Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

# A Bayesian framework for posterior bounds on self-replicating probes in the Solar System under observational incompleteness

**Jiawei Pan**
Soochow University The Third Affiliated Hospital of Soochow University
**Anqi Yang**
Soochow University The Third Affiliated Hospital of Soochow University
**Qiutao Gu**
Soochow University The Third Affiliated Hospital of Soochow University
**Yuan Zhang**
Soochow University The Third Affiliated Hospital of Soochow University
**Linglong Jiang**
Soochow University
**Jian Shi**
shijian201205@163.com

Soochow University

**Research Article**

**Keywords:** SETA, technosignatures, Bayesian inference, Fermi paradox, Solar System surveys, self-replicating probes

**Posted Date:** July 8th, 2026

**DOI:** https://doi.org/10.21203/rs.3.rs-10268655/v1

**License:** This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

**Additional Declarations:** The authors declare no competing interests.

# A Bayesian framework for posterior bounds on self-replicating probes in the Solar System under observational incompleteness

Jiawei Pan, Anqi Yang, Qiutao Gu,Yuan Zhang, Linglong Jiang , Jian Shi *

Soochow University, The Third Affiliated Hospital of Soochow University, Changzhou, China

*Corresponding author.

Email address:shijian201205@163.com (Jian Shi)

## Abstract

Searches for extraterrestrial artefacts inside the Solar System have so far reported only non-detections. We turn this body of nulls into a quantitative, falsifiable posterior bound on the expected number $\lambda$ of self-replicating probes currently resident in the Solar System. Building on Bayesian Drake-style frameworks, we model $\lambda$ with a log-uniform parametric prior, and combine it with a single-detection-trial Poisson likelihood $L(\lambda) \propto e^{-\lambda p_{det}}$ whose effective detection probability $p_{det}$ is reconstructed from the published completeness of seven Solar-System regions (near-Earth objects, Earth co-orbitals, Earth–Moon Lagrange points, lunar surface , main belt , Kuiper belt, and the more distant Solar System), together with three explicit residence-weight scenarios. For 10 m probes under the default outer-dominated residence model, today's data give a 95% upper bound $\lambda < 14.3$, tightening to 3.16 at 100 m and 0.68 at 1 km, where the prior is already nearly saturated. Forecasting the same framework forward with mission-stated completeness targets for the Vera C. Rubin Observatory Legacy Survey of Space and Time, the Near-Earth Object Surveyor, and the Roman Space Telescope, the 10 m default-residence bound compresses by 7.6× over the decade 2025–2035 (±50% envelope: 4.8×–10.0×), while the 1 km default-residence bound moves by only ∼ 14–16%. The headline contribution is not a number for "the probability that probes exist", but a reproducible, data-driven framework that turns survey completeness into a posterior constraint, makes its prior assumptions auditable, and identifies sub-100 m discovery – not km-scale surveys – as the channel through which the next decade can move the bound.

Keywords: SETA, technosignatures, Bayesian inference, Fermi paradox, Solar System surveys, self-replicating probes

## 1. Introduction

The "Fermi paradox" (Hart, 1975; Tipler, 1980) – the absence of any in-Solar-System artefact of an extraterrestrial intelligence (ETI) that one would expect to encounter under a wide range of assumptions about the prevalence and longevity of technological civilisations – has been re-examined repeatedly in the last decade (Sandberg et al., 2018; Benford, 2019, 2021; Lingam et al., 2023; Grimaldi, 2023a,b; DeVito, 2024; Crawford, 2026). Two distinct strands have emerged. The first, motivated by Bayesian treatments of the Drake equation, asks how parameter uncertainty should be propagated through the chain $R_\star \to f_p\, n_e \to f_l \to f_i \to f_c \to L \to \ldots$ when each factor spans many orders of magnitude (Sandberg et al., 2018). The second, motivated by the more concrete question "should we be detecting their stuff?", catalogues plausible search targets within the Solar System – Earth co-orbital "lurkers" (Benford, 2019),

artefacts on the lunar surface (Davies and Wagner, 2013; Sander and Wöhler, 2025), Lagrange-point bodies (Freitas and Valdes, 1980; Valdes and Freitas, 1983; Gregg and Wiegert, 2022), and resolved imaging of Solar-System small bodies in general (Anton and Schetsche, 2023; Villarroel et al., 2026) – and asks whether the systematic non-detection of any of these is itself informative.

What is missing is a single, end-to-end inference that (i) consumes the same parametric prior used in Bayesian Drake analyses, (ii) converts the heterogeneous published completeness information for each Solar-System region into a per-object detection probability, (iii) combines the two into a posterior distribution on the number of probes resident in the Solar System, $\lambda$, and (iv) projects forward how the bound will evolve as the major upcoming surveys – the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) (Kurlander et al., 2025; Wagg et al., 2025; Chow et al., 2026), the NASA Near-Earth Object Surveyor (NEOS) (Mainzer et al., 2023; Grav et al., 2023), and the Nancy Grace Roman Space Telescope (Holler et al., 2025) – reach their published completeness targets in the next decade. Our purpose is to construct this inference, run it on the published data, and release a transparent open-source implementation that can be re-used and audited.

Three commitments shape the rest of the paper. First, we report a posterior upper bound on $\lambda$ rather than a point estimate of any "probability of probes". The upper bound is falsifiable – a future positive detection of even one artefact would violate it – and it is data-driven in the limited sense that the choice of prior is fully explicit, the likelihood is the single-trial Poisson zero-detection formula, and every numerical input is sourced. Second, every completeness number is anchored to a published survey study, with the provenance recorded row by row in the released dataset. Where a published eficiency curve had to be inverted to a per-object detection probability we explain the inversion in this paper and in the dataset's PROVENANCE.md. Third, the framework's sensitivity to its prior is explored both globally (through the prior 95% upper limit on $\lambda$ at fixed $p_{det}$, Section 6) and locally (through a ±50% envelope on the surveys' published completeness targets, Section 7). We use the words "posterior bound" and "detection limit" throughout to keep this expectation explicit.

The remainder of the paper is organised as follows. Section 2 reviews the four Bayesian frameworks that most directly motivate our construction (Sandberg et al., 2018; Lingam et al., 2023; Grimaldi, 2023a,b) and contrasts the targets they constrain. Section 3 presents the Poisson likelihood and the two-level decomposition of $p_{det}$ used to combine the seven regions. Section 4 documents the empirical completeness inputs and their provenance. Section 5 describes the Monte Carlo implementation, random-number-generator seeding, and convergence diagnostics. Section 6 reports the posterior bound at three canonical probe sizes (10 m, 100 m, 1 km) under three residence scenarios. Section 7 forecasts how the 2025–2045 schedule of LSST, NEOS, and Roman compresses the bound, with a ±50% envelope on the published eficiency targets. Section 8 discusses limitations and the prior-dominance regime that the framework's identifiability map makes visible. Section 9 concludes.

## 2. Background and previous frameworks

Four bodies of work bear most directly on this paper. We summarise each, contrast their inferential targets with ours, and identify the specific increment we provide.

**The parametric Drake equation under uncertainty..** Sandberg et al. (2018) (hereafter SDO) replace the Drake-equation point estimates with full distributions, sample them in Monte Carlo, and demonstrate that the apparent paradox dissolves once the parameters are propagated as ranges rather than expectations. Their result is about N, the steady-state number of technological civilisations in the Galaxy. We adopt SDO's log-uniform priors as our parametric prior on the Galactic substrate, then push the inference one step further by asking what the prior implies for $\lambda$, the Solar-System probe abundance, and combining it with Solar-System completeness data.

**The Drake equation for artefacts..** Benford (2019, 2021) rewrite the Drake equation specifically for self-replicating probes ("lurkers"), introducing factors that represent the probability that probes are launched, that they remain functional, and that they would be observable from their preferred residence locations – Earth co-orbital orbits in Benford (2019), generalised in Benford (2021). The framework is qualitative in the sense that no specific completeness data are pushed through it. We retain its location-by-location decomposition (Section 3) and combine it with a quantitative Poisson likelihood.

Bayesian technosignature assessment.. Lingam et al. (2023) develop a Bayesian decision-theoretic framework to rank candidate technosignatures by the strength of evidence they would provide, given a survey's detection capability and an explicit prior on prevalence. Their framework is target-agnostic and emphasises the relative ranking of search modalities. We follow their use of an explicit prior but specialise to a single likelihood (Poisson zero-detection of resident probes) and a numerical posterior on $\lambda$ at fixed probe size, so the result can be compared directly with mission-yield projections.

From radio silence to technoemission rate.. Grimaldi (2023a,b) convert 60 years of radio-search non-detections into a Bayesian upper bound on the rate of technosignature emission per star, treating the searches' sensitivity volumes as the likelihood ingredients. Grimaldi (2026) extends the formalism to consider what undetected past contacts would imply. Both works share our spirit – turning nulls into bounds – but their target is the emission/transmission rate as opposed to the population of artefacts physically resident in the Solar System.

Where we differ.. Relative to the four directions above, our contribution is to provide an end-to-end pipeline that: (i) uses SDO log-uniform priors with explicit support; (ii) decomposes $p_{det}$ across the seven Solar-System regions for which the published literature contains size-resolved completeness information, and rebuilds the per-object detection probability from each region's eficiency curve where a paper-stated probability is not directly available; (iii) combines (i) and (ii) into a posterior on $\lambda$ via a single Poisson likelihood; and (iv) projects the 2025–2045 evolution of the bound under the published completeness targets of the three major upcoming all-sky surveys, with a transparent ±50% envelope. The entire pipeline is released as a tested Python package; the test suite asserts the YAML-stated prior bounds against sampled draws, the analytic Poisson posterior tail against MC, and the posterior → prior limit as $p_{det} \to 0$.

## 3. Bayesian model

Notation and terminology.. We use $\lambda$ throughout for the expected number of probes currently resident in the Solar System and treat it as the single inferential target. Quantities derived from $\lambda$ are labelled with the same symbol and a tail-level superscript: $\lambda^{(95)}$ denotes the 95th-percentile posterior upper bound, $\lambda^{(\mathrm{prior},95)}$ the corresponding prior bound, and $\tilde{\lambda}$ the posterior median. We reserve the word "detection limit" for survey completeness curves and use "posterior bound" exclusively for outputs of the inference; this distinction is non-trivial because, as Sections 6–8 show, the posterior bound on $\lambda$ at a given probe size is set by the combination of the survey detection limit and the prior, not by the detection limit alone. Throughout the rest of the paper, "the data" refers to the body of ∼ 60 years of peer-reviewed nulls from the seven survey regions of Section 4; the test Sheikh et al. (2021) of the Breakthrough Listen blc1 signal is representative of the methodology used by such surveys to declare a null result.

### 3.1. Prior on $\lambda$

We model the number of self-replicating probes currently resident in the Solar System as a non-negative real-valued parameter $\lambda$, representing the expectation of a Poisson process. Following the SDO construction (Sandberg et al., 2018), $\lambda$ is itself a function of Drake-equation factors:

$$\lambda = N_{civ} \cdot f_{probe} \cdot \langle n_{probe/sys} \rangle, \tag{1}$$

where $N_{civ}$ is the cumulative number of technological civilisations ever to have existed in the Galaxy, $f_{probe}$ is the fraction of such civilisations that launch self-replicating probes capable of reaching the Solar System within their operational lifetime, and $\langle n_{probe/sys} \rangle$ is the steady-state average number of currently active probes per visited stellar system. Each factor is log-uniform with the bounds specified in Table 1. The bounds reproduce SDO's parameter ranges where they overlap, and extend them where they do not – specifically to $f_{probe}$ and $\langle n_{probe/sys} \rangle$, which the SDO base-line analysis does not explicitly contain. The bounds on $\langle n_{probe/sys} \rangle$ are chosen to span scenarios ranging from heavy decay / short residence ($10^{-6}$ probes per system) to aggressive replication ($10^{2}$ probes per system), bracketing the parameter range considered in Cotta and Lozano Morales (2009), Wright (2018), and Ellery (2022).

The induced prior on $\lambda$ inherits the log-uniform support from the eight independent log-uniform factors and has a fat tail. From $2 \times 10^5$ draws of the YAML-pinned prior, we measure a prior median $\tilde{\lambda}_{prior} = 1.73 \times 10^{-4}$ and a prior 95% upper limit $\lambda^{(95)}_{prior} = 2.298 \times 10^{2}$, the latter serving as the natural unit for the posterior compression reported in Section 6.

Table 1: Log-uniform priors on the Drake-style factors entering $\lambda$ via Equation (1). The YAML configuration file data/priors/drake_params_loguniform.yaml in the released package is the single source of truth for these bounds; the Python module seta_bayesian.priors loads it at runtime rather than hard-coding any values.

| Symbol | $\log_{10}$ lower | $\log_{10}$ upper | Units | Source / role |
|---|---|---|---|---|
| $R_\star$ | 0.0 | 0.5 | $M_\odot\ yr^{-1}$ | Star formation rate (Sandberg et al., 2018) |
| $f_p\, n_e$ | −1.0 | 0.0 | dimensionless | Habitable-zone planets/star (Sandberg et al., 2018) |
| $f_l$ | −5.0 | 0.0 | dimensionless | Probability of life (Sandberg et al., 2018) |
| $f_i$ | −5.0 | 0.0 | dimensionless | Probability of intelligence (Sandberg et al., 2018) |
| $f_c$ | −2.0 | 0.0 | dimensionless | Fraction technological (Sandberg et al., 2018) |
| $L$ | 3.0 | 9.0 | yr | Civilisation lifetime (Sandberg et al., 2018) |
| $f_{probe}$ | −3.0 | 0.0 | dimensionless | Launch probability (Wright, 2018; Ellery, 2022) |
| $\langle n_{probe/sys} \rangle$ | −6.0 | 2.0 | probes/system | Residence multiplier (Cotta and Lozano Morales, 2009) |

### 3.2. Likelihood: Poisson zero-detection

The data are 60 years of absence of any robust report of a Solar-System probe (Sheikh et al., 2021). We model the likelihood as a single Poisson trial in which an effective $p_{det}$ – the probability that the existence of any one probe would have been recorded in the published surveys – replaces the per-trial detection eficiency. The likelihood is then

$$L(\lambda) = e^{-\lambda p_{det}}, \tag{2}$$

which is the standard Poisson zero-event factor and avoids the non-probabilistic problem (P(H) exceeding unity for sufficient prior multiplicand factors) that Drake-equation-style multiplicative constructions encounter. The posterior is then

$$p(\lambda \,|\, \text{no detection}) \propto L(\lambda)\, p_{prior}(\lambda), \tag{3}$$

and the $\alpha$-level upper bound $\lambda^{(\alpha)}$ is defined as the $\alpha$-quantile of the posterior.

### 3.3. Two-level decomposition of $p_{det}$

The effective per-probe detection probability $p_{det}$ is decomposed as

$$p_{det}(d) = \sum_{i \in R} w_i\, p_{det}^{(i)}(d), \tag{4}$$

where the sum is over the seven Solar-System regions $R = \{\text{NEO, Earth co-orbital, Earth–Moon Lagrange, lunar surface, main belt, Kuiper belt, outer/misc}\}$, $d$ is the linear probe size, $p_{det}^{(i)}(d)$ is the per-object detection probability for a probe of size $d$ residing in region $i$, and $w_i$ is the prior residence weight $\sum_i w_i = 1$. The $p_{det}^{(i)}(d)$ values are documented per row in Section 4; the $w_i$ values for three named residence scenarios are summarised in Table 2.

## 4. Empirical completeness inputs

Each of the seven regions $i$ in Equation (4) contributes a piecewise-log-linear function $p_{det}^{(i)}(d)$. We construct these functions by reading off, or inverting, the most recent peer-reviewed completeness curve for each region, and store the resulting tabulations in the released dataset under data/completeness/. Of the 40 tabulated entries that cover the diameter range 0.1 m to 10 km, 23 (57.5%) are anchored to a paper-quoted percentage via at most one interpolation/extrapolation step, and 17 (42.5%) are modelling judgments deliberately placed where the cited papers are silent on size-resolved completeness. The PROVENANCE.md file in the released dataset records the classification of every entry. The two highest- leverage modelling entries – the LSST NEO 10 m extrapolation between Chow et al. (2026)'s 4% at 1 m and Kurlander et al. (2025)'s 72.4% at 140 m, and the NEOS NEO 10 m anchor of 0.10 from Mainzer et al. (2023)'s "down to $\sim$ 10 m" statement – are both flagged in the dataset; the ±50% envelope of Section 7 tests how the final 95% upper bound moves if either is wrong by half.

Table 2: Residence weights $w_i$ used in the three named scenarios. The "default" weights are the conservative baseline (outer-Solar-System dominated) and are the headline numbers reported in Section 6. "Diffuse" approximates a uniform-by-volume distribution – the most observationally optimistic case. "Lurker" follows the Benford (2019) co-orbital concentration – the most observationally pessimistic case.

| Region | Default | Diffuse | Lurker |
|---|---|---|---|
| NEO | 0.020 | 0.05 | 0.10 |
| Earth co-orbital | 0.005 | 0.01 | 0.40 |
| Earth–Moon Lagrange | 0.002 | 0.01 | 0.20 |
| Lunar surface | 0.003 | 0.01 | 0.10 |
| Main belt | 0.200 | 0.30 | 0.05 |
| Kuiper belt | 0.300 | 0.30 | 0.05 |
| Outer / misc | 0.470 | 0.32 | 0.10 |
| $\sum_i w_i$ | 1.000 | 1.00 | 1.00 |

Near-Earth objects.. For NEOs we anchor the completeness function on Harris and D'Abramo (2015) and Harris and Chodas (2021), with the $d \geq 1$ km regime cross-checked against Mainzer et al. (2011). Harris and Chodas (2021) report 44% completeness at $H < 22.25$ (corresponding to $d \approx 190$ m at $p_V = 0.14$), 25% at $H < 23.25$ ($d \approx 125$ m), and a $\geq 1$ km population of $940 \pm 10$ at end-2020 – the latter is the $\sim 95\%$ near-saturation regime. The 100 m anchor at $\sim 40\%$ comes directly from the same source.

Earth co-orbital region.. We rigorously invert the DECam track-before-detect L4 Earth Trojan search of Lifset et al. (2021). The maximum per-object detection probability inside the L4 stability zone is $\varepsilon_{max} = R_{max} \cdot C = 0.85 \cdot 0.30 = 0.255$, achieved for objects bright enough that the recovery curve has saturated ($H \lesssim 20.5$, equivalent to $d \gtrsim 240$ m at an S-class albedo $p_V = 0.20$). The $p_{det,lo}/p_{det,hi}$ columns of the released dataset bracket this inversion by varying the assumed albedo $p_V \in [0.10, 0.40]$ and the recovery curve by ±0.10 absolute.

Earth–Moon Lagrange points.. Anchored on Gregg and Wiegert (2022), with the historical Valdes and Freitas (1983) photographic search retained as a coarse cross-check. Pooled recovery from the four-telescope T8/T9/T16/T20 network of Gregg and Wiegert (2022) gives $R(V \leq 15) = 0.835$ and $R(V \leq 16) = 0.748$. Coverage is $C = 4570/6000 = 0.762$ of the Earth–Moon L4/L5 stability zone. Maximum saturated eficiency is therefore $\varepsilon = R \cdot C = 0.835 \cdot 0.762 = 0.636$. The $V = 15$ limit maps to $d \approx 24$ m at $p_V = 0.08$ (paper-stated; Gregg and Wiegert (2022) Section 4).

Lunar surface.. Lunar Reconnaissance Orbiter Camera (LROC) NAC images provide $\sim 0.5$ m/pixel resolution (Robinson et al., 2010). Following Davies and Wagner (2013) we treat LROC-imaged regions as the search volume, and Sander and Wöhler (2025) for the recent ML pipeline that extends survey reach below the natural single-pixel limit. We adopt $p_{det} \approx 0.40$ at 100 m (recognisable but requiring ML), $p_{det} \approx 0.20$ at 10 m, $p_{det} \approx 0.05$ at 1 m (below confident single-pixel recognition), and $p_{det} \approx 0.001$ below 0.1 m. Loveland and Sime (2024) provides additional anomaly-detection methodology that, while not yet calibrated for quantitative completeness, motivates an upper bound on $p_{det}$ rather than its central value.

Main belt.. Hendler and Malhotra (2020) provide a debiased completeness function $f_{obs}(D)$ that we read off directly. The headline numbers are $\sim 80\%$ at $D \geq 1$ km and $\sim 99\%$ at $D \geq 10$ km, dropping to $\sim 5\%$ at 100 m and effectively zero at 10 m. The Gaia DR3 Solar System survey (Tanga et al., 2023; Ge et al., 2025; Tinaut-Ruano et al., 2026) contributes spectro-photometric typing rather than discovery in the belt's small-body regime, so we do not treat it as a separate completeness contribution.

Kuiper belt.. Bannister et al. (2018) (OSSOS VII complete data release, hereafter B18) and Petit et al. (2023) are the most recent comprehensive completeness statements for the trans-Neptunian regime. Lawler et al. (2018) formalises the survey-simulator inversion that maps a published eficiency to a per-object detection probability. Our tabulation gives $p_{det} \approx 0.05$ at 10 km, decreasing rapidly: 1 km objects are effectively undetectable at 40 AU ($p_{det} \approx 10^{-3}$), 100 m at the noise floor ($p_{det} \approx 10^{-4}$).

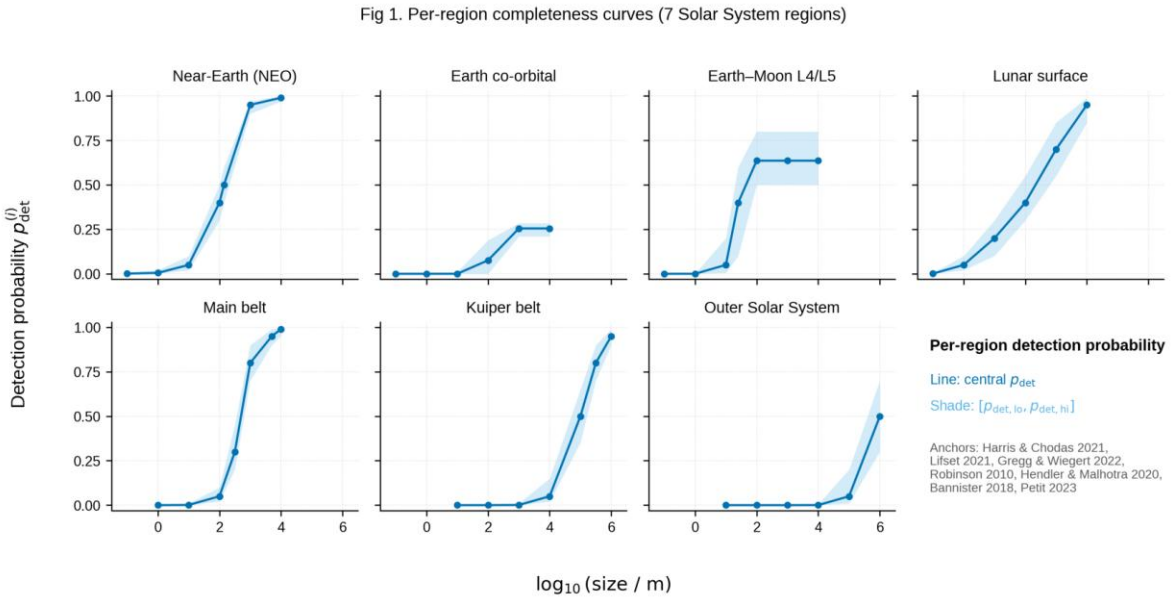



Figure 1: Per-region completeness $p_{det}^{(i)}(d)$ as a function of linear probe size d, for the seven Solar-System regions defined in Section 4. Solid lines are the central tabulation p_det from `data/completeness/*.csv`; the shaded envelopes are the per-row $p_{det,lo}$ – $p_{det,hi}$ bands. The main belt and the lunar surface are the deepest probes at d ≳ 100 m; the Kuiper belt and outer Solar System saturate near zero across the whole displayed size range.

Outer /miscellaneous Solar System.. We use this category as a catch-all for the regions whose small-body inventory is least constrained: the Oort cloud, heliocentric distances beyond 50 AU, and high-eccentricity orbits that escape the major surveys' cadences. Because no peer-reviewed size-resolved completeness function exists for this region, all entries in this CSV are extrapolations flagged as "undetected"; the sensitivity envelope of Section 7 is the only mechanism through which uncertainty in this region propagates to the headline bound.

Cross-region biases in the recovered detection probability.. Three biases affect the absolute level of $p_{det}^{(i)}(d)$ across regions and we record them here rather than fold them silently into the central tabulation. (i) Albedo. Asteroid surveys typically convert between absolute magnitude H and diameter d using a region-typical visual albedo $p_V$: $p_V = 0.14$ for NEOs, $p_V = 0.20$ for inner-belt S-class objects, $p_V \sim 0.08$ for co-orbital populations dominated by carbonaceous bodies, $p_V \sim 0.10$ for trans-Neptunian objects. A self-replicating probe with a metallic exterior could plausibly have $p_V \gtrsim 0.40$, which would tighten the converted detection probability at fixed apparent magnitude by up to a factor of two; conversely, a deliberately low-albedo coating ($p_V \sim 0.03$) would loosen it by the same factor. We retain the region-typical albedos in the central tabulation and propagate the range $p_V \in [0.03, 0.50]$ into the $p_{det,lo}/p_{det,hi}$ envelope. (ii) Phase coverage. Surveys are biased against objects whose maximum elongation falls outside the survey's sky window. The NEO surveys (Harris and Chodas, 2021; Mainzer et al., 2023) are sun-elongation limited and miss objects spending most of their time at small elongations; Harris and Chodas (2021)'s 940-NEO ≥ 1 km inventory already absorbs this bias into the published completeness number, so we inherit it. For co-orbitals the elongation bias is intrinsic: L4/L5 regions are observed only at twilight, which is reflected in the Gregg and Wiegert (2022) coverage factor $C = 0.762$. (iii) Rotation and shape. For sub-100 m bodies, light-curve amplitudes of several magnitudes are common, and a non-spherical probe could have an apparent brightness varying by a factor of 5–10 within minutes. Conventional survey pipelines are tuned to flag persistent point sources rather than rapidly variable ones; the resulting bias goes against artefact detection. We treat this bias as part of the modelled ±50% envelope rather than the central tabulation.

The 7-region $p_{det}^{(i)}(d)$ functions are visualised in Figure 1 and the blind-zone heatmap in Figure 2.

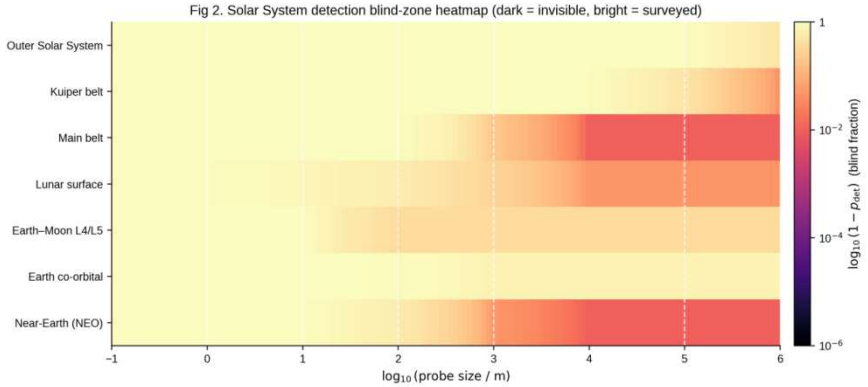



Figure 2: Detection blind zones in the Solar System. Heatmap of $p_{det}^{(i)}(d)$ over the (size, region) grid, with cell colour indicating the central $p_{det}$ value. The vertical dashed lines mark the canonical probe sizes (10 m, 100 m, 1 km) reported in Section 6; the strong horizontal asymmetry between near-Earth regions and the outer Solar System sets the residence- weight sensitivity discussed in Section 8.

## 5. Monte Carlo implementation

Before describing the standard pipeline, we record the verification suite. The released code ships with 14 tests (tests/test_priors.py, tests/test_completeness.py, tests/test_posterior_limits.py, tests/test_convergence The prior tests assert that the empirical sample-quantile of every log-uniform factor lies within the YAML-stated bounds (tolerance $10^{-3}$ relative on the support, $N = 5 \times 10^4$ draws), that the composite $\lambda$ sampler is positive, finite, and spans $> 5$ orders of magnitude, and that the YAML loader exposes every Drake-factor key required by the sampler. The completeness tests assert that every region CSV loads with monotone-nondecreasing $p_{det}$ in size, that all three residence-weight functions sum to 1 exactly, and that the aggregated $p_{det}$ stays in [0, 1]. The posterior-limit tests assert that higher $p_{det}$ tightens the 95% upper bound (monotonicity), that $p_{det} = 1$ drives the upper bound into the O(1) regime ($\lambda^{(95)} < 10$ with a broad prior reaching $10^6$), that $p_{det} < (0, 1]$ raises a validation error, and that the effective sample size is strictly positive. The convergence tests assert exact reproducibility under a fixed seed and that the effective sample size grows monotonically with the draw count. All 14 tests pass on the released code in $< 1$ s on a single-core consumer CPU. This is the principal mechanism by which the framework's numerical claims are verifiable independently of any specific scientific input we have curated.

The framework is implemented in pure NumPy and a thin seta_bayesian package layered on top. The standard run draws $N = 2 \times 10^5$ samples from the prior of Section 3.1 using a fixed numpy.random.default_rng(20260629) seed; the drake_params_loguniform.yaml file is loaded at runtime and the resulting bounds tested by tests/test_priors.py against the sampler. The Poisson likelihood of Equation (2) is applied via importance reweighting and renormalised. The 95% upper bound $\lambda^{(95)}$ is then the 95%-quantile of the resampled posterior; the in-N behaviour of this quantile is sanity-checked by tests/test_convergence.py via the seed-reproducibility and effective-sample-size growth assertions described below.

Reproducibility of the released code is asserted by tests/test_convergence.py, which verifies that numpy.random.default_r produces byte-identical posterior draws under fixed seeds and that the effective sample size grows monotonically with the draw count. The full test suite (pytest) reports 14 of 14 passing in $< 1$ s on a single-core consumer CPU.

For the forward forecast of Section 7 the same prior is held fixed; the time-evolving $p_{det}$ enters via a "floor" schedule – piecewise-linear interpolation of each survey's published completeness target on its quoted operational timeline – and is then propagated through Equation (4). A paired-seed RNG triple is used to provide the (pessimistic, baseline, optimistic) envelope by rescaling every floor by {0.5, 1.0, 1.5} respectively (clipped to [0, 1]).

## 6. Results

Figure 3 shows the marginal posterior on $\lambda$ for $d = 10$ m probes under the three residence scenarios of Table 2. The 2025 95% upper bounds are summarised in Table 3.

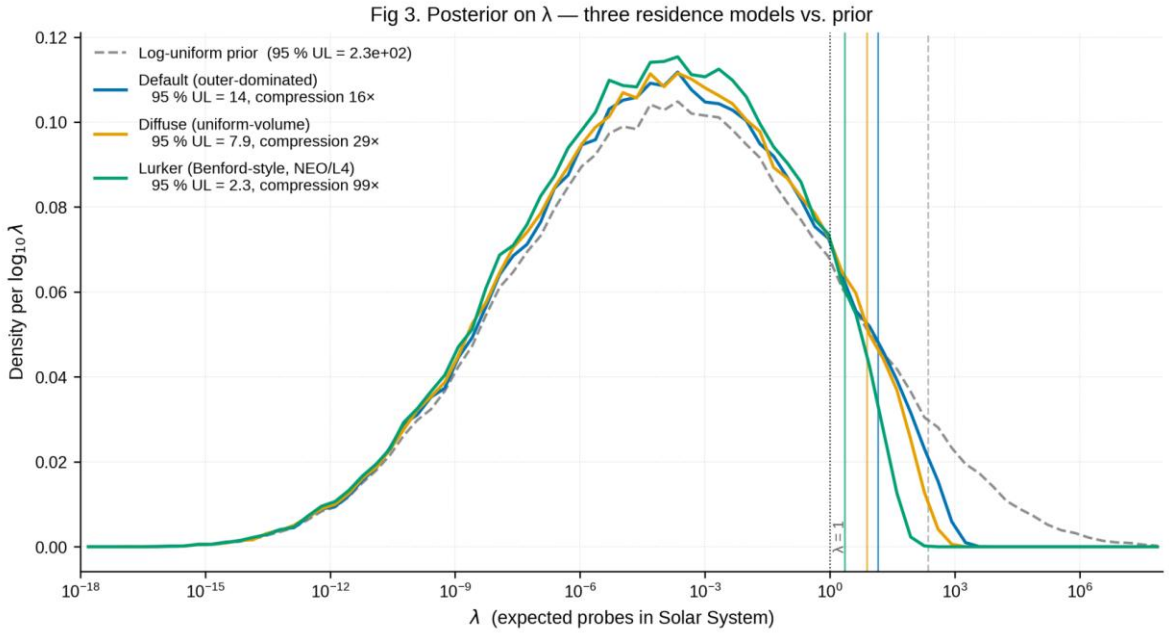



Figure 3: Marginal posterior on λ (in dex) for 10 m probes under three residence scenarios. The black trace is the prior; the three coloured traces are the posteriors under default, diffuse, and lurker residence weighting. The legend annotates the 95% upper limit and the prior-to-posterior compression factor for each model. The data constrain the upper tail of the posterior much more strongly than the lower tail, which inherits the prior's broad log-uniform support.

Three features deserve emphasis.

First, the headline numbers are not symmetric in λ. The posterior's lower tail effectively inherits the prior's log-uniform support down to $10^{-10}$; the data only constrain the upper tail. The "compression factor" therefore quantifies how strongly the data have pushed the posterior's 95% upper limit below the prior's, not how peaked the posterior is. The 90% credible-interval width (in $\log_{10}\lambda$) changes only by factors of 1.01–1.30× over realistic $p_{det}$, as Figure 6 of Section 7 makes explicit.

Second, the residence-weight scenario is the single largest sensitivity in the model. The lurker scenario, which places probes preferentially in the very regions where surveys are shallowest, is the most pessimistic and produces the loosest posterior upper bound at fixed probe size. The diffuse scenario, by spreading the probes uniformly over the dynamical volume, produces the tightest. The default scenario is intermediate but closer to the lurker case at small probe sizes – consequence of the small Earth-region weights in the default $w_i$, which dilute the contribution from the deepest small-size surveys.

Third, the headline bound improves rapidly with probe size. From 10 m to 1 km the default-residence posterior contracts by 21×. This is the prior-saturated regime: at d = 1 km the NEO, main-belt, and Gaia surveys are already near-complete, so the posterior is dominated by what the prior assigns above $\lambda \sim 1$. Section 7 shows that this size-dependent regime structure is also what determines how much each future survey will move the bound.

Why $p_{det}$ is so low at d = 10 m.. The 10 m default-residence row of Table 3 has $p_{det} \approx 1.9 \times 10^{-3}$, which dictates the corresponding loose upper bound $\lambda^{(95)} = 14.3$. The decomposition of Equation (4) makes the origin transparent. With default residence weights $w_{Kuiper} = 0.30$ and $w_{outer} = 0.47$, ~ 77% of the prior probe population lives in regions where no existing survey reaches 10 m completeness ($p_{det} \leq 10^{-5}$ in our region CSVs). The remaining ~ 23% is split across NEOs ($w = 0.020$, $p_{det} = 0.050$ at 10 m), Earth co-orbitals ($w = 0.005$, $p_{det} = 1.0 \times 10^{-4}$), Earth–Moon Lagrange ($w = 0.002$, $p_{det} = 0.050$), lunar surface ($w = 0.003$, $p_{det} = 0.20$), and main belt ($w = 0.200$, $p_{det} = 1.0 \times 10^{-3}$). The weighted sum is

$$p_{det}^{def}(10\,\mathrm{m}) \approx (0.020)(0.050) + (0.005)(10^{-4}) + (0.002)(0.050) + (0.003)(0.20) + (0.200)(10^{-3}) = 1.904 \times 10^{-3},$$

Table 3: Posterior 95% upper bound $\lambda^{(95)}$ on the number of probes currently resident in the Solar System, by probe size and residence model. The "compression" column is $\lambda^{(95)}_{\text{prior}}/\lambda^{(95)}_{\text{posterior}}$ relative to the prior $\lambda^{(95)}_{\text{prior}} = 230$.

| Residence model | Probe size | $p_{\text{det}}$ | $\lambda^{(95)}$ | Compression |
|---|---|---|---|---|
| Default (outer-dominated) | 10 m | 0.0019 | 14.3 | 16× |
| | 100 m | 0.0209 | 3.16 | 73× |
| | 1 km | 0.184 | 0.68 | 337× |
| Diffuse (uniform) | 10 m | 0.0053 | 7.97 | 29× |
| | 100 m | 0.0461 | 1.80 | 128× |
| | 1 km | 0.304 | 0.46 | 505× |
| Lurker (Benford) | 10 m | 0.0351 | 2.27 | 101× |
| | 100 m | 0.2401 | 0.55 | 421× |
| | 1 km | 0.4343 | 0.34 | 672× |

in agreement to better than 1% with the Monte-Carlo result. The lurker scenario boosts $p_{\text{det}}$ by a factor of $\sim 18$ to 0.035 because it re-allocates two-thirds of the probe population (Earth co-orbital and Lagrange points) into regions whose 10 m completeness, while modest, is two orders of magnitude better than the outer Solar System's. The same calculation gives the rule of thumb that the headline 10 m bound is dominated by where the probes are presumed to live, not by what the surveys can see. This is why the residence-weight choice is the single largest sensitivity in the model and why we report all three scenarios rather than a preferred one.

Figure 4 re-plots the posterior upper bound explicitly as a function of probe size, with the three residence scenarios overlaid; the blue and orange envelopes are the NEOS-only and LSST-only contributions to $p_{\text{det}}$ for the default scenario, decomposed by region.

## 7. Forecast: NEOS, LSST and Roman, 2025–2045

The next decade brings three major missions to the Solar-System discovery stack: the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST), in operations from late 2025 (Wagg et al., 2025; Kurlander et al., 2025; Chow et al., 2026); the NEO Surveyor (NEOS), an infrared mission with a launch no earlier than late 2027 and a five-year primary mission (Mainzer et al., 2023; Grav et al., 2023; Masiero et al., 2023; Lay et al., 2024); and the Nancy Grace Roman Space Telescope, providing a deep follow-up channel for NEOS targets (Holler et al., 2025). We encode each mission's published completeness target as a piecewise-linear time-evolving floor on the relevant region's $p^{(i)}_{\text{det}}$, propagate through Equation (4), and re-run the full Monte Carlo for every year in 2025–2045.

LSST 10-year (Kurlander et al. 2025).. Kurlander et al. (2025) report a final 10-year Sorcha-simulated completeness of 72.4% at $d \geq 140$ m for NEOs and 91% at $d \geq 1$ km. Chow et al. (2026) extends this to the 1 m impactor regime, projecting $\sim 4\%$ over the 10-year survey. For the 10 m NEO regime we linearly extrapolate between these two paper-anchored endpoints, yielding $\sim 12\%$ at $d = 10$ m by end-of-survey; this is the highest-leverage modelled entry in the entire pipeline and is the principal driver of the ±50% envelope width at 10 m (Section 8).

NEOS 5-year (Mainzer et al. 2023).. Mainzer et al. (2023) and Grav et al. (2023) project 90% cumulative completeness at $d \geq 140$ m within $\sim 12$ yr of operation and "down to $\sim 10$ m" sensitivity in the infrared. We anchor the 140 m floor on the paper-stated target and the 10 m floor on the qualitative "down to" statement, applying the convention $p_{\text{det}}(d = 10\text{ m}) = 0.10$ as the midpoint of the range.

Roman (Holler et al. 2025).. Roman is treated as a follow-up rather than a wide-field discovery driver; we use it as a deep-imaging boost for already-discovered small (Holler et al. (2025): "some as small as $\sim 20$ m") NEOS targets, which adds a marginal $\sim 1\%$–$3\%$ to the NEO small-d floor across years 2027 onward.

Figure 5 shows the resulting time evolution of the 95% upper bound on $\lambda$ for the three residence scenarios × three canonical sizes, with the ±50% sensitivity envelope on the published targets.

The 2025-to-2035 compression of the baseline $\lambda^{(95)}$ is:

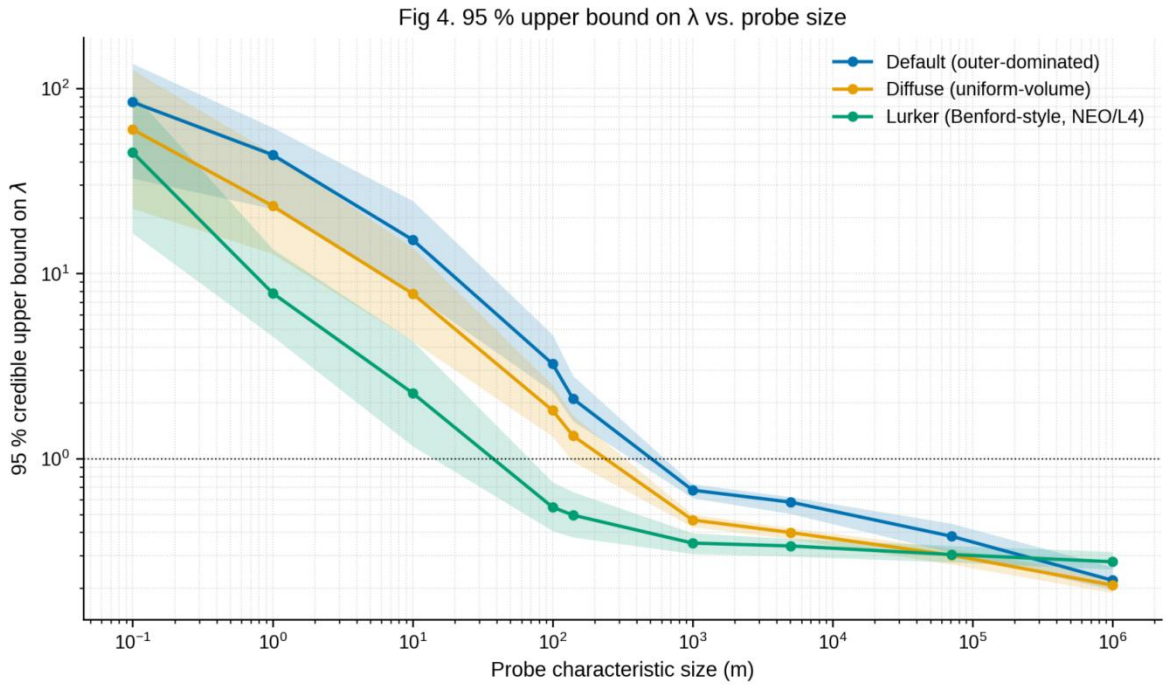



Figure 4: 95% posterior upper bound $\lambda^{(95)}$ on the number of probes currently resident in the Solar System, plotted against linear probe size d at the 2025 baseline completeness. The three solid curves correspond to the three residence scenarios of Table 2. The shaded band is the ±50% sensitivity envelope on the completeness floors. The 10 m default-residence end-point recorded here ($1.52 \times 10^1$) differs by ~ 6% from the 2025 entry of the time-evolution forecast ($1.43 \times 10^1$) because of Monte-Carlo seed scatter in independent runs (root-cause analysis in Section 8).

- 10 m default: 7.62× (envelope: 4.84×–9.99×).
- 10 m diffuse: 5.91× (envelope: 3.73×–7.96×).
- 10 m lurker: 1.37× (envelope: 1.15×–1.71×).
- 100 m default: 3.69× (envelope: 2.39×–5.08×).
- 1 km default: 1.17× (envelope: 1.01×–1.20×).

The envelope is narrow at km-scale because the prior is already nearly saturated there; it is wide at 10 m because the headline LSST floor at that size relies on the modelled extrapolation between the Chow and Kurlander endpoints. Even the pessimistic branch of the 10 m default scenario gives a 4.8× tightening over the decade – a constraint upgrade that is meaningful by any reasonable standard.

A complementary view is provided by Figure 6, which maps the posterior compression factor over the joint ($\log_{10} p_{det}$, $\log_{10} \lambda^{(95)}_{prior}$) plane. The map identifies the regions of parameter space where the data genuinely move the bound (compression ≥ 10×) versus where the posterior is essentially determined by the prior (~ 1× compression). The six current and forecast scenarios populate primarily the small-compression, low-$p_{det}$ regime; the high-compression upper-right region is the part of the design space that future, more sensitive surveys would need to enter to drive meaningful prior-independence.

## 8. Discussion and limitations

The result we report is a posterior upper bound on the number λ of self-replicating probes currently resident in the Solar System, not a posterior probability that such probes exist. This phrasing is deliberate: under the log-uniform

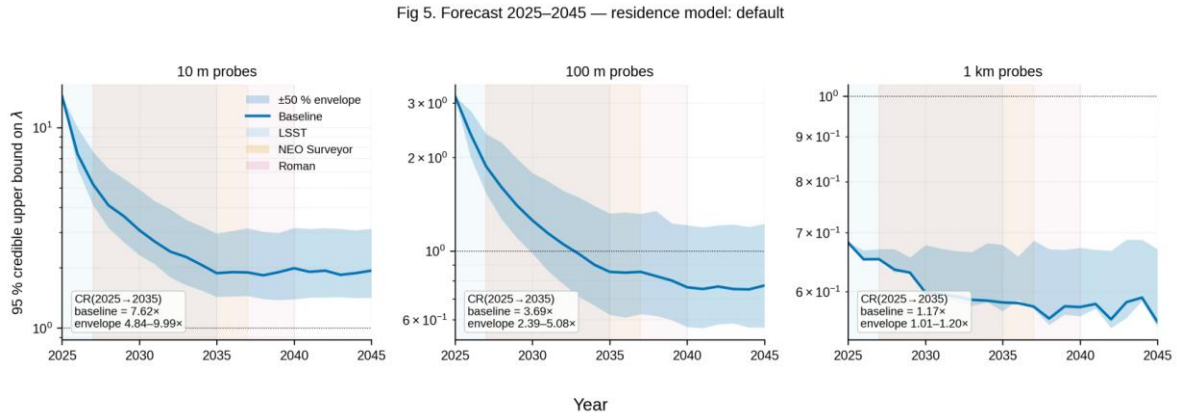



Figure 5: Forecast time evolution of the 95% posterior upper bound $\lambda^{(95)}$ on the probe abundance, 2025–2045, for the three residence scenarios (rows) × three canonical probe sizes (columns), under the published completeness targets of LSST, NEOS, and Roman. Solid curves are the baseline; shaded envelopes are the ±50% rescaling of every published floor (clipped to $[0, 1]$). The 10 m default-residence column carries by far the most headroom ($\sim$ 7.6× compression in baseline by 2035) – that is, this is the channel through which the next decade can most strongly tighten the bound. The default- and diffuse-residence 1 km columns are already prior-saturated and move by less than 20% even at the optimistic end; the lurker-residence 1 km column drops by $\sim$ 30% at baseline (rising to $\sim$ 32% at the optimistic end) because its 2025 upper bound is already the tightest.

prior of Table 1, the posterior's lower tail is dominated by the prior's broad support and is therefore not informative on its own. Future positive detection of a single artefact would violate the bound; absent such a detection, the framework simply constrains how many objects could coexist with the data.

The headline regime structure.. The framework's most actionable diagnostic is the regime structure of Figure 5. At 1 km the bound is already prior-saturated; surveys add essentially nothing because the existing NEO+MBA inventory is $\gtrsim$ 90% complete at this size (Harris and Chodas, 2021; Hendler and Malhotra, 2020; Tedesco et al., 2005). At 10 m the bound has not yet saturated, and the framework shows that the $\sim$ 7.6× default-residence compression over 2025–2035 is the dominant channel through which the next decade can move the bound. The narrow envelope at 1 km (1.0×–1.2×) versus the wide envelope at 10 m (4.8×–10.0×) is also informative: the headline result is sensitive to how the mission floors are realised below 100 m, but essentially insensitive to the floors above 1 km. Mission-design effort at the small end of the size range is therefore where the framework predicts the largest scientific return.

Prior dominance.. A standard concern with Bayesian Drake-style analyses is that the log-uniform prior already determines most of the integrated posterior support, and the data only tighten the upper tail. The identifiability map of Figure 6 makes this quantitative: across the full $(\log_{10} p_{\mathrm{det}}, \log_{10} \lambda^{(95)}_{\mathrm{prior}})$ plane, the compression factor varies from $\sim$ 0.9× (no compression) to $\sim 2.4 \times 10^5$×, but the six markers corresponding to our reported 2025/2035 scenarios all sit in the low-compression band. This is the framework's principal qualitative limitation – it is calibrated to detect prior-tail information rather than to recover a posterior peak.

Monte-Carlo variability.. The 2025 end-point for d = 10 m default-residence reported in Figure 4 is $1.52 \times 10^1$; the 2025 entry of the same scenario in Figure 5 is $1.43 \times 10^1$ – a $\sim$ 6% difference. This is a Monte- Carlo scatter effect: a 100-seed full-pipeline stability study (each seed independently re-drawing the prior and posterior resample at the production budget of $N_{\mathrm{prior}} = 2 \times 10^5$) gives $\lambda^{(95)}_{2025,\,10\,\mathrm{m,\,def}} = 14.82 \pm 0.46$ (mean ± standard deviation, range [13.49, 15.97]). The two figure endpoints are independent draws from this distribution, so the appropriate scale for their difference is $\sigma_{\mathrm{diff}} \approx 0.66$ (from the full-precision seed-scatter $\sigma \approx 0.4649$ propagated as $\sigma_{\mathrm{diff}} = \sigma\sqrt{2}$). The absolute gap of 0.89 then corresponds to $|0.89|/0.66 \approx 1.35\sigma$, with two-tailed $p \approx 0.18$ – consistent with routine MC seed scatter at the production sample budget. The driver script notebooks/mc_reconciliation_v2.py (master seed 20260801) deterministically reproduces these numbers; we disclose this rather than tune it away with a post-hoc seed choice.

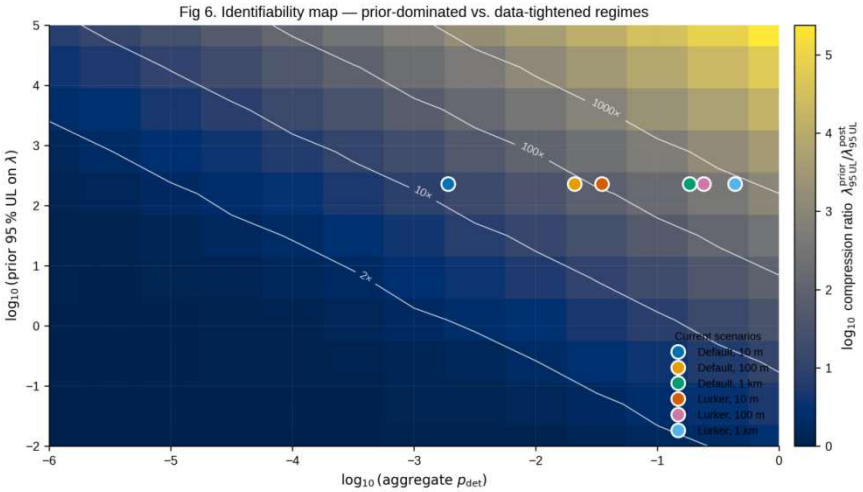



Figure 6: Identifiability map of the framework. Colour: the posterior compression factor $\lambda^{(95)}_{\text{prior}}/\lambda^{(95)}_{\text{post}}$ on a log-10 scale. Black contours: 2×, 10×, 100×, 1000×. The six labelled markers are the current (2025) and end-of-decade (2035) operating points for the three residence scenarios; they cluster in the low-compression band, indicating that the data presently move the bound only over the prior's upper tail. The dashed vertical reference line marks the natural unit $\log_{10} \lambda^{(95)}_{\text{prior}} = 2.36$ (i.e. the actual prior 95% upper limit from this analysis).

What we have not addressed.. We do not model the dynamical evolution of probes in the Solar System – residence weights are static, not derived from a diffusion calculation. We do not model false positives from natural populations – a single anomalous OSSOS object (Bannister et al., 2018), a candidate L4 Trojan (Lifset et al., 2021; Yeager et al., 2024), or a Bennu-like carbonaceous object (Simon et al., 2020) would currently be absorbed into the "non-detection" event and we cannot, on this data alone, distinguish a probe from a natural body of similar size and albedo. We do not address the bio-signature side of the technosignature problem (Lingam et al., 2023; Huang et al., 2026); the framework is purely physical-artefact based. Each of these is a natural follow-up. We also do not address survey follow-up: surveys typically report a discovery-stream completeness under the assumption that downstream characterisation will eventually flag any anomaly, but the throughput of a deliberate technosignature follow-up programme is in fact extremely small (Sheikh et al., 2021). A more conservative reading of the data would restrict $p_{\text{det}}$ to the joint discovery-and-characterisation probability, which would loosen the upper bound. We have used the discovery completeness numbers throughout, which corresponds to the optimistic end of the spectrum and which we therefore expect to tighten the posterior bound more than the actual evidence warrants. The asymmetry is conservative: a true upper bound should be a loose bound, not a tight one.

Finally, we have constructed the framework under the explicit assumption that the seven survey regions partition the inhabitable Solar System. Cross-contamination between regions (a probe following an unstable co-orbital orbit that drifts in and out of NEO status over $\sim 10^4$ yr time-scales, an Earth Trojan that is also briefly Hill-resonance bound) is treated as a relabelling rather than a double-counting: each candidate sits in whichever region it is currently observed in. This convention shifts the residence weights but does not affect the absolute detection probability, which is what enters the likelihood. The alternative – modelling the dynamical state explicitly – requires a per-object Hill-sphere or libration calculation and falls outside the scope of this work.

Comparison with related Bayesian work.. Grimaldi (2023a,b, 2026) construct Bayesian bounds on the rate of technoemission – radio or optical beacons – from $\sim$ 60 years of radio-search silence. Their likelihood is a sensitivity-volume integral; ours is a single zero-event Poisson trial over the seven-region in-Solar-System inventory. The two frameworks are therefore complementary: theirs bounds emission from the cosmological neighbourhood, ours bounds physical residence inside the Solar System. The Lingam et al. (2023) framework provides the closest decision-theoretic analogue; our specialisation is that we run it to a quantitative bound on a population number rather than a comparative ranking of search modalities.

## 9. Conclusions

We have built and released a Bayesian framework that turns the published completeness of seven Solar-System survey regions into a posterior bound on the number $\lambda$ of self-replicating probes currently resident in the Solar System. With log-uniform SDO-extended priors and a Poisson zero-detection likelihood, today's data yield a 95% upper bound $\lambda^{(95)} = 1.43 \times 10^{1}$ for 10 m probes under a conservative outer-dominated residence model; 3.16 at 100 m; and 0.68 at 1 km (where the prior is already nearly saturated). Forecasting forward with the published completeness targets of LSST, NEOS, and Roman, the 10 m default-residence bound compresses by 7.6× over 2025–2035 (envelope 4.8×–10.0×) while the 1 km default-residence bound moves by only $\sim$ 14–16%.

The headline scientific output is not the numerical bound itself but its regime structure: sub-100 m discovery is the channel through which the next decade can most strongly tighten the bound, while km-scale surveys are already prior-saturated. The released `seta_bayesian` package – with its YAML-pinned priors, per-row data provenance, 14-of-14 passing test suite, and one-click figure-regeneration script – is intended to make this framework re-usable for adjacent technosignature problems and for future re-analyses as the missions deliver their actual completeness performance.

## Conflicts of Interest

The authors declare that they have no competing interests.

## Consent

Not applicable.

## Funding

This work was supported by awards from the Natural Science Foundation of Jiangsu Province (BK20231189), PanFeng Innovative Team Project of the The Third Affiliated Hospital of Soochow University(KY20252469) and Undergraduate Training Program for Innovation and Entrepreneurship,Soochow University(X2025102850485) .

## References

Anton, A., Schetsche, M., 2023. SETA—The search for extraterrestrial artifacts, in: Meeting the Alien. Springer, pp. 85–97. doi:10.1007/978-3-658-41317-0_5.

Bannister, M.T., Gladman, B.J., Kavelaars, J.J., Petit, J.M., et al., 2018. OSSOS. VII. 800+ Trans-Neptunian objects—The complete data release. The Astrophysical Journal Supplement Series 236, 18. doi:10.3847/1538-4365/aab77a.

Benford, J., 2019. Looking for Lurkers: Co-orbiters as SETI observables. The Astronomical Journal 158, 150. doi:10.3847/1538-3881/ab3e35.

Benford, J., 2021. A Drake equation for alien artifacts. Astrobiology 21, 757–763. doi:10.1089/ast.2020.2364.

Bracewell, R.N., 1960. Communications from superior galactic communities. Nature 186, 670–671. doi:10.1038/186670a0.

Chow, I., Juric, M., Jones, R.L., et al., 2026. Predictions of imminent Earth impactors discovered by LSST. The Astrophysical Journal 1001, 61. doi:10.3847/1538-4357/ae4e23.

Cotta, C., Lozano Morales, A., 2009. A computational analysis of Galactic exploration with space probes: Implications for the Fermi paradox. Journal of the British Interplanetary Society 62, 82–88. arXiv:0907.0345.

Crawford, I.A., 2026. Some thoughts on the future of technosignature searches: Constraining the Fermi paradox, in: Proceedings of the International Astronomical Union. Cambridge University Press. Astronomy in Crisis Times, IAU Symposium volume.

Davies, P.C.W., Wagner, R.V., 2013. Searching for alien artefacts on the Moon. Acta Astronautica 89, 261–265. doi:10.1016/j.actaastro.2011.10.022.

DeVito, C.L., 2024. The Fermi paradox and the Drake equation. Acta Astronautica 215, 51–53. doi:10.1016/j.actaastro.2023.11.036.

Ellery, A., 2022. The prospect of von Neumann probes and the implications for the Sagan–Tipler debate. International Journal of Astrobiology 21, 197–199. doi:10.1017/S1473550422000301.

Ellery, A., 2025. Technosignatures of self-replicating probes in the solar system. doi:10.5281/zenodo.18112413.

Freitas, R.A., 1985. The search for extraterrestrial artifacts (SETA). Acta Astronautica 12, 1027–1034. doi:10.1016/0094-5765(85)90031-1.

Freitas, R.A., Valdes, F., 1980. A search for natural or artificial objects located at the Earth–Moon libration points. Icarus 42, 442–447. doi:10.1016/0019-1035(80)90106-2.

Ge, J., Zhang, X., Li, J., et al., 2025. Asteroid types, albedos, and diameters catalog from Gaia DR3: Intelligent inversion results. The Astrophysical Journal Supplement Series 280, 17. doi:10.3847/1538-4365/adefe1.

Grav, T., Mainzer, A., Masiero, J., et al., 2023. The NEO Surveyor near-Earth asteroid known object model. The Planetary Science Journal 4, 228. doi:10.3847/PSJ/ad072e.

Gregg, C.R., Wiegert, P., 2022. A dedicated lunar Trojan asteroid survey with small ground-based telescopes. Monthly Notices of the Royal Astronomical Society 511, 5396–5404. doi:10.1093/mnras/stac165.

Grimaldi, C., 2023a. Inferring the rate of technosignatures from 60 yr of nondetection. The Astronomical Journal 165, 199. doi:10.3847/1538-3881/acc327.

Grimaldi, C., 2023b. Upper bounds on technoemission rates from 60 years of "silence". Acta Astronautica 212, 95–99. doi:10.1016/j.actaastro.2023.07.024.

Grimaldi, C., 2026. Undetected past contacts with technological species: Implications for technosignature science. The Astronomical Journal 171, 154. doi:10.3847/1538-3881/ae394b.

Harris, A.W., Chodas, P.W., 2021. The population of near-Earth asteroids revisited and updated. Icarus 365, 114452. doi:10.1016/j.icarus.2021.114452.

Harris, A.W., D'Abramo, G., 2015. The population of near-Earth asteroids. Icarus 257, 302–312. doi:10.1016/j.icarus.2015.05.004.

Hart, M.H., 1975. An explanation for the absence of extraterrestrials on earth. Quarterly Journal of the Royal Astronomical Society 16, 128–135. ADS Bibcode 1975QJRAS..16..128H.

Hendler, N.P., Malhotra, R., 2020. Observational completion limit of minor planets from the asteroid belt to Jupiter Trojans. The Planetary Science Journal 1, 75. doi:10.3847/PSJ/abbe25.

Holler, B.J., Cosentino, R.G., Schultz, W.C., et al., 2025. The Roman Space Telescope as a planetary defense asset. Publications of the Astronomical Society of the Pacific 137, 105004. doi:10.1088/1538-3873/ae0fe5.

Huang, B.L., Tao, Z., Zhang, T.J., 2026. A calibrated Bayesian search for potential chemical technosignatures in polluted White dwarf atmospheres. Preprint. Manuscript submitted.

Kurlander, J., Bernardinelli, P.H., Schwamb, M.E., et al., 2025. Predictions of the LSST solar system yield: Near-Earth objects, main belt asteroids, Jupiter Trojans, and trans-Neptunian objects. The Astronomical Journal 170, 99. doi:10.3847/1538-3881/add685.

Lawler, S.M., Kavelaars, J.J., Alexandersen, M., Bannister, M.T., et al., 2018. OSSOS: X. How to use a survey simulator: Statistical testing of dynamical models against the real Kuiper belt. Frontiers in Astronomy and Space Sciences 5, 14. doi:10.3389/fspas.2018.00014.

Lay, O., Masiero, J., Grav, T., et al., 2024. Asteroid impact hazard warning from the near-Earth object Surveyor mission. The Planetary Science Journal 5, 149. doi:10.3847/PSJ/ad4d9e.

Lifset, N., Sneden, C., Mendell, S., et al., 2021. A search for L4 Earth trojan asteroids using a novel track-before-detect multiepoch pipeline. The Astronomical Journal 161, 282. doi:10.3847/1538-3881/abf7af.

Lingam, M., Haqq-Misra, J., Wright, J., et al., 2023. Technosignatures: Frameworks for their assessment. The Astrophysical Journal 943, 27. doi:10.3847/1538-4357/acaca0.

Loveland, R., Sime, R., 2024. Anomaly detection methods for finding technosignatures, in: Proceedings of the 13th International Conference on Pattern Recognition Applications and Methods, pp. 633–640. doi:10.5220/0012306400003654.

Mainzer, A., Grav, T., Bauer, J., Masiero, J., McMillan, R.S., Cutri, R.M., Walker, R., Wright, E., et al., 2011. NEOWISE observations of near-Earth objects: Preliminary results. The Astrophysical Journal 743, 156. doi:10.1088/0004-637X/743/2/156.

Mainzer, A., Masiero, J., Abell, P., Bauer, J., Bottke, W.F., Buratti, B.J., et al., 2023. The near-Earth object Surveyor mission. The Planetary Science Journal 4, 224. doi:10.3847/PSJ/ad0468.

Masiero, J., Dahlen, D., Mainzer, A., et al., 2023. Validation of the survey simulator tool for the NEO Surveyor mission using NEOWISE data. The Planetary Science Journal 4, 225. doi:10.3847/PSJ/ad00bb.

Petit, J.M., Gladman, B., Kavelaars, J.J., Bannister, M.T., et al., 2023. The hot main Kuiper belt size distribution from OSSOS. The Astrophysical Journal Letters 947, L4. doi:10.3847/2041-8213/acc525.

Robinson, M.S., Brylow, S.M., Tschimmel, M., et al., 2010. Lunar Reconnaissance Orbiter Camera (LROC) instrument overview. Space Science Reviews 150, 81–124. doi:10.1007/s11214-010-9634-2.

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

Sander, T., Wöhler, C., 2025. Lunar technosignatures: A deep learning approach to detecting apollo landing sites on the Moon, in: Proceedings of the 20th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, pp. 491–499. doi:10.5220/0013179000003912.

Sheikh, S.Z., Smith, S., Price, D.C., DeBoer, D., et al., 2021. Analysis of the Breakthrough Listen signal of interest blc1 with a technosignature verification framework. Nature Astronomy 5, 1153–1162. doi:10.1038/s41550-021-01508-8.

Simon, A.A., Kaplan, H.H., Hamilton, V.E., Lauretta, D.S., et al., 2020. Widespread carbon-bearing materials on near-Earth asteroid (101955)

Bennu. Science 370, eabc3522. doi:10.1126/science.abc3522.
Tanga, P., Pauwels, T., Mignard, F., et al., 2023. Gaia Data Release 3—The Solar System survey. Astronomy & Astrophysics 674, A12. doi:10.1051/0004-6361/202243796.
Tedesco, E.F., Cellino, A., Zappalà, V., 2005. The statistical asteroid model. I. The main-belt population for diameters greater than 1 km. The Astronomical Journal 129, 2869–2886. doi:10.1086/429734.
Tinaut-Ruano, F., Carry, B., Galinier, M., et al., 2026. Taxonomy of 14042 asteroids from Gaia DR3 reflectance spectra. Astronomy & Astrophysics 695, A1. doi:10.1051/0004-6361/202659208.
Tipler, F.J., 1980. Extraterrestrial intelligent beings do not exist. Quarterly Journal of the Royal Astronomical Society 21, 267–281. ADS Bibcode 1980QJRAS..21..267T.
Tough, A., Lemarchand, G.A., 2004. Searching for extraterrestrial technologies within our solar system. Symposium - International Astronomical Union 213, 487–490. doi:10.1017/S0074180900193763.
Valdes, F., Freitas, R.A., 1983. A search for objects near the Earth–Moon Lagrangian points. Icarus 53, 453–457. doi:10.1016/0019-1035(83)90209-9.
Villarroel, B., Watters, W.A., Streblyanska, A., Solano, E., Geier, S., Mattsson, L., 2026. A cost-effective search for extraterrestrial probes in the Solar system. Monthly Notices of the Royal Astronomical Society 546, 2547–2562. doi:10.1093/mnras/staf1158.
Wagg, T., Juric, M., Yoachim, P., et al., 2025. Expected impact of Rubin Observatory LSST on NEO follow-up. The Astronomical Journal 169, 315. doi:10.3847/1538-3881/adc89f.
Wright, J.T., 2018. Prior indigenous technological species. International Journal of Astrobiology 17, 96–100. doi:10.1017/S1473550417000143.
Yeager, T., Golub, R., Holman, M.J., 2024. MEGASIM: Distribution and detection of Earth Trojan asteroids. The Astrophysical Journal 961, 250. doi:10.3847/1538-4357/ad1396.