

A Monte Carlo Forecast for the Detection of Planet Nine

Working paper · Version 3.0.0 · 2026-06-05 · Not peer reviewed

Citation and status boundary: this PDF is a stable public working-paper edition. It is not a peer-reviewed journal article. Consult the canonical web edition for interactive material, linked resources, corrections, and the current machine-readable citation record.

Where, when, and by which instrument an undetected trans-Neptunian super-Earth would most plausibly be found

Architect & Curator: Mayone Maha Rajan AI Synthesis Instrument: Google Antigravity (agentic model)

Revision 3 June 2026

This forecast was produced through human-directed AI synthesis. The human architect curated the inquiry, selected the source models, and is responsible for all claims; the AI instrument assisted with drafting, simulation, and literature synthesis under that direction. All quantitative results are reproducible from the accompanying simulation.py (seed 42); parameter values quoted from secondary sources should be confirmed against the primary papers before being relied upon.

Abstract

We forecast the detectability of the hypothesized Planet Nine by propagating a weighted three-model parameter ensemble drawn from Batygin et al. (2019), Brown & Batygin (2021), and Siraj et al. (2025) through a Monte Carlo orbit simulation ($N = 100,000$). For each sample we solve Kepler's equation, project to geocentric equatorial coordinates with astropy, estimate an r-band apparent magnitude, apply existing survey null-detection masks (ZTF, Pan-STARRS, DES), and model future detection by the Vera C. Rubin Observatory/LSST, Subaru/HSC, and infrared surveys through 2036. The current heliocentric distance dominates all detection uncertainty. Under a corrected treatment of LSST's northern footprint (Revision 3), the combined cumulative detection probability rises to approximately 0.72 by 2032 under plausible northern-cadence coverage up from the 0.61 reported in Revision 2, which modelled the northern boundary as a hard wall. The exact value is contingent on an undecided Rubin survey-scheduling choice; the honest figure is a range, $\sim 0.61\text{--}0.72$, whose width is the value of that decision. We treat the planet's existence as an unproven hypothesis and report how a sustained null result revises a stated prior. This report does not claim Planet Nine exists; it quantifies what its detection would look like if it does.

Revision note (what changed since Revision 2)

Revision 3 corrects a load-bearing error in the detection-timing model that Revision 2 flagged in principle but did not act on. Revision 2's Limitations section noted that per-facility detection years were drawn from plausible binned schedules, not a cadence-level survey simulator. On examination, one of those heuristics the treatment of LSST's northern footprint was not merely coarse but materially wrong, and correcting it changes both a headline number and a central limitation claim.

1. The LSST northern boundary: a hard wall replaced with a soft, survey-strategy parameter. Revision 2 modelled LSST coverage as a hard cutoff at Dec +15: every sample north of +15 received an infinite detection year (never detected). This is incorrect on two counts. First, the LSST Wide-Fast-Deep baseline_v2.0 footprint spans Dec 72 to +12, not 65 to +15 (arXiv:2303.02355). Second, and more importantly, the northern boundary is not a hardware limit but an editable survey-strategy choice: the telescope at Cerro Pachon can physically observe to Dec < +33.5 at airmass 2.2 (lsst.org),

the Northern Ecliptic Spur mini-survey already extends coverage to ecliptic latitude +10 (arXiv:1812.01149), and a proposed northern extension could fill nearly to Dec +30 (RubinEuclid report). Revision 3 replaces the hard wall with a graded northern-coverage model: full cadence to +12, reduced (NES-extension) cadence with a coverage penalty from +12 to +30, low-priority coverage to +33.5, and no coverage beyond. The northern coverage probabilities in this graded model (0.5 between +12 and +30, 0.15 above) are reasoned estimates of plausible NES-extension coverage, not sourced values; they depend on Rubin scheduling decisions that are not yet finalized, and are tagged accordingly.

2. LSST limiting depth aligned to documented figures. Revision 2 used a flat single-tier limit of $m_r < 25.0$. Revision 3 uses the documented Rubin depths single-visit $r \approx 24.0$, ten-year co-add $r \approx 26.9$ with a graded detection schedule, so bright objects are held to a slightly stricter early-survey limit and the faint tail becomes reachable only via deep late-survey co-addition. This produces a more physically realistic curve shape (lower in the first two seasons, higher in the back half of the survey).

3. Albedo prior aligned to sources. Revision 2 sampled geometric albedo from $\mathcal{U}(0.3, 0.5)$. The companion repository and the cited source (Fortney et al. 2016) span 0.20.75. Revision 3 uses $\mathcal{U}(0.2, 0.75)$. This has negligible effect on the headline detection probability (the distance term dominates, as Revision 2's sensitivity analysis established) but removes an inconsistency between the code and the stated sources.

Effect on the results

The corrections move the forecast's two headline outputs, both in the direction of greater instrument capability. Cumulative detection probability by 2036 rises from 61.3% to 71.9% (stable across random seeds to within Monte Carlo noise of ~ 0.3 points). The increment appears in 20302032, in the northern aphelion region the hard wall had previously zeroed out. The null-detection posterior strengthens accordingly: because the instruments are more capable than Revision 2 assumed, continued non-detection is stronger evidence against existence, pulling the stated prior of 70% down to $\sim 39.6\%$ (Revision 2: $\sim 47.4\%$). The qualitative conclusion is unchanged a null result erodes but does not refute the hypothesis but the erosion is roughly 8 percentage points deeper.

Correction to the stated cause of the residual

Revision 2's Limitations section described the residual non-detection probability as structural a consequence of the faint, distant tail requiring a 10 m-class facility to close. Revision 3 revises this: a substantial part of the residual was not faintness but the hard northern coverage cutoff a survey-scheduling artifact, not a fundamental sensitivity limit. The genuinely irreducible residual (faint, deep, off-ecliptic, galactic-plane-confused) is correspondingly smaller, and the primary recommendation shifts from 'a new large telescope is needed' to 'most of the apparent gap is closable by Rubin northern-cadence allocation; new hardware is warranted only for a narrow, atmosphere- and geometry-specific residual corner.' This correction is developed in the companion paper, Planet Nine Detectability Gap Analysis: 20262035, which is the source of the present revision.

1. The hypothesis and the inference logic

Planet Nine is a proposed distant super-Earth invoked to explain the anomalous orbital clustering of extreme trans-Neptunian objects (eTNOs) bodies whose elongated orbits share a similar orientation in space to a degree unlikely to arise by chance. The planet has never been directly imaged; it is inferred from its gravitational effect on visible bodies, the same inferential route by which Neptune was predicted from perturbations in the orbit of Uranus before it was ever seen.

The observed clustering of distant eTNOs has been assessed as significant at the 99.6% confidence level, corresponding to a chance-occurrence probability below $\sim 0.007\%$ (Brown & Batygin 2021). Other analyses using

different object samples and bias treatments have reported weaker or stronger figures; the existence of the planet remains debated rather than established. This forecast is conditioned on the hypothesis being true and separately tracks the probability that it is not (Section 7).

2. Parameter priors and provenance

We capture two distinct kinds of uncertainty:

- Aleatoric the planet's current position along its orbit is unknown, so the mean anomaly is sampled uniformly over $[0, 2\pi]$.
- Epistemic different research groups infer materially different orbits. Rather than averaging them into a false consensus, we sample from a mixture of three source models, weighted toward the more recent analyses.

Each Monte Carlo sample is assigned to one model, then its orbital elements are drawn from that model's distributions.

| Parameter | Batygin et al. (2019) [a] | Brown & Batygin (2021) [b] | Siraj et al. (2025) [c] | | :--- | :--- | :--- | :--- | | Model weight | 0.20 | 0.40 | 0.40 | | Mass M_p ($M_{\oplus}$) | $\mathcal{U}(5.0, 10.0)$ | $\text{SplitNormal}(6.2, 1.3, 2.2)$ | $\mathcal{N}(4.4, 1.1)$ | | Semi-major axis a (AU) | $\mathcal{U}(400, 800)$ | $\text{SplitNormal}(380, 80, 140)$ | $\mathcal{N}(290, 30)$ | | Eccentricity e | $\mathcal{U}(0.2, 0.5)$ | derived from q : $e = 1 - q/a$ | $\mathcal{N}(0.29, 0.13)$, clipped to $[0, 0.9]$ | | Perihelion q (AU) | derived: $q = a(1 - e)$ | $\text{SplitNormal}(300, 60, 85)$ | derived | | Inclination i (deg) | $\mathcal{U}(15, 25)$ | $\mathcal{N}(16, 5)$ | $\mathcal{N}(6.8, 5.0)$ [c] | | Long. of asc. node Ω (deg) | $\mathcal{N}(105, 15)$ | $\mathcal{N}(105, 10)$ | $\mathcal{N}(105, 20)$ | | Arg. of perihelion ω (deg) | $\mathcal{N}(145, 15)$ | from $\varpi \sim \mathcal{N}(250, 14)$, $\omega = \varpi - \Omega$ | from $\varpi \sim \mathcal{N}(230, 30)$, $\omega = \varpi - \Omega$ | | Mean anomaly M | $\mathcal{U}(0, 2\pi)$ | $\mathcal{U}(0, 2\pi)$ | $\mathcal{U}(0, 2\pi)$ |

Sources. [a] Batygin et al. (2019), Physics Reports "The Planet Nine Hypothesis": $510 M_{\oplus}$, $a \approx 400800$ AU, $i \approx 1525^\circ$. [b] Brown & Batygin (2021), AJ "The Orbit of Planet Nine" (arXiv:2108.09868): $a = 380^{+140}_{-80}$ AU, $m = 6.2^{+2.2}_{-1.3} M_{\oplus}$, perihelion $q \approx 300$ AU, $i \approx 16^\circ$, clustering significant at 99.6%. [c] Siraj et al. (2025): $a = 290 \pm 30$ AU, $e = 0.29 \pm 0.13$, $i = 6.8 \pm 5.0^\circ$, $m = 4.4 \pm 1.1 M_{\oplus}$, predicted apparent magnitude 21. This is the most recent and lowest-distance, lowest-inclination estimate, and the disagreement with [a]/[b] is genuine.

> Clustering significance (a property of the observations, not of any single model): 99.6% confidence, chance probability $< 0.007\%$ (Brown & Batygin 2021). Other studies report different values depending on sample and bias treatment.

A modelling assumption, not a literature value, is the mass-radius relation $R = R_{\oplus} \sqrt{M_p/M_{\oplus}}$. This is a first-order scaling only; it does not closely track empirical mini-Neptune mass-radius data and is used purely to convert mass into a reflecting cross-section. Radius enters magnitude weakly (Section 5), so this approximation has limited impact on the forecast.

3. Orbital mechanics

For each sample, Kepler's equation relates the mean anomaly M to the eccentric anomaly E :

$$M = E - e \sin E$$

solved by ten Newton-Raphson iterations from the starting guess $E_0 = M + e \sin M$:

$$E_{n+1} = E_n - \frac{E_n - e \sin E_n - M}{1 - e \cos E_n}$$

Position in the orbital plane:

$$x_{\text{orb}} = a(\cos E - e), \quad y_{\text{orb}} = a \sqrt{1 - e^2} \sin E, \quad z_{\text{orb}} = 0$$

Rotation into the heliocentric ecliptic frame via $\mathbf{R}_z(-\Omega)\mathbf{R}_x(-i)\mathbf{R}_z(-\omega)$:

$$\begin{aligned} x_{\text{ecl}} &= x_{\text{orb}}(\cos\omega\cos\Omega - \sin\omega\sin\Omega\cos i) - \\ y_{\text{orb}}(\sin\omega\cos\Omega + \cos\omega\sin\Omega\cos i) & \\ y_{\text{ecl}} &= x_{\text{orb}}(\cos\omega\sin\Omega + \sin\omega\cos\Omega\cos i) - y_{\text{orb}}(\sin\omega\sin\Omega - \\ \cos\omega\cos\Omega\cos i) & \\ z_{\text{ecl}} &= x_{\text{orb}}\sin i + y_{\text{orb}}\cos i \end{aligned}$$

The heliocentric ecliptic Cartesian coordinates are passed to `astropy.coordinates (HeliocentricMeanEcliptic)` and transformed to the geocentric frame (GCRS) at epoch 2026-06-05, yielding geocentric right ascension α , declination δ , and geocentric distance Δ . Galactic latitude b is computed for stellar-crowding screening; samples with $|b| < 15$ lie in the confused galactic plane.

4. Brightness model

Radius from the assumed massradius scaling:

$$R = R_{\oplus} \sqrt{M_p/M_{\oplus}}, \quad R_{\text{km}} = R \cdot 6371$$

Geometric albedo $p_V \sim \mathcal{U}(0.2, 0.75)$ (Revision 3: widened from $\mathcal{U}(0.3, 0.5)$ to match Fortney et al. 2016). Absolute magnitude:

$$H_V = 14.113 - 5\log_{10}(R_{\text{km}}) - 2.5\log_{10}(p_V)$$

Apparent magnitude, using both heliocentric distance d_{hel} and geocentric distance Δ (the d^{-4} reflected-light dependence), with a fixed solar-color offset to r-band:

$$m_V = H_V + 5\log_{10}(d_{\text{hel}}\Delta), \quad m_r = m_V - 0.2$$

5. Existing survey constraints

Before forecasting future facilities, samples already excluded by completed shallow surveys are down-weighted (not deleted, reflecting incomplete sky/depth coverage):

- ZTF ($\delta > -30$, $m_r < 20.5$): surviving weight 0.05 (95% complete)
- Pan-STARRS ($\delta > -30$, $m_r < 21.0$): surviving weight 0.10 (90% complete)
- DES ($-65 < \delta < -35$, $-60 < \alpha < 60$, $m_r < 21.5$): surviving weight 0.05 (95% complete)

Weights are then normalized to sum to one. This implements the published result that prior surveys have already excluded a substantial fraction of the bright, nearby parameter space.

6. Results

6.1 Marginal credible intervals (16th / 50th / 84th percentile, prior-weighted)

Quantity	Median	+84th	16th	10 ⁻⁵	10 ⁻⁹⁵	10 ⁻¹	10 ⁻⁹⁹	Mass ($M_{\oplus}$)	6.51	+2.20	1.46	Semi-major axis a (AU)
556.8	+150.9	113.6	Eccentricity e	0.390	+0.112	0.129	Inclination i (deg)	18.1	+4.6	5.4	Current	
heliocentric distance d (AU)	660.1	+199.2	128.4	Apparent magnitude m_r	22.03	+1.14	0.82					

The median predicted brightness near $m_r \approx 22$ sits squarely in the regime the literature has long anticipated (roughly $22 < V < 25$ near aphelion) and within LSST's single-epoch reach, but the faint tail beyond $m_r \approx 24.5$ is

where detection becomes hard. (Credible intervals are essentially unchanged from Revision 2; the albedo-prior widening shifts them negligibly, consistent with the distance-dominated sensitivity.)

[figure available in web edition] Figure 1. Prior-weighted distribution of predicted r-band apparent magnitude. The bulk lies between 21 and 24, with the faint tail ($m_r > 24.5$) representing the hardest-to-detect orbits.

6.2 Where to look

[figure available in web edition] Figure 2. On-sky probability density (geocentric equatorial). The highest-probability region follows the ecliptic and concentrates in the northern winter sky, roughly RA 3h6h, Dec +10 to +30 the aphelion direction where a body on an eccentric orbit spends most of its time. The white curves mark the galactic plane and its $|b| = 15$ confusion zone. Note that much of this high-probability northern region lies near and above the Dec +12 LSST baseline footprint edge which is precisely why the northern-cadence scheduling choice corrected in Revision 3 is load-bearing for the detection probability.

The probability mass tracks the ecliptic (as expected for a planet) and piles up toward aphelion. Where that arc crosses the galactic plane, detection is degraded a real observational handicap, not a modelling artifact.

6.3 Sensitivity: what drives the uncertainty

| Driver | Raw Pearson r with m_r | Interpretation | | :--- | ---: | :--- | | Heliocentric distance d | +0.95 | Dominant. Sets brightness via d^{-4} . | | Semi-major axis a | +0.75 | Drives d , hence correlated. | | Eccentricity e | +0.33 | Modulates current d given phase. | | Radius R | +0.24 | Weak; massradius scaling is shallow. | | Albedo p_V | 0.09 | Near-zero raw see note. |

To visualize how the orbital inclination parameters affect the detection probability across the LSST survey masks, run the simulation below:

> On the albedo result. The near-zero raw correlation is not a bug and not evidence that albedo is irrelevant. Albedo enters H_V directly and spans ~ 0.55 mag across the sampled range, but heliocentric distance varies so much more that it dwarfs albedo in an unconditioned correlation. Holding distance and radius fixed, the partial correlation of albedo with m_r is 0.998 i.e. at fixed distance, albedo almost entirely determines the residual brightness, exactly as the physics requires. The lesson is that distance uncertainty must be reduced before albedo matters observationally, not that albedo is unimportant.

6.4 When, and by which instrument

Detection year is modelled per facility:

– Vera C. Rubin Observatory / LSST (10-yr survey): [Revision 3] a graded northern-coverage model replaces the Revision 2 hard Dec +15 cutoff. Full WFD cadence applies to Dec +12 (baseline_v2.0); reduced NES-extension cadence with a coverage penalty applies from +12 to +30; low-priority coverage extends to the telescope's physical limit of +33.5; nothing is observable beyond. Depth uses documented Rubin limits (single-visit $r \approx 24.0$, ten-year co-add $r \approx 26.9$): bright targets found in the first 12 seasons, mid-range by years 35, the faintest only via multi-year co-addition late in the survey. Northern (Dec $> +12$) detections carry a reduced-cadence delay; galactic-plane samples carry a 70% per-pass miss penalty.

– Subaru / Hyper Suprime-Cam (targeted deep fields, $m_r < 26.0$): covers the aphelion box (RA 2.5h5.5h, Dec +5 to +25) with $\sim 80\%$ completeness outside the galactic plane, $\sim 30\%$ inside, on a 20262029 timescale.

– Infrared (archival IRAS/AKARI plus NEO Surveyor): models thermal detection for $d_{\text{hel}} < 450$ AU; NEO Surveyor is baselined for launch around late 2027 (no later than June 2028), with a five-year survey running approximately mid-2028 through ~2033.

The combined curve takes the earliest detection across facilities.

[figure available in web edition] Figure 3 (Revision 3). Cumulative detection probability by year, under the corrected northern-coverage model. The combined probability (black) climbs as LSST reaches depth and asymptotes near 0.72 by 2032 up from the 0.61 of Revision 2 with the increment appearing in 2030. In the northern aphelion region the previous hard wall had excluded. The 0.72 value is contingent on plausible NES-extension northern coverage; under no northern extension it remains near the Revision 2 value, so the honest figure is the range ~0.61-0.72.

The combined cumulative probability reaches ~72% by 2032 under plausible northern coverage, then flattens. The genuinely irreducible residual corresponds to orbits that are simultaneously faint ($m_r > 25.5$), deep, off-ecliptic, and galactic-plane-confused a much smaller parameter space than the Revision 2 residual, most of which was the northern coverage artifact rather than a fundamental sensitivity limit.

7. If nothing is found: the null-detection posterior

Treating existence as uncertain, we start from a stated prior $P(\text{exists}) = 0.70$ and apply Bayes' rule given continued non-detection. With $P(\text{null} \mid \text{exists}) = 1 - P_{\text{combined}}(\text{year})$ and $P(\text{null} \mid \text{not exists}) = 1$:

$$P(\text{exists} \mid \text{null}) = \frac{P(\text{null} \mid \text{exists})P(\text{exists})}{P(\text{null} \mid \text{exists})P(\text{exists}) + 1 \cdot (1 - P(\text{exists}))}$$

Through end of	$P(\text{detect} \mid \text{exists})$	Posterior $P(\text{exists} \mid \text{null})$	---	---	---	---		
2026	~0%	70.0%	2028	~50%	53.7%	2030	~67%	43.7%
2032	~71%	40.4%	2036	~72%	39.6%			

[Revision 3] A decade of LSST silence would pull the posterior from 70% down to ~39.6% deeper erosion than the ~47.4% of Revision 2, because the corrected model makes the instruments more capable and therefore makes non-detection stronger evidence against the hypothesis. The qualitative conclusion is unchanged: a null result substantially erodes but does not refute the hypothesis, because a residual fraction of the orbit stays beyond reach. The hypothesis is hard to kill with these instruments even though it is reasonably easy to confirm if the planet sits in the favorable bright/accessible region. Note that this revision strengthens the disconfirmation: continued non-detection argues against Planet Nine more decisively than Revision 2 reported.

8. Limitations and open questions

– Existence is unproven. Everything here is conditional. The clustering signal is statistically suggestive but contested, and observational bias remains a live alternative explanation.

– Distance dominates everything. The single largest lever on detectability is the planet's current heliocentric distance, which the data constrain only weakly. Until that tightens, brightness predictions carry a wide spread and secondary parameters (albedo, radius) cannot be observationally isolated.

– The residual is smaller than Revision 2 stated, and partly a scheduling artifact. [Revision 3] Revision 2 described the residual as structural faintness requiring a 10 m-class facility. A substantial part of it was instead the hard northern coverage cutoff an editable survey-strategy parameter. The genuinely irreducible residual (faint, deep, off-ecliptic, galactic-plane-confused) is smaller, and is closable in large part by Rubin northern-cadence allocation rather than new

hardware.

- Northern-coverage probabilities are reasoned, not sourced. [Revision 3] The graded coverage values (0.5 from +12 to +30, 0.15 above) estimate plausible NES-extension coverage and depend on undecided Rubin scheduling; the ~0.72 detection figure is contingent on them, and the honest headline is the range ~0.610.72.
- Detection-timing model is heuristic. Per-facility detection years are drawn from plausible binned schedules, not a cadence-level survey simulator (e.g. SORCHA/rubin_sim). A full opsim-based treatment would refine the CDF shape.
- Mass-radius relation is a placeholder. $R \propto \sqrt{M}$ is a convenience scaling, not a calibrated relation.
- Static epoch. Positions are evaluated at a single date; parallax motion over a survey baseline (a key real-world discriminator) is not modelled dynamically.

9. Reproducibility

All results derive from a single script, `simulation.py` ($N = 100,000$, numpy seed 42), using numpy, astropy, and matplotlib. The script samples the ensemble, propagates orbits, applies survey masks, and emits the three figures and all printed credible intervals, correlations, and posteriors.

The script, the three output figures, and dependency/run instructions are available at:
<https://github.com/Maha-Strategies/planet-nine-forecast>

References

- Batygin, K., Adams, F. C., Brown, M. E., & Becker, J. C. (2019). The Planet Nine Hypothesis. *Physics Reports*, 805, 153. [arXiv:1902.10103](https://arxiv.org/abs/1902.10103).
- Brown, M. E., & Batygin, K. (2021). The Orbit of Planet Nine. *The Astronomical Journal*. [arXiv:2108.09868](https://arxiv.org/abs/2108.09868).
- Brown, M. E., & Batygin, K. (2016). Observational Constraints on the Orbit and Location of Planet Nine in the Outer Solar System. [arXiv:1603.05712](https://arxiv.org/abs/1603.05712).
- Batygin, K., & Brown, M. E. (2016). Evidence for a Distant Giant Planet in the Solar System. *The Astronomical Journal*.
- Siraj, A., Chyba, C. F., & Tremaine, S. (2025). Orbit of a Possible Planet X. *The Astrophysical Journal*, 978, 139. <https://doi.org/10.3847/1538-4357/ad98f6>
- Fortney, J. J., et al. (2016). The Hunt for Planet Nine: Atmosphere, Spectra, Evolution, and Detectability. *ApJL*, 824, L25. [arXiv:1604.07424](https://arxiv.org/abs/1604.07424).
- Vera C. Rubin Observatory / LSST: survey-strategy and footprint documentation ([arXiv:2303.02355](https://arxiv.org/abs/2303.02355); [arXiv:1812.01149](https://arxiv.org/abs/1812.01149); lsst.org survey-cadence drivers).
- Vera C. Rubin Observatory / LSST documentation, Rubin Observatory.
- NSFDOE Vera C. Rubin Observatory first-imagery release (2025).
- Reproducibility repository (simulation code and figures): <https://github.com/Maha-Strategies/planet-nine-forecast>

Citations are provided for traceability. The Siraj et al. (2025) and Batygin et al. (2019) parameters, and the Brown & Batygin (2021) parameters used in column [b], have been confirmed against their primary publications. The LSST footprint and northern-coverage figures introduced in Revision 3 are sourced to the Rubin survey-strategy literature cited above; the graded northern-coverage probabilities are reasoned estimates and are flagged as such.

Canonical web edition and structured metadata: <https://research.mahastrategies.com/papers/planet-nine-forecast>