted rather than omnibus.* Stop asking “is there four-way structure” — the Sun says yes — and ask “is there structure at the specific alignment a sender would impose”, calibrated against decoy alignments. Cheap, and it narrows the hypothesis to something a null can be written for.
- *An ensemble null from physics.* Below.

Phase 1 — forms, and one is already excluded. §4.16's theorem: adding $_c_k s(t)$ to channel k gives a linear four-body form $(_c_a - 3_c_\beta + 3_c_\gamma - _c_\delta)s$ and a pair form $(c_i - c_j)s$, so invisibility to every pair forces all $_c_k$ equal, whereupon $(1 - 3 + 3 - 1) = 0$. **A linear four-body form cannot carry a signal pairs cannot also see.** Only the multiplicative form is a genuine four-body carrier, which cuts the space by three quarters before any compute is spent.

Phase 2 — the sweep. Eighteen days on 88 cores, or under a second on a leadership machine, once Phases 0 and 1 are settled. It is the cheapest part and should be scheduled last.

The HPC project: an ensemble null for solar higher-order statistics

The question, stated without reference to technosignatures. *What is the null distribution of higher-order statistics — bispectra, trispectra, cross-channel phase alignment — of solar output?* Nobody knows, and it is not an idle question: it sets the false-alarm rate of every search for non-linear coupling in solar and heliospheric data, and it is the reason three results in this paper are void.

Why observation cannot answer it. We have one Sun and one realization of it. A null distribution requires an ensemble, and the only ensemble available is a synthetic one.

What the project would be. A large ensemble of independent global solar convective-dynamo simulations — Rayleigh, ASH or MURaM class — each integrated over several simulated activity cycles, from which synthetic disc-integrated irradiance and Doppler-velocity series are extracted through a forward model matched to the instruments actually used (VIRGO SPM and TSI, GOLF, GOES EUVS). The ensemble is the null: the distribution of any higher-order statistic across realizations is what “no imposed signal” looks like for a star that generates its own structure.

Scale. At literature cost of 5×10^5 to 2×10^6 core-hours per realization, an ensemble of 50–200 members is **0.4 to 1.8 million node-hours** — within a single INCITE-class award, and modest by the standards of the simulations themselves.

The gate this project must pass, stated first because it decides whether it is worth running. A null built from simulations that do not reproduce the Sun *in the statistics being tested* is worse than no null: it would license exactly the false confidence this paper spends §5.6 documenting. Global convection simulations are known to struggle with observed large-scale flows — the convective conundrum — and **the ensemble must be validated against the real Sun on the same higher-order statistics before it is used as a null, not after.** If it fails that validation the project still answers a real solar-physics question — *how far do our dynamo models depart from the Sun in their non-linear structure?* — which is worth knowing independently and is arguably the more interesting result.

What it unlocks if it passes. The triple and quadruple spaces become searchable, the two void results of §4.9 and §4.15–19 become limits, and any future search for non-linear coupling in solar data inherits a characterized false-alarm rate. **That is the sequence: the ensemble is the expensive part, and the sweep that follows it is 0.2 seconds.**

Needs instruments that do not exist. Four columns are blank across the whole seventy-year record: core g-modes, the interplanetary electric field, continuous disc-integrated polarimetry, and high-latitude solar wind, the last of which existed only while Ulysses flew. These are unmeasured quantities rather than unsearched archives, and they are the only lever on the roughly one-in-five chance that we happen to measure whichever combination was chosen — a probability no instrument on the observables we already have can improve.

5.7 The receiver’s own position, and why coverage is not simply a matter of time

A gate keyed to capability invites an obvious thought about our own position: we have monitored the Sun for about five decades, the fast band for less than one, and the compute for a combination sweep has existed for perhaps fifteen years. Our capability is still rising steeply, and a null obtained early on a rising curve is worth less than the same null obtained later.

That reasoning is comfortable, and the data does not support the comfortable version of it. Coverage is not a monotonic function of elapsed time. Of the 53 pairs the sweep cannot reach, we asked how many would close if we simply waited:

Of 435 pairs	N	Prospect
reachable now	382	examined in §4.8
unreachable, both records still running	0	would close by waiting
unreachable, a record has stopped	53	cannot close by waiting on that record

*The blocking channels are the Sac Peak Ca II K indices, which end in 2015 and appear in eleven blocked pairs each, and the ACE SWICS charge-state and abundance ratios, which end in 2011 and appear in five each. Successors exist — Solar Orbiter’s SWA/HIS has returned heavy-ion charge states since 2020, and Ca II K synoptic programs continue — but **a successor is not a continuation**: closing these pairs requires cross-calibrating a new record onto a retired one. (Several other channels are flagged as ended only because our retrieved copy stops — total irradiance, F10.7 and the neutron monitors are still produced. Refreshing those would not change the 53.)*

Our capability is not rising in every dimension: it is rising in compute and cadence, and falling in continuity. The last 12% of the pair space is waiting on somebody flying a solar-wind charge-state spectrometer again. **The binding constraint on this space is instrument continuity** — a decade-long gap removes pairs from the reachable set until somebody does the cross-calibration. The same arithmetic leaves 1,126 of 4,060 triples unreachable.

****The compute dimension, quantified.** The paper measured 3.0×10^4 statistic evaluations per second on 80 cores (§5.2), which with the 26-hour triple sweep implies about 3×10^7 FLOP per evaluation on daily-cadence series and about 10^9 on one-minute series. Six tiers of the combination space follow, using direct exceedance counting with the surrogate budget the paper’s own rule requires, $M > N/\alpha$ — the conservative design, since a designer sizing a gate cannot assume the receiver has the tail-fit screen of §5.2, which brings the larger tiers down by 10^3 – 10^4 but must not be used as a verdict. Sixty observables is the thirty of §3.2 plus the thirty candidates named in §5.5 and §3.6; ten forms is the four of §3.3 plus arrangement, timing, profile-ratio and polarization forms. T6 is everything we can currently name, not a claim that the space ends there.

Tier	Space	Tests	Shifts	Evaluations	FLOP
T1	pairs, 30 observables, 4 forms, slow band — done (\$4.8)	1,300	26,000	3.4×10^7	1×10^{15}
T2	triples, 30 observables — done, void (\$4.9)	12,000	240,000	2.9×10^9	9×10^{16}
T3	quadruples, 30 observables	82,000	1.6×10^6	1.3×10^{11}	4×10^{18}
T4	triples, 60 observables, 10 forms, slow band	2.5×10^5	5×10^6	1.25×10^{12}	4×10^{19}
T5	as T4, fast band	2.5×10^5	5×10^6	1.25×10^{12}	1.25×10^{21}
T6	quadruples, 60 observables, 10 forms, fast band — the enumerable maximum	3.75×10^6	7.5×10^7	2.8×10^{14}	2.8×10^{23}

Machine	Year	Sustained FLOP/s	T1	T2	T3	T4	T5	T6
Cray-1	1976	8×10^7	148 d	34 yr	1,500 yr	15,000 yr	5×10^5 yr	1×10^8 yr
Cray X-MP/4	1984	4×10^8	30 d	7 yr	309 yr	3,000 yr	1×10^5 yr	2×10^7 yr
Cray Y-MP/8	1988	1.4×10^9	9 d	2 yr	92 yr	880 yr	29,000 yr	7×10^6 yr
TMC CM-5	1993	1.8×10^{10}	16 h	56 d	7 yr	66 yr	2,200 yr	5×10^5 yr
ASCI Red	1997	3.2×10^{11}	53 min	3 d	141 d	4 yr	123 yr	28,000 yr
Earth Simulator	2002	1.1×10^{13}	2 min	2.2 h	4 d	40 d	4 yr	824 yr
BlueGene/L	2005	8.4×10^{13}	12 s	17 min	13 h	5 d	172 d	106 yr
Roadrunner	2008	3.1×10^{14}	3 s	5 min	3.5 h	1 d	47 d	29 yr
K computer	2011	3.2×10^{15}	< 1 s	28 s	21 min	3.3 h	5 d	3 yr
Tianhe-2	2013	1.0×10^{16}	< 1 s	9 s	6 min	1 h	1 d	319 d
Summit	2018	4.5×10^{16}	< 1 s	2 s	1 min	14 min	8 h	72 d
Fugaku	2020	1.3×10^{17}	< 1 s	1 s	29 s	5 min	2.6 h	24 d
Frontier	2022	3.3×10^{17}	< 1 s	< 1 s	12 s	2 min	1.1 h	10 d
El Capitan	2024	5.2×10^{17}	< 1 s	< 1 s	7 s	1 min	40 min	6 d
this paper, 80-core workstation	2025	9×10^{11}	19 min	1 d	50 d	1 yr	44 yr	9,900 yr
10 ⁵ -H100 cluster, FP64	2024	2.0×10^{18}	< 1 s	< 1 s	2 s	19 s	10 min	2 d

Machine	Year	Sustained FLOP/s	T1	T2	T3	T4	T5	T6
4 × 10 ⁵ -Blackwell cluster, FP64	2026	4.8 × 10 ¹⁸	< 1 s	< 1 s	1 s	8 s	4 min	16 h

Machine rates are TOP500 Rmax, or FP64 peak for GPU clusters, × 0.3 for an FFT-heavy statistic (× 0.5 of peak for the Cray era). GPU clusters are quoted at double precision; the FP8 and BF16 figures used for AI training are 30–100 × larger and do not apply to this workload. FLOP per evaluation is calibrated from one workstation and one statistic; a third-order surrogate of the kind §4.9 calls for could cost 10 × more per evaluation, and the tiers move together. Compute is necessary and not sufficient: every error in §5.6 was made with adequate compute, and the gate is passed by a receiver that can search and validate.




Figure 5. The compute gate against real machines. Left: sustained FLOP/s of the leading supercomputer of each era (circles), two large GPU clusters at FP64 (squares), and the workstation used in this paper (triangle, measured); dashed lines mark the rate at which each tier completes in one year. Right: wall-clock time to close each tier on eight of those machines, with one day, one year, a forty-year career and 5,000 years marked. The pair space crossed the one-year line with the Cray X-MP, the triple space with the Y-MP, and the enumerable maximum between Tianhe-2 and Summit in the mid-2010s; T6 is a working day on a current GPU cluster and ten thousand years on the machine that produced this paper, which is why rows 3–7 of §3.6 remain unsearched.

None of this is offered as an account of the nulls, and §5.3 forbids using it that way.** What it bears on is the *weight* a reader should give a coverage figure: **a coverage fraction that cannot be improved by waiting is a different quantity from one that can**, and this paper should not have quoted the first as though it were the second.

5.9 Scored against the nine axes of merit

Sheikh (2020) organizes technosignature searches along nine axes — four functions of us, five of the technology sought — and the framework was developed at the same 2018 workshop whose taxonomy §1.1 sets this work against. Scoring against it is the honest way to answer *where does this sit*, and two of the answers are “worse”.

axis		
Cost	very high	Twenty-nine searches, no new observations, one workstation. No telescope time was requested and none is needed
Observing capability	high, with a correction	The data exists now — but of §3.6’s estimated reaches, all three that have since been measured came in short: row 5 unreachable, rows 6 and 7 at margins 0.5 and 0.069 against an estimated ~ 1
Ancillary benefits	moderate	The sweeps recover real heliophysics unprompted, and two bands are now measured as unsearchable at part-per-million level rather than merely unsearched. The durable product is methodological
Detectability	mixed	Deep absolutely — 1.4×10^{-6} in fractional Lyman- α at two minutes, 2.0×10^{-7} in broadband irradiance. Marginal against the level the gating argument predicts: the best margin achieved anywhere is 0.5
Duration (<i>L</i>)	very high	A passive modulator needs no power, no consumables and no maintenance (§2.7). A beacon must be <i>run</i> ; this must only persist, and <i>L</i> could be geological
Ambiguity	high in principle, poor in practice	A dimensionless relationship among independently measured observables is not something natural processes have reason to impose. But see below
Extrapolation	moderate	$\sim 1.5 \times 10^{12}$ m ² of film for one part per million: far beyond us, but passive structure rather than exotic physics
Inevitability	low	The weak axis, and we concede it rather than argue it
Information content	very high	An infrared excess reports that something is there; this architecture carries a message by construction (§2.5)

The two concessions, stated plainly.

Inevitability is the weakest axis and cannot be argued away. This architecture requires a civilization to want to signal, to prefer a capability gate to a broadcast, **and** to choose a combination we happen to measure. §2 argues each is the *efficient* choice; efficient is not inevitable. Radio leakage and waste heat score higher because they require no intent at all, and the arithmetic of §B.1 puts the chance that we measure a chosen pair at about one in five. What is offered in exchange is that the search costs nothing — the archives are already collected — so a low prior is affordable here in a way it would not be in a proposal for observing time.

Ambiguity is limited by the analysis, not by the physics, and §5.6 is the evidence. The claim that a dimensionless carrier is hard for nature to imitate is true and is not the binding constraint. In the course of this program four distinct artifact families were identified — instrument lines carrying sidebands sixty times wider than the veto written for them, solar-cycle harmonics, harmonics of the analysis’s own detrend window, and the sampling Nyquist — and three were unanticipated. Every control that failed was a control incapable of returning “no”. **The pipeline generates candidate structure faster than the sky does**, and a search of this kind is bounded by how well that is policed rather than by how exotic the carrier is.

6. Conclusions

We have asked what a technosignature search looks like when the channel is not known in advance — no dedicated radiating apparatus assumed, the carrier possibly a relationship among observables. Two constraints make the question finite: a sender sharing no units can encode only in dimensionless quantities, and a sender minimizing energy gates the message on demonstrated capability. Against the resulting combination space we ran 29 searches of public archives, none of it data collected for SETI.

	Principal result
1	Twenty-five nulls, three void, one detection. Thirteen limits are injection-verified; one sensitivity remains analytic. The deepest verified limit is 1.4×10^{-6} in fractional Lyman- α irradiance at two minutes, at 95% recovery
2	Cadence dominates sensitivity. The same instrument, days and pipeline are $245 \times$ more sensitive at two minutes than at one day (the limits table as a whole spans 10^5 , but across different observables and statistics). Most high-cadence solar data has never been searched at the cadence at which it is most sensitive
3	Coverage, not sensitivity, is the binding constraint. Seventeen further observables took the pair space from 18% to 88% and found nothing across the Sun–heliosphere boundary in 510 tests. The remaining 12% of pairs are blocked by retired records and cannot close by waiting (§5.7); 2,749 of the 2,934 reachable triples completed and returned void for want of a matched three-body null (§4.9)
4	Completing the enumeration is affordable, but not linearly. A $52 \times$ larger search costs only 25% in amplitude; the compute scales as N^2 because the surrogate count must scale with the test count. The full pair space is under an hour on one machine, the triple space 26 h by direct counting or about an hour with a tail-fitted null (§5.2)
5	Physics-supplied controls earn their place. Controls that cannot be tuned — acoustic cutoff, ecliptic pole, anti-sidereal frequency, twin instrument — removed two apparent detections that had passed every amplitude test, including one at 4.9σ whose anti-sidereal counterpart stood at 6.6σ
6	Both voids are failures of the null, not of the data. A surrogate that preserves what the statistic measures, or a two-body null applied to a three-body form, constrains nothing. Neither is repaired by more data
7	The first line-of-sight test is null. Earth against Mars in Lyman- α : no viewpoint-specific modulation to 8.9×10^{-4} at 3 d and 4.2×10^{-3} at 307 s, the latter limited $580 \times$ by the MAVEN monitor and not by the geometry (§4.12). The first result here that bounds the mechanism §2.7 proposes rather than modulation in general

The framework's own best guess is now tested, and it does not reach. Of §3.6's unsearched rows 3–7, **four are searched here**. Row 6 is null at an injection-verified 2.0×10^{-7} , **margin 0.5** — the closest this program comes to a designer level. Row 7 is null at 1.46×10^{-5} , **margin 0.069**. Row 3 is null. Row 5 is **void**: the six-hour noise floor is 31–108 ppm and a 76 ppm Venus transit is not recoverable. **All three rows whose reach was estimated and has since been measured came in worse than the estimate**, and the §3.6 margin column should be read as an upper bound on capability rather than as capability

The one detection is a re-detection: solar p-modes in GOES EXIS Mg II irradiance, comb spacing 135.1 and 135.0 μHz on the two spacecraft against an accepted 134.9. Priority belongs to Eden et al. (2024), and we claim none. Its value is as a positive control whose answer was fixed in advance and corroborated by an independent group — a stronger check on the pipeline than any self-designed injection, and it correspondingly strengthens the twenty-five nulls.

We do not claim that the framework of §§2–3 is supported by these results, and we do not invoke it to explain them. A gating hypothesis that absorbs negative results has stopped being a hypothesis. The nulls are reported as nulls: each carrier was searched at a stated sensitivity and was not found. What the framework is credited with here is generating the searches, not surviving them.

The limits stand independently of the argument that motivated them. A reader who rejects §§2–3 entirely is still left with eleven injection-verified bounds and one analytic sensitivity on dimensionless modulation of solar output, at cadences and in combinations not previously examined — and with the observation that the archives are far more informative at two minutes than at one day, which is a fact about solar physics and not about SETI.

Three things would advance the problem, in ascending order of cost. **Search the fast band on channels already recorded**, in the order §3.6 gives: VIRGO/SPM and TIM at minute cadence, first as a transit-timing-sequence search after Arnold (2005) and then coherently at 10^{-7} ; sub-minute EUV from PROBA2/LYRA and SDO/EVE ESP; RSTN spectral indices at one second; and BiSON and GOLF mode frequencies. **Complete the pair enumeration and begin the triples**, pre-registered — completeness is only credible when it is declared in advance. **Measure the five quantities that have no continuous record at all**: core g-modes, the interplanetary electric field, continuous solar polarimetry, high-latitude solar wind, and a disc-integrated irradiance monitor off the Earth line (§4.12).

Wright et al. (2018) established that the value of a SETI result lies in the fraction of a defined space it excludes; this paper applies that standard to a space whose axes are quantities rather than pointings. By that measure the honest summary is not that nothing was found, but that **the pair space is now essentially closed and the triple space was attempted and returned void** — closing the pairs required no new instrument, only seventeen archives that already existed — that **the first viewpoint test bounds the mechanism §2.7 proposes rather than modulation in general** (§4.12), and that the framework's own ordering names where the next searches belong: in archives that exist, at sensitivities already achieved, and not yet examined (§3.6).

Appendix A. All 29 searches

Twenty-nine searches, one row each: **null** (no detection), **void** (the detector or its null failed a positive control, so the row constrains nothing), **detection**. Numbering is the order in which the searches were run within the wider program, retained for traceability to code and data; the grouping is by domain.

Why each search ended where it did. “Null” is not one thing, and a reader deciding where to spend effort next needs the distinction. Across the twenty-nine, the reasons sort into six:

failure mode	n	what it means	where
nothing above threshold	22	the detector worked, the control passed, and the sky was empty at the stated level. This is a result	most rows
underpowered	3	the search ran and could not have found the effect even if present — power measured at 36–53%	§A.2 rows 6, 13, 14
detector failed its own control	1	the statistic could not fire, or fired on unmodified data. Constrains nothing	§4.5
no defensible null	1	the surrogate is correct and the Sun rejects it, so “no signal” cannot be specified	§4.9
the Sun is too loud	1	the physical noise floor sits above the level required. Not fixable by analysis	§4.17
the archive removed the observable	—	a level-2 product had already decided what to discard; outside the twenty-nine, the observable having gone before a search could be scored	§4.19
detection	1	a known signal recovered as a positive control	§4.6

Only the first is an astronomical statement. The next three are statements about the analysis, and the last two about the data. That distinction matters more than the tally: twenty-five nulls sound like broad coverage, and the honest reading is that twenty-two channels were searched competently, three searches were too weak to count, two were defeated by their own machinery, and one was defeated by the data.

A.1 Solar, fast band (14)

#	Search	Data	Population	Result	Why	Null quality / limit
17	Solar p-modes, EXIS MgII	GOES-16 + 17, 1-min	2 spacecraft, 5.3 yr	detection	known signal, positive control	$\Delta\nu = 135.1$ (G16) and 135.0 (G17) μHz , accepted 134.9 — the program’s positive control
18	EUVS 1-min, seven lines	GOES-16 + 17, 8 channels	1.4M bins x 8, 5.3 yr	null	nothing above threshold	Lyman-alpha to 1.4 ppm at 2 min, 95% recovery — deepest verified limit here
19	Multi-scale, six time bases	Oulu 1-min, 26 yr	12 tests, 13.7M samples	null	nothing above threshold	both designs excluded above 0.5%; both detectors calibrated

#	Search	Data	Population	Result	Why	Null quality / limit
20	Neutron monitor fast search	Oulu + Kiel, 1-min	6.8M bins, 26 yr	null	nothing above threshold	163 candidates, 0 survive the twin; 0.0063%
21	Blind narrowband, 1-min X-ray (cf. Hippke & Forgan 2017)	GOES-16 + GOES-17 XRS-B	2.5M bins, 9.68 yr	null	nothing above threshold	0 non-diurnal candidates; 0.027% at 2 min
22	EXIS line ratios, twin-gated	10 channels, one instrument	135 tests, 5.3 yr	null	nothing above threshold	9 survivors on one spacecraft, 0 replicate on two
23	Achromatic spectral modulation	7 UV lines + XRS-B	972 d, frozen filter	null	nothing above threshold	< 5% achromatic modulation excluded at 99%
24	Coronal hardness (photon parameter)	XRS-A / XRS-B	2,233 d, 3 spacecraft	null	nothing above threshold	occultation-vs-emission discriminator; noise-limited
25	Instrument-handover control	GOES-15 vs GOES-16 X-ray	1,031 dual days	null	nothing above threshold	detector output shifts 27.7σ across the handover; the shift-null p-value does not move
30	Radio spectral index, 1 s	RSTN Learmonth, Palehua, San Vito	84–89 d each, 28 pairs per site	null	nothing above threshold	no excess in the dimensionless carrier; the multi-site 5.9 s peaks are per-channel and the RSTN twin gate cannot separate a per-design artifact (§4.14)
31	Broadband irradiance, coherent	SOHO/VIRGO SPM blue + green + red, 60 s	7.16M bins x 3, 27.3 yr	null	nothing above threshold	0 bins above threshold on any channel; 0 pass the three-photometer gate (§4.15)
32	p-mode frequency structure	SOHO/GOLF Doppler velocity, 20 s	25.9 yr, 98.4% duty; PM1 + PM2, 180 d segments	null	nothing above threshold	1.46×10^{-5} fractional at 95% recovery — 15 × short of the designer level (§4.16)
33	Occultation dips, transit timing	SOHO/VIRGO TSI, 60 s, 27.0 yr	6 known planetary transits	void	see §4.5 / §4.9 / §4.17 / §4.19	the detector cannot recover Venus at 76 ppm; six-hour noise floor measured at 31–108 ppm (§4.17)

#	Search	Data	Population	Result	Why	Null quality / limit
34	Sub-minute EUV, four bands + dark diode	SDO/EVE ESP, 0.25 s, 120 d	20.7M bins x 4, Nyquist 2 Hz	null	nothing above threshold	all candidates at exactly Nyquist; p-modes recovered in light, absent from dark (§4.18)

A.2 Solar, slow band (10)

#	Search	Data	Population	Result	Why	Null quality / limit
2	Cross-channel coherence	TSI + ACE + OMNI, 8 yr	4 channels, 728 days	null	nothing above threshold	FP 3.5%, 100% power at 0.40σ
3	Self-keyed spread spectrum	ACE MAG + F10.7, 8 yr	15.3M field samples	null	nothing above threshold	FP 6.3%, 99.7% power at 0.10σ ; alignment excess, resolved in §4.10
4	Dimensionless combination sweep	13 solar/heliospheric channels	229 tests, 13–44 yr	null	nothing above threshold	126 tests across the Sun–heliosphere boundary, none survive
48	Thirty-channel pair sweep	30 observables, 5 new archives	1,074 tests, 382 of 435 pairs	null	nothing above threshold	0 of 510 across the Sun–heliosphere boundary; all 22 survivors known physics or shared-instrument
49	Thirty-observable triple sweep	30 observables, 2 three-body forms	10,996 tests, 2,934 of 4,060 triples	void	see §4.5 / §4.9 / §4.17 / §4.19	neither form has a calibrated null — one is $9\times$ heavy-tailed, the other cannot fire
6	Sunspot cycle sequence	24 cycles, 1755–2019	7 statistics	null	nothing above threshold	underpowered: misses lag-1 = 0.4 two times in three
13	Cycle-resolved 14C, 976–1894	Usoskin 2021 + SILSO	85 cycles, 36 reliable	null	nothing above threshold	29 of 85 unidentifiable; power only 36% → 53%
14	Cosmogenic sequence, 11,350 yr	Solanki + Usoskin	1,113 decadal points	null	nothing above threshold	no structure beyond spectrum + distribution, two reconstructions
27	Occultation dips, quiet epochs	LASP TSI + SILSO, 46 yr	4,919 quiet days	null	nothing above threshold	FP 7.2%, excludes 30-day dips > 62 ppm

#	Search	Data	Population	Result	Why	Null quality / limit
28	Aperiodic structure in solar output	LASP TSI, 46 yr	16,801 daily samples	void	see §4.5 / §4.9 / §4.17 / §4.19	invalid — wrong channel, detector fails control

A.3 Frame & geometry (4)

#	Search	Data	Population	Result	Why	Null quality / limit
7	Sidereal fold, EXIS	GOES-16 + 17, 3 channels	1-min, 5.3 yr	null	nothing above threshold	sidereal 0.2σ against a 246σ solar artifact
15	Sidereal fold, anti-sidereal gated	Oulu 1-min, 26 yr	1440 bins, 9,530 folds	null	nothing above threshold	sidereal 4.9σ but anti-sidereal 6.6σ — leakage, not sky
16	Barycentric-frame sky scan	Oulu 1-min, 26 yr	60 directions, 3.2M bins	null	nothing above threshold	no direction beats topocentric; artifacts die $14\times$ pole to plane
29	Cross-viewpoint coherence, Earth vs Mars	GOES-16 EUVS + MAVEN/EUVM L2B Ly- α	1,946 d, 94.3% MAVEN coverage	null	nothing above threshold	geometry recovered at $p = 5\times 10^{-4}$; 8.9×10^{-4} fractional at 3 d, no limit above ~ 10 d (§4.12)

A.4 Outer solar system (1)

#	Search	Data	Population	Result	Why	Null quality / limit
26	Radial coincidence, 20–160 AU	Voyager 1 + 2, 48 yr	32k hourly samples	null	nothing above threshold	FP 3.0%, 100% power at 50% injection

29 rows: 25 null, 3 void, 1 detection. Seventeen of the twenty-five nulls state a sensitivity — **13 injection-verified amplitude limits, 3 measured false-alarm-and-power pairs and 1 analytic sensitivity** (the sidereal fold) — and the other eight are enumeration or control results that bound structure without quoting an amplitude. Per-search method, code and data provenance are in the supplementary material.

Appendix B. The datasets

A bird’s-eye inventory of every archive processed in this program: who collects it, when it starts, how many samples it holds, at what cadence, and — where an archive was not used — why.

Every row count, start date, end date and duty cycle below was measured from the files on disk on 2026-09-12, not quoted from documentation. Provenance (instrument, agency, archive) is from the acquisition code and the Data Availability table. Where a figure could not be measured it is left blank rather than estimated.

B.1 At a glance

distinct archives drawn on	19
independently measured channels	30 daily + 7 EUV bands + 3 photometric + 3 velocity + 6 ESP
total time samples held locally	~275 million
longest record	daily sunspot area, 1874–2016 (142 yr)
longest continuous high-cadence record	SOHO/GOLF, 25.9 yr at 20 s, 98.4% duty
finest cadence	SDO/EVE ESP, 0.25 s
all data public	yes — no proprietary, embargoed or author-collected data

What has been asked of them, and what came back

outcome	datasets	meaning
detection	GOES-16 + 17 EUVS	solar p-modes, $\Delta\nu = 135.0/135.1$ μHz against 134.9 predicted — a known signal, recovered as the program's positive control
null, limit stated	15 searches	amplitude bounded; 11 verified by injection. Deepest: 1.4 ppm in Lyman- α at 2 min
null, no limit	7 searches	constrains nothing; reported anyway so coverage is not overstated
void	NANOGrav step search; TSI aperiodic; triple sweep; quadruple sweep	the detector or the null failed its own validation. A null from an uncalibrated detector is a statement about the analysis, not the sky
not yet searched	—	all acquired archives have now been searched

Four voids, and they are not all the same failure. Two are detector failures. Three are *null* failures — the circular shift assumed pairwise structure was irrelevant, the common-phase surrogate assumed stationarity, the envelope surrogate assumed the nonlinearity was pure amplitude modulation, and all three assumptions are false. The distinction matters to anyone picking this up: a void from a bad null is repaired by a better control, never by more data.

B.2 Daily solar and heliospheric channels

These thirty form the grid the combination space is enumerated over. All are resampled to a common daily grid spanning 1980-01-01 to 2025-12-31 (16,802 days); “rows” is the number of days actually present, so the fraction present is rows/16,802.

channel	instrument / collector	archive	first	last	rows	cadence	searched in	outcome
TSI	composite radiometry	LASP LISIRD	1980-01-01	2023-12-30	16,070	1 d	§4.8 pair, §4.9 triple; searches 27, 28	null (pair); dip bound 62 ppm; search 28 void

channel	instrument / collector	archive	first	last	rows	cadence	searched in	outcome
F10.7	Penticton 2.8 GHz radio telescope	NRCan / NOAA	1980-01-01	2023-12-30	16,050	1 d	§4.8 pair, §4.9 triple, search 3	null
sunspot number	visual counts, global network	SILSO, Royal Obs. Belgium	1980-01-01	2025-12-31	15,209	1 d	§4.8 pair; searches 6, 13	null , both underpowered
Mg II core-to-wing	SBUV/GOME/SCIAMACHY composite	LASP LISIRD	1980-01-01	2013-07-15	12,109	1 d	§4.8 pair, §4.9 triple	null
cosmic ray	Oulu neutron monitor	NMDB	1980-01-01	2023-12-30	16,048	1 d	searches 15, 16, 19, 20	null ; 163 candidates, 0 survive twin
IMF B	multi-spacecraft merged	NASA GSFC/SPDF OMNI2	1980-01-01	2023-12-30	14,321	1 d (from 1 h)	§4.8 pair, §4.9 triple	null ; the four solar-wind pairs that survive are textbook heliophysics
proton temperature	”	”	1980-01-01	2023-12-30	16,070	1 d (from 1 h)	§4.8 pair, §4.9 triple	null ; the four solar-wind pairs that survive are textbook heliophysics
wind density	”	”	1980-01-01	2023-12-30	14,044	1 d (from 1 h)	§4.8 pair, §4.9 triple	null ; the four solar-wind pairs that survive are textbook heliophysics
wind speed	”	”	1980-01-01	2023-12-30	14,146	1 d (from 1 h)	§4.8 pair, §4.9 triple	null ; the four solar-wind pairs that survive are textbook heliophysics
alpha/proton ratio	”	”	1980-02-19	2023-12-30	12,541	1 d (from 1 h)	§4.8 pair, §4.9 triple	null ; the four solar-wind pairs that survive are textbook heliophysics
X-ray background 1–8 Å	GOES XRS	NOAA NCEI	1983-05-19	2025-04-05	15,185	1 d	§4.8 pair; searches 24, 25	null ; handover control passed
proton flux > 10 MeV	GOES SEM	NOAA NCEI	1985-12-31	2019-12-30	12,401	1 d	§4.8 pair	null
Lyman-α	GOES EUVS	NOAA NCEI	2006-07-03	2025-04-05	4,866	1 d	§4.8 pair; search 18	null , 1.4 ppm at 2 min
sunspot area	RGO + USAF/NOAA network	NASA MSFC Greenwich	1980-01-01	2016-10-31	13,454	1 d	§4.8 pair, §4.9 triple	null

channel	instrument / collector	archive	first	last	rows	cadence	searched in	outcome
hemispheric asymmetry	”	”	1980- 01-01	2016- 10-02	11,860	1 d	§4.8 pair, §4.9 triple	null
WSO mean field	Wilcox Solar Observatory magnetograph	Stanford WSO	1980- 01-01	2024- 03-19	13,065	1 d	§4.8 pair, §4.9 triple	null
Ca II K emission index	Sacramento Peak spectroheliograph	LASP LISIRD	1980- 02-21	2015- 09-30	4,042	1 d	§4.8 pair, §4.9 triple	null
Ca II K2V/K3	”	”	1980- 02-21	2015- 09-30	4,042	1 d	§4.8 pair, §4.9 triple	null
Ca II K3	”	”	1980- 02-21	2015- 09-30	4,042	1 d	§4.8 pair, §4.9 triple	null
Ca II ΔK1	”	”	1980- 02-21	2015- 09-30	4,042	1 d	§4.8 pair, §4.9 triple	null
CME rate	SOHO/LASCO coronagraph	CDAW universal catalog	1996- 01-11	2025- 12-31	9,951	1 d	§4.8 pair, §4.9 triple	null
CME mean speed	”	”	1996- 01-11	2025- 12-31	9,943	1 d	§4.8 pair, §4.9 triple	null
O ⁷⁺ /O ⁶⁺	ACE SWICS 1.1	ACE Science Center	1998- 02-04	2011- 08-21	4,783	1 d	§4.8 pair, §4.9 triple	null ; charge states implausible as a carrier
C ⁶⁺ /C ⁵⁺	”	”	1998- 02-04	2011- 08-21	4,783	1 d	§4.8 pair, §4.9 triple	null
mean Fe charge	”	”	1998- 02-04	2011- 08-21	4,782	1 d	§4.8 pair, §4.9 triple	null
mean Si charge	”	”	1998- 02-04	2011- 08-21	4,799	1 d	§4.8 pair, §4.9 triple	null
mean O charge	”	”	1998- 02-04	2011- 08-21	4,783	1 d	§4.8 pair, §4.9 triple	null
Fe/O	”	”	1998- 02-04	2011- 08-21	4,782	1 d	§4.8 pair, §4.9 triple	null
He/O	”	”	1998- 02-04	2011- 08-21	4,783	1 d	§4.8 pair, §4.9 triple	null
C/O	”	”	1998- 02-04	2011- 08-21	4,783	1 d	§4.8 pair, §4.9 triple	null

Why the grid starts in 1980 even though several records run far longer: the grid is the intersection window chosen so that the majority of channels are live. Sunspot area and sunspot number extend to 1874 and 1700 respectively and are truncated here, not lost. **Retired records are the binding constraint on pair coverage** — ACE SWICS ended 2011, Ca II K 2015, sunspot area 2016, Mg II 2013 — which is why 12% of the pair space cannot be closed by waiting.

B.3 High-cadence solar — the fast band

dataset	instrument / collector	archive	start	rows	channels	cadence	searched in	outcome
GOES-16 EUVS 1-min	EXIS/EUVS, GOES-R	NOAA NCEI	2018-12-14	2,800,800	7 bands + Mg II	60 s	searches 7, 17, 18, 22, 23, 29	DETECTION (p-modes, positive control); else null
GOES-17 EUVS 1-min	”	”	2018-12-21	1,614,240	7 bands + Mg II	60 s	searches 7, 17, 18, 22	DETECTION replicated on 2nd spacecraft
GOES-16 XRS 1-min	EXIS/XRS	”	2016-01-01	5,091,840	2 (A, B)	60 s	searches 21, 24, 25	null ; 0 non-diurnal candidates
GOES-17 XRS 1-min	”	”	2017-01-01	3,155,040	2 (A, B)	60 s	searches 21, 24, 25	null ; twin gate
Oulu neutron monitor	ground NM64	NMDB	—	13,675,680	1	60 s	searches 15, 16, 19, 20	null ; 0.0063%
Kiel neutron monitor	ground NM64	NMDB	—	13,675,680	1	60 s	search 20 (twin gate)	null
Oulu, hourly	”	”	—	6,522,048	1	1 h	\$4.8 pair	null
Kiel, hourly	”	”	—	6,522,048	1	1 h	\$4.8 pair	null
SDO/EVE ESP	EUV SpectroPhotometer	LASP	2014-001	41,391,240	5 + dark	0.25 s	\$3.6 row 3, \$4.18	NULL ; all candidates at Nyquist; dark diode clean

EUVS bands: 25.6, 28.4, 30.4, 117.5, 121.6 (Lyman- α), 133.5, 140.5 nm. **GOES-16 and GOES-17 overlap for 1,121 days**, which is what makes the twin-instrument gate possible: a candidate must appear on both spacecraft. EVE ESP carries **CH_D**, a **dark channel** processed through the identical pipeline as a null control.

B.4 Helioseismology — SOHO

dataset	instrument / collector	archive	first	last	rows	cadence	duty	searched in	outcome
GOLF MEAN	resonant-scattering spectrophotometer, Doppler velocity	NASA SOHO archive (IAS)	1996-04	2022-02-28	40,845,600	20 s	98.4%	\$3.6 row 7, \$4.16	NULL , 1.46×10^{-5} at 95% recovery — $15 \times$ short of the designer level

dataset	instrument / collector	archive	first	last	rows	cadence	duty	searched in	outcome
GOLF PM1	photomultiplier 1	"	"	"	40,845,600	20 s	98.4%	§3.6 row 7, twin gate	NULL; PM1 vs PM2 $r =$ 0.982
GOLF PM2	photomultiplier 2	"	"	"	40,845,600	20 s	98.4%	§3.6 row 7, twin gate	NULL; twin gate passed
VIRGO SPM BLUE (402 nm)	sun photometer	NASA SOHO archive (IAC/VDC)	1996-01- 23	2023-04- 30	14,342,400	60 s	90.5%	§3.6 row 6, §4.15	NULL, 0.20 ppm at 307 s (margin 0.5)
VIRGO SPM GREEN	"	"	"	"	14,342,400	60 s	94.3%	§3.6 row 6, §4.15	NULL; three- photometer gate passes nothing
VIRGO SPM RED	"	"	"	"	14,342,400	60 s	95.4%	§3.6 row 6, §4.15	NULL; three- photometer gate passes nothing
VIRGO TSI minute	PMO6 radiometer	"	1996-01	2023-02- 20	14,199,837	60 s	—	§3.6 row 5, §4.17	VOID — 6 h noise floor 31– 108 ppm; Venus not recovered

Two features make this group unusually well controlled. **GOLF PM1/PM2 are two photomultipliers inside one instrument**, so the twin-detector gate needs no cross-calibration. And **GOLF measures Doppler velocity while VIRGO measures photometry** — genuinely different observables of the same oscillation, so a candidate can be required to appear in both.

Two caveats that must travel with any limit from VIRGO SPM. The L2 header states a *seven-degree polynomial fit and a two-month highpass filter*, so this product is **blind by construction to modulation slower than ~2 months**. It also states correction for *orbit, degradation, outliers and “attractors”* — and an outlier step can remove exactly the impulsive or quantized structure these searches look for. The first is a hard sensitivity limit; the second is testable by injection and must not be assumed either way.

VIRGO SPM also carries strong instrument lines at exactly 180.0000 s and 360.0000 s (3 and 6 minutes, being 2/3 and 1/3 of the 60 s Nyquist), reaching $R \approx 19,000$. They are identifiable a priori as exact multiples of the sample interval and are excluded from the search band.

Rows 6 and 7 both pass all their gates (§4.15, §4.16). The comb came out at $\Delta v = 135.00 \mu\text{Hz}$ against **134.9 predicted from $\sqrt{(M/R^3)}$** with $r \approx 0.90$ on all three GOLF channels; the known solar-cycle frequency shift came out at $0.584 \mu\text{Hz}$ peak-to-peak (fractional 1.89×10^{-4} against a literature $\sim 1.3 \times 10^{-4}$); and PM1 and PM2 agree at $r = 0.982$ over 104 segments. Both gates initially **failed on bugs of my own** — a continuum filter of 501 bins = $0.61 \mu\text{Hz}$ against a mode linewidth, which flattened the modes it was meant to normalize against and returned $r = 0.004$, and an unguarded parabolic interpolation that produced shifts of order $10^{29} \mu\text{Hz}$. Row 6 (VIRGO SPM) still fails gate 1 on GREEN and RED for the same reason in a different place: a 2,001-bin envelope filter is $2.3 \mu\text{Hz}$ against an envelope $\sim 1,000 \mu\text{Hz}$ wide. Its

search is clean — zero of 7,164,744 bins above threshold on every single channel, gate 2 passing nothing — but **no null is claimed until the gate passes**.

The failure class is worth naming. Three times in one session: a smoothing or masking window chosen without checking it against the scale of the feature it operates on. §5.6 records one instance (“a re-implementation of the mode-comb detector returning 127 μHz where the original gave 135”) as a single incident. It is a class, and it is the most productive thing the gates have caught.

B.5 Beyond the Sun

dataset	instrument / collector	archive	rows	cadence	searched in	outcome
NANOGrav 15 yr	68 millisecond pulsars, GBT + Arecibo + VLA	Zenodo 16051178	2,278	irregular	timed spin-up step search	VOID — 4 nulls tried, none calibrated; the 10^{-14} limit was withdrawn
Lunar laser ranging	APOLLO, OCA, McDonald; Apollo/Lunokhod retroreflectors	80 files	—	irregular	within-session + inter-reflector	null ; an h4/H4 case bug forced withdrawal of an earlier claim
RSTN radio	US Air Force solar radio network, 4 stations	NOAA NGDC	273 files, 80 MB	1 s / 5 s	search 30, spectral index	null ; twin gate cannot separate a per-design artifact
Voyager 1 & 2	plasma wave + trajectory, 20–160 AU	NASA/JPL	—	—	search 26, radial coincidence	null ; FP 3.0%, 100% power at 50%
^{14}C , cycle-resolved	tree-ring radiocarbon	VizieR J/A + A/649/A141	—	~11 yr	searches 13, 14	null , but underpowered (36%→53%)

B.6 Cross-viewpoint

dataset	instrument / collector	archive	cadence	searched in	outcome
MAVEN EUVM L2b	EUV Monitor, diode C (Lyman- α), Mars orbit	PDS PPI mirror, UCLA	orbit-merged	search 29, Earth vs Mars	null ; geometry recovered at $p = 5 \times 10^{-4}$, 8.9×10^{-4} at 3 d
Wind MFI L2	magnetic field magnitude, L1	NASA GSFC/SPDF	daily medians	search 3 follow-up (§4.10)	null ; resolved the alignment excess

The level-3b MAVEN product is deliberately not used — it is FISM-M *model output*, not a measurement. The PDS PPI mirror at UCLA is used rather than the LASP SDC because it serves **9,000 days/hour against 40** — a $245\times$ difference that turned a week of acquisition into an afternoon.

B.7 Archives not used, and why

Stating these matters: a coverage claim is only meaningful if what is missing is named.

archive	wanted for	status	reason
SOHO/VIRGO	§3.6 rows 5–6	✓ NOW AVAILABLE	recorded blocked for four revisions on a <i>wrong inference</i> . PMOD's <code>ftp.pmodwrc.ch</code> genuinely is IPv6-only and genuinely fails from here — retried 2026-09-12, all three variants. But NASA mirrors the entire mission over plain HTTPS . The primary source was dead; the dataset never was
SOHO/GOLF	§3.6 row 7	✓ NOW AVAILABLE	same mirror. Supersedes BiSON and GONG for this row, and is a <i>different observable</i> (velocity, not photometry)
BiSON	row 7	× blocked	data portal returns 403/404 on every path; only Zenodo record is a PDF
GONG	row 7	× needs a human	interactive query form, serves full-disk FITS rather than a disk-integrated series
PROBA2/LYRA	§3.6 row 3	× unreachable	(EVE ESP serves row 3 instead)
LASP LISIRD, sub-daily irradiance	rows 5–6	× not served	nothing faster than six-hourly
Amateur-radio propagation archive	withdrawn	⚠ removed	10.1 M spots, but not public , so not reproducible; and the medium is HF ionospheric propagation, whose diurnal structure is itself the dominant signal

The generalisable lesson, which cost four revisions: a dead primary source is not an unavailable dataset. Rows 5, 6 and 7 sat marked “capable but unsearched” across four drafts on the strength of one true fact and one wrong inference drawn from it. Before recording an archive as blocked, check for a mirror.

B.8 Notes for a researcher picking this up

What the coverage actually is. Pair space is 88% closed (382 of 435 pairs); the residual 12% is limited by *record overlap* between retired instruments, not by sensitivity, and cannot close by waiting. Triple space returned void. Quadruple space returned void for a different and sharper reason.

The three-clause regularizer is the load-bearing constraint on which channels may be combined: a candidate quantity must be dimensionless, **independently measured**, and must not share an instrument, spectrum or denominator with its partner. Plasma β and the Alfvén Mach number were dropped under this rule despite being dimensionless — they are *functions* of density, temperature, speed and field, all already in the set, and testing $|B|$ against β tests $|B|$ against a formula containing $|B|$. Including them produced sixteen spurious Bonferroni survivors. Mg II from EXIS is excluded from EUVS band combinations for the same reason.

Cadence is the single biggest lever on sensitivity. The same instrument, the same days, the same pipeline is **245× more sensitive at two minutes than at one day**. Any future work should start at the fast end, which is why the 0.25 s ESP and 20 s GOLF records are the most valuable rows in this table despite being the shortest.

Software

Tool	Use
Claude Code (Anthropic)	analysis environment: wrote and ran every search pipeline, designed and repaired the surrogate tests, executed the enumerations ($\sim 10^8$ statistic evaluations), retrieved and parsed the archives, derived the analytic results of §2, produced the figures, drafted the manuscript
Python 3.10, NumPy, SciPy	numerics; <code>scipy.ndimage</code> for running medians, <code>scipy.stats</code> for the generalized-Pareto tail fit
pyhdf	ACE SWICS level-2 Vdata tables
cdflib	ACE MAG, Wind MFI and MAVEN EUVM CDF files
Astropy	heliocentric ephemerides of Earth and Mars for the viewpoint geometry of §4.12

On the scale of the tooling, and its reliability. The program as reported is closer in scale to what a small team would undertake than to one investigator’s hand-work, and a reader assessing this many results is entitled to know how they were produced. Claude Code is a tool and not an author: Robert Griffin is responsible for every number here, and that responsibility is not nominal. Errors made by the tool and recorded in this paper include most of the §5.6 table, a statistic that sat at 0.954 for data and surrogates alike and so could not fail (§4.9), and a re-implementation of the mode-comb detector returning 127 μHz where the original gave 135. **Each was found by checking rather than by inspection:** a result is worth what its controls are worth, and a pipeline that produces limits faster than they can be checked produces limits that have not been checked. **The audit reported here was run against commit 0aa39cd**, and it checked 20 quoted numbers against the result files that produced them: 20 reconcile, and the one manuscript error it found — a triple count that gave the reachable figure where the completed one was meant — is corrected in §4.9, restated at submission against whatever commit is submitted; the repository README carries the claim-to-script table, and two of its checks execute directly from a clean clone.

Acknowledgments

The analysis, figures and manuscript were produced using Claude Code (Anthropic); its scope and documented failures are set out under Software. This work used public archives maintained by NOAA/NCEI, NASA GSFC and JPL, the ACE Science Center, LASP, the Royal Observatory of Belgium, NRCan, NMDB, the Wilcox Solar Observatory, the CDAW Data Center and CDS/VizieR. **None of these data were gathered with SETI in mind** — which is the point of §1.5.

Data availability

Every result in this paper derives from a public archive. No proprietary, embargoed or author-collected data is used, and no observation was taken for this work.

Archive	Products used	Sections
NOAA NCEI, GOES-R EXIS	XRS 1-minute and daily background (<code>xrsf-l2-avg1m</code> , <code>xrsf-l2-bkd1d</code>); EUVS 1-minute and daily (<code>euvs-l2-avg1m</code> , <code>euvs-l2-avg1d</code>), GOES-16 through 19	§4.2–4.6, A.1
NOAA NCEI, GOES 13/14/15	EUVE daily irradiance, 2006–2016; X-ray background; > 10 MeV integral proton flux	§3.2, A.1–A.2
NASA GSFC / SPDF	OMNI2 hourly solar wind and interplanetary field, 1980–2023	§3.2, A.2

Archive	Products used	Sections
SILSO, Royal Observatory of Belgium	international sunspot number, daily and cycle-resolved	§3.2, A.2
NRCan / NOAA	Penticton F10.7 solar radio flux	§3.2, A.2
NMDB	Oulu and Kiel neutron monitors, 1-minute and 5-minute	§4.2, A.1
LASP LISIRD	total solar irradiance composite; Mg II core-to-wing index	§3.2, A.2
VizieR	Usoskin et al. 2021 cycle-resolved ^{14}C (J/A + A/649/A141)	A.2
NASA/JPL	Voyager 1 and 2 trajectory and plasma-wave data, 20–160 AU	A.4
LASP MAVEN SDC	MAVEN/EUVM level-2b orbit-merged band irradiances (<code>mvn_euv_l2b_orbit_merged</code>), diode C (Lyman- α), 1 AU-normalized. The level-3b product is FISM-M model output and is deliberately not used (§4.12)	§4.12, A.3
NASA GSFC / SPDF, Wind	MFI level-2 magnetic field magnitude (<code>wi_h0_mfi</code>), reduced to daily medians	§4.10
ACE Science Center	SWICS 1.1 level-2 daily charge states and abundance ratios (<code>ssv4_data_1day</code>), quality-flag filtered	§3.2, §4.8
NASA MSFC	RGO + USAF daily sunspot areas, hemispheric, 1874–	§3.2, §4.8
Wilcox Solar Observatory	mean line-of-sight solar magnetic field, 1975–	§3.2, §4.8
CDAW / SOHO LASCO	universal CME catalog; daily rate and mean linear speed	§3.2, §4.8
LASP LISIRD (Sac Peak)	Ca II K emission index, K2V/K3, K3, $\Delta K1$	§3.2, §4.8

Per-search provenance — exact file versions, retrieval dates and the quality flags applied — is given in the supplementary material, together with the analysis code and the injection curve underlying each limit.

One dataset was removed. An earlier version included three searches of a privately collected amateur-radio propagation archive (10.1 million spots). They are withdrawn: the archive is not public, so the searches were not reproducible, and the medium is HF ionospheric propagation, whose diurnal and seasonal structure is itself the dominant signal. **Removing them costs no measured limit** — all three fell among the nulls that carried none.

References

Aasi, J., et al. (LIGO Scientific Collaboration) 2015	<i>Class. Quantum Grav.</i> 32 , 074001 — Advanced LIGO
Abbott, B. P., et al. 2016	<i>Phys. Rev. Lett.</i> 116 , 061102 — Observation of gravitational waves from a binary black hole merger
Amaro-Seoane, P., et al. 2017	arXiv:1702.00786 — Laser Interferometer Space Antenna
Arnold, L. F. A. 2005	<i>ApJ</i> 627 , 534 — Transit light-curve signatures of artificial objects

Ball, J. A. 1973	<i>Icarus</i> 19 , 347 — The zoo hypothesis
Banaszek, K., et al. 2025	arXiv:2501.13356 — Communicating at a record 14.5 bits per received photon through a photon-starved channel
Benford, G., Benford, J., & Benford, D. 2010	<i>Astrobiology</i> 10 , 475 — Messaging with cost-optimized interstellar beacons
Benford, J. 2019	<i>AJ</i> 158 , 150 — Looking for lurkers: co-orbiters as SETI observables
Biswas, A., et al. 2018	<i>Proc. IEEE ICSOS 2017</i> , 23 — Status of NASA's deep space optical communication technology demonstration
Boroson, D. M., et al. 2014	<i>Proc. SPIE</i> 8971 , 89710S — Overview and results of the Lunar Laser Communication Demonstration
Bracewell, R. N. 1960	<i>Nature</i> 186 , 670 — Communications from superior galactic communities
Brin, D. 2014	<i>JBIS</i> 67 , 8 — The search for extraterrestrial intelligence (SETI) and whether to send 'messages' (METI): a case for conversation, patience and due diligence
Chennamangalam, J., et al. 2015	<i>New Astronomy</i> 34 , 245 — Jumping the energetics queue
Cocconi, G., & Morrison, P. 1959	<i>Nature</i> 184 , 844 — Searching for interstellar communications
Cohen, M., Wheaton, W. A., & Megeath, S. T. 2003	<i>AJ</i> 126 , 1090 — Spectral irradiance calibration in the infrared. XIV. The absolute calibration of 2MASS
Davies, P. C. W., & Wagner, R. V. 2013	<i>Acta Astronautica</i> 89 , 261 — Searching for alien artifacts on the moon
DeVito, C. L., & Oehrle, R. T. 1990	<i>JBIS</i> 43 , 561 — A language based on the fundamental facts of science
Eden, T. D., et al. 2024	<i>ApJL</i> 973 , L18 — Solar atmospheric oscillations as measured by GOES-R EXIS EUVS-C
Farr, W. H., Choi, J. M., & Moision, B. 2013	<i>Proc. SPIE</i> 8610 , 861006 — 13 bits per incident photon optical communications demonstration
FCC 2026	Space Bureau, DA 26-113, 4 Feb 2026 — public notice accepting for filing SpaceX's application for a non-geostationary system of up to one million satellites (Orbital Data Center system)
Freitas, R. A. 1980	<i>JBIS</i> 33 , 251 — A self-reproducing interstellar probe
Freudenthal, H. 1960	North-Holland — <i>Lincos: design of a language for cosmic intercourse</i>
Gertz, J. 2016	<i>JBIS</i> 69 , 31 — Reviewing METI: a critical analysis of the arguments
Guardiani, A., et al. 2024	<i>Proc. SPIE</i> 12877 , 128770S — Superconducting nanowire single-photon detectors for laser communication
Hao, H., et al. 2024	<i>Light Sci. Appl.</i> 13 , 25 — A compact multi-pixel superconducting nanowire single-photon detector array supporting gigabit space-to-ground communications
Haqq-Misra, J., Busch, M. W., Som, S. M., & Baum, S. D. 2013	<i>Space Policy</i> 29 , 40 — The benefits and harm of transmitting into space
Hemmati, H. (ed.) 2006	Wiley — <i>Deep Space Optical Communications</i> (JPL Deep-Space Communications and Navigation Series)
Hippke, M., & Forgan, D. H. 2017	arXiv:1712.06639 — Interstellar communication VI: searching X-ray spectra for narrowband communication

Hippke, M., & Learned, J. G. 2018	arXiv:1802.02180 — Interstellar communication. IX. Message decontamination is impossible
Hippke, M. 2017a	arXiv:1706.03795 — Interstellar communication I: maximized data rate for lightweight space-probes (<i>Int. J. Astrobiology</i>)
Hippke, M. 2017b	arXiv:1712.05682 — Interstellar communication V: introduction to photon information efficiency
Hoang, T., Lazarian, A., Burkhart, B., & Loeb, A. 2017	<i>ApJ</i> 837 , 5 — The interaction of relativistic spacecrafts with the interstellar medium
Hogben, L. 1952	<i>JBIS</i> 11 , 258 — Astraglossa, or first steps in celestial syntax
Hyper-Kamiokande Proto-Collaboration 2018	arXiv:1805.04163 — Hyper-Kamiokande design report
IceCube Collaboration 2021	<i>Nature</i> 591 , 220 — Detection of a particle shower at the Glashow resonance with IceCube
Jiang, J. H., et al. 2022	<i>Galaxies</i> 10 , 55 — A Beacon in the Galaxy: updated Arecibo message for potential FAST and SETI projects
KISS 2019	Data-driven approaches to searches for the technosignatures of advanced civilizations (arXiv:2308.15518)
Learned, J. G., et al. 2008	arXiv:0809.0339 — The Cepheid galactic internet
Learned, J. G., Pakvasa, S., & Zee, A. 2009	<i>Phys. Lett. B</i> 671 , 15 — Galactic neutrino communication
Moision, B., & Hamkins, J. 2005	<i>IPN Progress Report</i> 42-161 — Coded modulation for the deep-space optical channel: serially concatenated pulse-position modulation
Ollongren, A. 2013	Springer — <i>Astrolinguistics: design of a linguistic system for interstellar communication based on logic</i>
OpenAI 2026a	Planar point sets with many unit distances (proof document, 20 May 2026); see also Quanta Magazine, 3 Aug 2026, on the Erdős problems
OpenAI 2026b	announcement of a Lean-formalized blow-up solution to the Navier–Stokes problem by $\sim 10^4$ coordinated agents, 8 Sep 2026; unverified by the community and subject to a priority dispute at the time of writing (Science, Nature, 8–9 Sep 2026)
Perakis, N., & Hein, A. M. 2016	<i>Acta Astronautica</i> 128 , 13 — Combining magnetic and electric sails for interstellar deceleration
Rivest, R. L., Shamir, A. & Wagner, D. A. 1996	MIT/LCS/TR-684 — Time-lock puzzles and timed-release crypto
Sagan, C., Sagan, L. S., & Drake, F. 1972	<i>Science</i> 175 , 881 — A message from Earth
Sawin, W. 2026	arXiv:2605.20579 — An explicit lower bound for the unit distance problem
Sheikh, S. Z. 2020	<i>Int. J. Astrobiology</i> 19 , 237 — The nine axes of merit for technosignature searches
Silagadze, Z. K. 2008	<i>Acta Phys. Polon. B</i> 39 , 2943 — SETI and muon collider
SpaceNews 2026	J. Foust, “SpaceX files plans for million-satellite orbital data center constellation,” 31 Jan 2026
Staff at the National Astronomy and Ionosphere Center 1975	<i>Icarus</i> 26 , 462 — The Arecibo message of November, 1974
Tipler, F. J. 1980	<i>QJRAS</i> 21 , 267 — Extraterrestrial intelligent beings do not exist
Willmer, C. N. A. 2018	<i>ApJS</i> 236 , 47 — The absolute magnitude of the Sun in several filters

Wright, J. T., Kanodia, S., & Lubar, E. 2018

AJ **156**, 260 — How much SETI has been done?

Zubrin, R. M., & Andrews, D. G. 1991

J. Spacecraft Rockets **28**, 197 — Magnetic sails and interplanetary travel

Data references (Solanki 2004; Usoskin et al. 2021; SILSO; NMDB; NOAA/NCEI GOES-R; ACE Science Center; LASP LISIRD) are given in the data availability statement.

Draft for review. Analysis pipelines, surrogate tests and the claim-to-script audit are in the repository named in the Software section. Charts on this page are drawn from the tables in the text; Figures 2–4 are reproduced in the repository.