

Calibration of CNEOS Fireball Velocities and the Robustness of Nominally Hyperbolic Events

Volkan Duran¹ and Avi Loeb²

¹ Iğdır University, Türkiye

ORCID: 0000-0003-0692-0265 | volkan.duran@igdir.edu.tr

² Astronomy Department, Harvard University, Cambridge, Massachusetts, USA

ORCID: 0000-0003-4330-287X | aloeb@cfa.harvard.edu

Abstract

The NASA/JPL CNEOS Fireball Data catalogue provides a near-global record of bright atmospheric entries but does not report per-event velocity covariance. We calibrate its kinematic accuracy using 19 events with independent reference trajectories and use optical, satellite, radar, and infrasound catalogues as source-specific controls. For nine modern high-quality comparisons, the CNEOS-minus-reference speed residual has mean -1.00 km s^{-1} and standard deviation 0.86 km s^{-1} ; variance decomposition assigns 0.77 km s^{-1} to the CNEOS component. The radiant residuals have a 0.74° per-axis core plus one 16.3° outlier, so the extreme tail remains poorly determined. Seven of 354 complete CNEOS states are nominally hyperbolic without correction. A $+0.9935 \text{ km s}^{-1}$ atmospheric-stage shift produces five additional crossings, but these are treated only as sensitivity cases. Polar-IM (2026 April 1) is the strongest post-2018 CNEOS anomaly, whereas independent GLM-LI stereo solutions yield 57.3 ± 2.0 and $56.7 \pm 3.0 \text{ km s}^{-1}$, close to their direction-specific parabolic thresholds. We conclude that CNEOS can identify priority candidates, but secure interstellar classification requires an independent, uncertainty-bearing three-dimensional trajectory for the same event.

Unified Astronomy Thesaurus concepts: Meteors; Fireballs; Interstellar objects; Orbit determination; Astronomical data analysis

1. Introduction

The discovery of 1I/‘Oumuamua (Meech et al. 2017), 2I/Borisov (Guzik et al. 2020), and 3I/ATLAS (Seligman et al. 2025) demonstrates that macroscopic interstellar bodies traverse the inner Solar System. At meteoroid sizes, atmospheric entry provides a large effective collecting area, but hyperbolic classifications are sensitive to measurement error (Baggaley 2000; Hajdukova & Kornos 2020; Hajdukova et al. 2024). The CNEOS fireball catalogue, derived from U.S. Government sensors, is the only near-global public record of metre-scale atmospheric

impactors and reports three-component Earth-fixed velocity vectors for a subset of events (Brown et al. 2002, 2016). It does not provide per-event velocity uncertainties.

A fireball is classified as interstellar when its reconstructed heliocentric speed exceeds the local solar escape speed, and the margin above that boundary is typically a few kilometres per second — the same order as the plausible measurement error. Classification therefore cannot be read off the catalogue; it is an inference conditional on an assumed error distribution. Peña-Asensio et al. (2025) supplied the first empirical distribution by cross-matching eighteen CNEOS events with independently observed ground counterparts, identifying a post-2018 low-discrepancy regime and quoting $\sigma v = 0.55 \text{ km s}^{-1}$, $\sigma \text{RA} = 1.35^\circ$ and $\sigma \text{Dec} = 0.84^\circ$. Cloete & Loeb (2026a, 2026b) adopted those values, propagated them through 10^6 Monte-Carlo realizations, and reported three post-2018 candidates with $\text{pbound} < 3 \times 10^{-6}$. Concerns about CNEOS velocity accuracy were raised early by Devillepoix et al. (2019), and the difficulty of separating genuinely hyperbolic orbits from measurement error is long established in the meteor literature (Hajduková & Kornoš 2020; Hajduková et al. 2024).

Probabilities of that magnitude are strongly conditioned by the assumed error kernel. Eighteen calibration events can constrain the core of an error distribution but only weakly constrain its extreme tail, and a core-only Gaussian assigns negligible probability to the type of large vector discrepancy that is present in the calibration record. The distinction between a very small formal bound probability and a more modest tail probability materially changes the evidential interpretation of a candidate. Hajduková et al. (2024) reach a different population-level conclusion, finding no evidence for interstellar fireballs in CNEOS, so both event-level and population-level claims require explicit treatment of measurement uncertainty.

Here we recalibrate CNEOS kinematics and evaluate nominally hyperbolic events under an explicit hierarchy of evidence. The analysis (1) stratifies CNEOS-minus-reference residuals by reference provenance and epoch; (2) treats the approximately 1 km s^{-1} atmospheric-stage offset as a sensitivity term rather than a universal correction; (3) constrains the rare-error tail only with direct velocity and radiant comparisons, using GLM geolocation and infrasound back azimuth as external diagnostics; and (4) separates kinematic hyperbolicity from interstellar provenance.

We apply this framework to a frozen CNEOS snapshot and an expanded 19-event orbit-bearing calibration cohort. Results are reported as primary nominal classifications, atmospheric-stage and tail-model sensitivity tests, and independent multi-sensor checks. The wider public-data census is used as a source audit, with each dataset tied to the observable it constrains rather than combined into a common exposure.

2. Data and Samples

2.1 The CNEOS sample

The refreshed public CNEOS API returns 1,069 entries; all event dates and twelve public source fields are unchanged from the preserved 15 August 2026 snapshot used in this study. Of these, 360 carry a three-component Earth-fixed velocity vector and 354 carry the complete combination of position, altitude and vector required for orbit determination. Public-vector coverage rises from 163/753 (21.6%) before 2018 to 197/316 (62.3%) from 2018 onward; a logistic model of vector availability gives an odds ratio of 4.15 (95% CI 2.93–5.89) per decade in reported impact energy and 5.68 (4.52–7.14) per elapsed decade. The published vectors are therefore a selected subset, and all population statements below are conditional on that selection. The Fireball Data API v1.2 labels v_x , v_y and v_z as entry-velocity components and states that all returned fields are relative to the peak-brightness event; this wording is retained as a documented ambiguity in the stage analysis rather than resolved by assumption.

The published 2008 TC3 z-component sign correction is applied only in the analysis field, preserving the source value. Eight of the 357 entries carrying both a scalar speed and a vector differ by more than 1% between the two; both values are retained.

2.2 Orbit-bearing calibration cohort

The calibration cohort comprises nineteen events for which an independent numerical reference orbit or state exists, is frame-compatible with the CNEOS reduction, and refers to an unambiguously identified common event. Eighteen are the frozen benchmark inherited from Peña-Asensio et al. (2025), recomputed here with a single Earth ephemeris; the nineteenth is the 2023 May 20 Queensland superbolide, whose public release (Silber et al. 2026; Zenodo 20739713) states a J2000 geocentric radiant and speed directly and so avoids reconstruction from rounded elements. Admission requires distinct common-event identity, independent reference kinematics, compatible geocentric frames and a numerical reference orbit; it does not require identical observing pipelines.

Each event is assigned to one of three reference classes according to how its reference kinematics were obtained. Class A (pre-impact astrometry, $n = 5$: 2008 TC3, 2018 LA, 2019 MO, 2022 EB5, 2024 RW1) uses a heliocentric orbit determined telescopically before atmospheric entry (Jenniskens et al. 2009; Borovička & Charvát 2009 for 2008 TC3); its velocity uncertainty at Earth is of order $10^{-2} \text{ km s}^{-1}$ and is negligible here. Class N (dedicated fireball network, $n = 9$: Košice, Kalabity, Romania, Baird Bay, Flensburg, Novo Mesto, Ådalen, Iberian, Queensland) uses EFN, DFN, GMN, WMPL or multi-network reductions with published formal errors typically below 0.1 km s^{-1} . Class C (casual or ad-hoc video, $n = 5$: Buzzard Coulee, Chelyabinsk (Brown et al. 2013), Sarıçiçek, Ozerki, Viñales) uses reconstructions from uncalibrated public video. Classes A

and N are jointly designated high-quality (HQ, $n = 14$).

Reference class is assigned solely from the provenance of the independent measurement and is defined before inspection of the residual magnitude.

2.3 Ancillary public data

Table 1 lists the principal numerical inputs of the full census. Their roles are deliberately unequal: only the orbit-bearing cohort calibrates CNEOS event by event, while the large unmatched catalogues characterise their own populations, constrain quality-selection effects, and bound the plausible astrophysical unbound fraction. A million unmatched video meteors are not a million calibrators of a satellite sensor observing a different size and brightness regime.

Product	Source rows / events	Role in this analysis
CNEOS API	1,069 entries; 360 vectors; 354 analysed	Target catalogue; reporting-selection model
Orbit-bearing cohort	19 events (11 from 2018 onward)	Event-level velocity calibration
Global Meteor Network	3,415,540 unique trajectories	Population tails; shower residual controls
SonotaCo annual files	565,902 records	Video-population control; event search
EFN824 / EFN179	824 + 179 (18 shared codes)	Atmospheric-stage measurement; precise-orbit control
PANSY radar Level 3	2,095,187 event IDs	Radar-population tails; covariance audit
IAU MDC (CAMSV3/SAAMER/AMOR)	12,400,435 records	Independent population controls
MAARSY head echoes	1,599,393 IDs	Orbital profile; time/station screening
NASA GLM catalogue	11,451 records; 130 CNEOS matches	Time/location diagnostics; metadata audit
FRIPON public release	1,340 events (1,334 usable)	Formal eccentricity tails
EDMOND-hosted releases	317,830 + 427 + 4,877	Separate orbit-population controls
CAMO/EMCCD release	386 rows	Full-state and covariance admissibility audit
IAU historical catalogues	908 DMS; 8,916 Hissar; 6,345 photographic	Error-field and publisher-flag audits
BLADE	124 events	Light-curve morphology and reporting selection

Product	Source rows / events	Role in this analysis
SOMN–ELFO	71 meteors; 90 acoustic arrivals	Optical consistency; acoustic geometry
Shober et al. Geminid release	584 events $\times$ 4 pipelines	Reported-uncertainty comparability
CEDAR / BRAMS releases	1,604 raw files; 10 paired cases	Source audits; stage-matching demonstration
IAU MDC shower database	1,736 solutions; 946 codes	Shower-anchored calibration test (Section 5.1)
GMN shower-association table	34,711 nodes; 386 showers	Chance-association control for the same test
Interstellar flux-gap constraints	$n_0 = 10^{-4}$ AU $^{-3}$; three published slopes	Astrophysical prior on provenance (Section 5.2)
Archived 2019 CNEOS snapshot	802 events; 208 vectors	Catalogue version stability (Section 4.11)
CNEOS–infrasound event database	362 detections; 137 events	Direction-channel test; impactor properties (Section 5.5)

Table 1. Numerical inputs and their distinct inferential roles. Nested and overlapping samples are never summed into a single exposure. The final five rows are public datasets not used in any previous CNEOS candidate analysis.

3. Analysis

3.1 Orbital reduction and independent validation

For each event the reported geodetic position and Earth-fixed velocity are assembled into an ITRS state at the event epoch and transformed to the Geocentric Celestial Reference System, accounting for Earth rotation, precession, nutation and polar motion. Earth’s gravity is removed with a two-body hyperbolic model: the incoming asymptote direction is reconstructed analytically from the eccentricity and angular-momentum vectors, and the hyperbolic excess speed follows from the geocentric specific energy. Earth’s heliocentric velocity at the event epoch is added, and the heliocentric specific energy is evaluated against the local solar escape speed. The classification statistic is the signed reciprocal semimajor axis of Equation (1), reported in AU $^{-1}$, in which $r_{\oplus}$ is Earth’s heliocentric distance, $v_{\odot}$ the impactor’s heliocentric speed and $\mu_{\odot}$ the solar gravitational parameter.

$$u = \left(\frac{2}{r_{\oplus}} - \frac{v_{\odot}^2}{\mu_{\odot}} \right) \text{ AU}, \quad u < 0 \iff v_{\odot} > v_{\text{esc}} = \left(\frac{2\mu_{\odot}}{r_{\oplus}} \right)^{1/2} \quad (1)$$

This reduction was implemented independently in two computational pipelines and cross-checked over all 354 vector-bearing events: the median absolute difference is 0.018 km s^{-1} in geocentric excess speed and 0.0028 AU^{-1} in u , and both implementations return the same seven nominally unbound states. The heliocentric Earth velocity for the 2026 April 1 epoch, $(+5.203, -26.920, -11.668) \text{ km s}^{-1}$, reproduces the JPL Horizons value tabulated by Cloete & Loeb (2026b) to the quoted precision. The reduction is therefore not the source of the differences discussed below.

3.2 Reference-class stratification

Let $\Delta v_i = v_{g,\text{CNEOS}} - v_{g,\text{reference}}$ for calibration event i . Under the stratified model

$$\Delta v_i = b + \varepsilon_i^{(C)} + \varepsilon_i^{(R,c(i))} \quad (2)$$

$$\varepsilon_i^{(C)} \sim \mathcal{N}(0, \sigma_C^2), \quad \varepsilon_i^{(R,c)} \sim \mathcal{N}(0, \sigma_{R,c}^2), \quad \sigma_{R,A} = 0.02 \text{ km s}^{-1} \quad (3)$$

the reference-side variance $\sigma_{R,c}^2$ is fixed at $(0.02 \text{ km s}^{-1})^2$ for class A and estimated freely for class N. Because class A contributes residuals that are almost pure CNEOS error while class N contributes the sum, the two variance components are separately identified. Parameters are obtained by maximum likelihood from multiple starts. Class C is excluded from the variance decomposition because an uncalibrated-video reference has no defensible parametric error scale; it is retained in all descriptive tables.

The direction channel is treated analogously. A radiant separation $\Delta\theta$ arising from independent two-dimensional Gaussian direction errors of per-axis width σ_θ follows a Rayleigh distribution, whose maximum-likelihood scale is Equation (4). We fit a two-component mixture, Equation (5): a core of width σ_{core} carrying weight $1 - \pi$, and a gross component of width σ_{gross} carrying weight π . With one gross event in nine, π is estimated as the mean of its Jeffreys posterior $\text{Beta}(k + \frac{1}{2}, n - k + \frac{1}{2})$ and its uncertainty is carried explicitly into every downstream probability rather than fixed at a point value.

$$\sigma_\theta = \sqrt{\frac{1}{2n} \sum_{i=1}^n (\Delta\theta_i)^2} \quad (4)$$

$$p(\Delta v, \Delta\theta) = (1 - \pi) \mathcal{N}(\Delta v; b, \sigma_C^2) \text{Ray}(\Delta\theta; \sigma_{\text{core}}) + \pi \mathcal{N}(\Delta v; b, \sigma_g^2) \text{Ray}(\Delta\theta; \sigma_{\text{gross}}) \quad (5)$$

3.3 Atmospheric-stage sensitivity

The CNEOS Fireball Data API states that all fields are relative to the peak-brightness event, while the velocity components are labelled as entry-velocity components. Independent fireball-network references generally report a pre-atmospheric speed. In the EFN824 catalogue (Borovička et al. 2022), 822 events with a valid Vmax give a median $V_\infty - V_{\text{max}}$ of 0.9935 km s^{-1} (95th percentile 4.103; maximum 10.665 km s^{-1}). The close numerical match between this value and the modern CNEOS-minus-reference mean offset is physically suggestive, but it does not establish an event-invariant correction. The primary candidate set therefore uses the uncorrected CNEOS vector. A $+0.9935 \text{ km s}^{-1}$ adjustment, together with nearby alternatives, is reported only as an atmospheric-stage sensitivity analysis.

3.4 Required-error statistic

Probability statements depend on the assumed kernel; a companion statistic should not. For an event with heliocentric speed $v_\odot$ and local escape speed v_{esc} , the boundary in geocentric-asymptote velocity space is a sphere of radius v_{esc} centred on minus Earth’s heliocentric velocity, so the minimum-norm change in the asymptotic velocity that would place the event exactly on the parabolic boundary is simply the margin $|v_\odot - v_{\text{esc}}|$. Mapping back through the gravity-removal step gives the minimum required error in the reported vector,

$$\|\Delta v\|_{\text{min}} = (v_\odot - v_{\text{esc}}) \cdot \frac{v_{\infty, \oplus}}{v_{\text{geo}}} \quad (6)$$

Equation (6) is model-free: it asks how large a CNEOS error would have to be, and it can be compared directly with the error magnitudes actually observed in the calibration cohort, Equation (7), which combines the speed residual with the chord length implied by the radiant separation. We report it alongside the pure-speed route obtained by solving the reduction numerically for the speed change that zeroes the margin.

$$\|\Delta \mathbf{v}_i\| = \sqrt{(\Delta v_i)^2 + \left[2v_{g,i} \sin\left(\frac{\Delta \theta_i}{2}\right) \right]^2} \quad (7)$$

3.5 Monte Carlo propagation

Because one gross radiant event in nine does not identify the extreme tail precisely, model-based tail probabilities are treated as sensitivity outputs rather than as primary evidence. Monte Carlo propagation uses the modern orbit-bearing speed/radiant core and a gross-direction component whose weight is drawn from the Jeffreys posterior of the direct cohort; the Peña-Asensio/Cloete-Loeb kernel and a several-percent cross-modal tail scale are evaluated only as comparisons. Candidate membership is never defined by a numerical $P(\text{bound})$: the primary

event-level quantities are the nominal orbital state, model-free required-error margins, and independent observations of the same event.

3.6 Population-count sensitivity

Population counts are used only as a secondary check. The existing whole-catalogue null simulation is evaluated under the same atmospheric-stage sensitivity scenario used to generate the twelve-state ladder; it is not treated as a test of the seven-event nominal catalogue. The exercise is retained because it shows how easily catalogue-level counts can be dominated by the adopted error model, but event-level and multi-sensor evidence remain the primary basis for inference.

4. Results

4.1 Residuals by reference class and epoch

Table 2 partitions the nineteen-event cohort. The pooled residual scatter is 3.45 km s^{-1} , with four residuals larger than 3 km s^{-1} . Those four occur either against casual-video references or in pre-2018 network comparisons. This pattern shows that the pooled tail cannot be interpreted as a stationary modern CNEOS error distribution. It does not, by itself, prove that every large residual originates on the reference side; rather, it motivates explicit separation by reference class and epoch.

Stratum	n	bias (km s^{-1})	scatter (km s^{-1})	RMS (km s^{-1})	max $ \Delta v $	median $\Delta\theta$ ($^\circ$)	max $\Delta\theta$ ($^\circ$)
all19	19	0.115	3.447	3.357	8.049	2.06	71.21
modern11	11	-0.938	0.781	1.197	2.146	1.27	16.31
pre2018_8	8	1.563	5.055	4.980	8.049	6.24	71.21
A_astrom	5	-0.891	1.209	1.401	2.146	1.27	16.31
N_network	9	0.883	3.091	3.045	8.049	1.30	15.59
C_casual	5	-0.259	5.490	4.917	8.000	6.82	71.21
HQ_all	14	0.249	2.666	2.581	8.049	1.28	16.31
HQ_modern	9	-1.001	0.857	1.286	2.146	1.02	16.31
HQ_pre2018	5	2.500	3.433	3.959	8.049	4.59	15.59
A_modern	4	-1.329	0.817	1.506	2.146	0.93	16.31
N_modern	5	-0.739	0.879	1.079	1.579	1.02	1.63

Table 2. Nineteen-event calibration cohort partitioned by epoch and reference class. The largest pooled residuals are confined to casual-video comparisons and pre-2018 network events; the modern high-quality subset is treated separately for primary calibration.

A variance-components fit to the all-epoch high-quality sample ($n = 14$) assigns 1.13 km s^{-1} to the CNEOS speed component and 3.06 km s^{-1} to the network-reference component. Because the sample is small and the CNEOS instrument regime itself changes with epoch, these values are interpreted as a decomposition of this calibration set rather than as universal instrument constants. Restricted to modern high-quality references ($n = 9$), the fit gives a mean residual of -1.03 km s^{-1} , $\sigma_{\text{CNEOS}} = 0.77 \text{ km s}^{-1}$ and $\sigma_{\text{reference}} = 0.33 \text{ km s}^{-1}$.




Figure 1. Speed and direction residuals by reference class and epoch. Filled symbols denote events from 2018 onward. The figure is used to motivate the modern high-quality calibration subset; it is not used to assign all pre-2018 residuals uniquely to the reference measurement.

4.2 Modern CNEOS error budget

The primary kernel is fitted to the nine modern high-quality events. The speed residual has mean $b = -1.001 \text{ km s}^{-1}$, standard deviation 0.857 km s^{-1} and RMS about zero 1.286 km s^{-1} . The direction residuals contain an eight-event core with $\sigma_{\theta, \text{core}} = 0.736^\circ$ per axis and one 16.31° event (2019 MO) against a pre-impact astrometric reference. That event requires a gross-direction channel, but one event in nine leaves its occurrence rate poorly determined: the Jeffreys 95% interval is 1.2-41.4%. This direct interval, not the cross-modal pooled value, defines the primary uncertainty on rare-event CNEOS-only probabilities.

Relative to the Gaussian kernel used in the recent candidate papers, the modern speed RMS is larger while the angular core is slightly narrower. More importantly, the direct calibration contains a 16.3° radiant outlier that a core-only Gaussian assigns negligible probability. The data therefore support modelling a gross-direction channel, while the small sample precludes a precise estimate of its frequency.




Figure 2. Calibrated kernels against the empirical residuals of the nine modern high-quality events. **Left:** the speed channel; the zero-bias $\sigma = 0.55 \text{ km s}^{-1}$ Gaussian does not reproduce the mean offset or the dispersion of the modern calibration subset. **Right:** the direction channel on a logarithmic axis; the modern calibrated core is narrower than the comparison kernel, while the mixture accommodates the observed 16° case.

No significant dependence of the residual on reported speed is detected within the calibrated range (Spearman $\rho = +0.29$, $p = 0.32$ for $|\Delta v|$ over the high-quality sample; $\rho = -0.23$, $p = 0.55$ over the modern subset). This is a statement about the interval $6.7\text{--}38.6 \text{ km s}^{-1}$ only, and Section 4.8 shows why that restriction matters more than the null result.

4.3 Atmospheric-stage sensitivity

The modern mean CNEOS-minus-reference speed residual (-1.001 km s^{-1}) is almost identical to the EFN824 median $V_\infty - V_{\text{max}}$ (0.9935 km s^{-1}), estimated from a disjoint sample of 822 fireballs. The sign and magnitude are consistent with a measurement-stage contribution. However, the EFN distribution is broad and event deceleration depends on altitude, entry geometry, mass, fragmentation and the precise definition of the reported velocity. We therefore describe the match as supporting evidence for a stage offset, not as an identification of a universal correction.

The primary orbital classification remains the reduction of the published CNEOS vectors without a global speed adjustment: seven of 354 complete states are nominally unbound. Applying $+0.9935 \text{ km s}^{-1}$ as a sensitivity scenario produces twelve nominally unbound states; the five additional events are designated stage-sensitive candidates throughout. Applying $+1.001$ or $+1.329 \text{ km s}^{-1}$ gives twelve and fourteen states, respectively, showing that the marginal population is sensitive to the adopted stage model.

This distinction matters most at high reported altitudes. Polar-IM is reported at 90.5 km and the 2026 January 30 event at 89 km, where the accumulated deceleration may be smaller than the all-event EFN median. A future altitude- and trajectory-resolved model could justify an event-specific adjustment; the present data do not.

4.4 Nominally hyperbolic events

Table 3 reports the seven nominally unbound CNEOS states together with five additional states that cross the parabolic boundary only under the $+0.9935 \text{ km s}^{-1}$ sensitivity adjustment. The latter are retained because they identify useful follow-up targets, but they are not counted as primary detections or used to claim an expanded interstellar population. The table deliberately omits tail probabilities; model-based probabilities are retained only in the separate sensitivity analysis and must be read alongside the required-error and multi-sensor tests.

Status	Event (UTC)	E_imp (kt)	v_geo	v$\odot$ nomi- nal	margin nomi- nal	v$\odot$ +0.9935	margin +0.9935
Nominal	2014-01-08 17:05:33	0.130	45.23	60.95	+18.47	61.84	+19.36
Nominal	2026-04-01 02:13:14	0.086	69.10	51.73	+9.59	52.65	+10.51
Nominal	2017-03-09 04:16:37	1.100	36.70	49.81	+7.53	50.64	+8.36
Nominal	2022-07-28 01:36:07	0.690	29.70	46.98	+5.18	47.89	+6.09
Nominal	2009-04-10 18:42:45	0.720	19.27	44.90	+2.82	46.05	+3.97
Nominal	2021-05-06 05:54:26	0.180	26.81	43.51	+1.58	44.36	+2.43
Nominal	2015-02-17 13:19:50	0.140	28.74	43.37	+0.99	44.15	+1.78
Stage sensi- tivity	2026-01-30 10:25:37	0.120	70.75	42.39	-0.05	43.37	+0.93

Status	Event (UTC)	E_imp (kt)	v_geo	v$\odot$ nomi- nal	margin nomi- nal	v$\odot$ +0.9935	margin +0.9935
Stage sensi- tivity	2022- 07-22 00:16:18	0.210	17.67	41.24	-0.55	42.38	+0.59
Stage sensi- tivity	2015- 03-11 06:18:59	0.230	19.92	41.69	-0.58	42.68	+0.41
Stage sensi- tivity	2025- 06-26 00:37:20	0.390	19.58	40.96	-0.82	42.03	+0.25
Stage sensi- tivity	2013- 07-31 03:50:13	0.220	17.62	40.88	-0.93	41.99	+0.18

Table 3. Primary nominal states and atmospheric-stage sensitivity crossings. The seven rows labelled Nominal are unbound in the uncorrected CNEOS reduction. The five rows labelled Stage sensitivity are bound in the primary reduction and cross the boundary only after a uniform $+0.9935 \text{ km s}^{-1}$ adjustment. The table intentionally omits tail probabilities because candidate membership is defined by the catalogue state and the atmospheric-stage scenario, not by a chosen extreme-tail model.

Among the seven nominal states, the 2014 January 8 event remains the largest-margin catalogue entry. Polar-IM is the most important post-2018 high-speed case because its CNEOS vector has a large nominal hyperbolic margin and an independent stereo measurement exists. CNEOS-22 (2022 July 28) is also reproduced from the current catalogue. The 2009 April 10 and 2015 February 17 events have smaller speed-route error requirements and are correspondingly more fragile.

The 2026 January 30 New Zealand event is not nominally hyperbolic: its uncorrected heliocentric margin is approximately -0.05 km s^{-1} . It becomes unbound only under the atmospheric-stage sensitivity adjustment. Its high geocentric speed, retrograde geometry and the possibility of ground observations nevertheless make it a high-value calibration target, not a primary interstellar candidate.




Figure 3. Conditional candidate-ranking sensitivity. The first series uses the PA25/Cloete & Loeb kernel on uncorrected vectors; the second uses the atmospheric-stage and cross-modal gross-rate sensitivity construction. The five rows labelled stage are bound in the uncorrected catalogue. This figure is diagnostic only and does not define the primary seven-event candidate set.

The 2026 January 30 event over New Zealand deserves separate mention. It is bound in the nominal reduction and crosses the parabolic boundary only under the atmospheric-stage sensitivity adjustment, with a retrograde ecliptic inclination of 168.9° . Public reports indicate that ground observations may exist, but no independent numerical state with reproducible covariance is available in the present dataset. A reduced trajectory would provide a high-speed common event, directly test the stage sensitivity, and extend the modern calibration beyond its current speed range.

Table 4 is retained as a kernel-sensitivity check. It shows that the ranking of the largest-margin events is more stable than the numerical tail probability, and that $P(\text{bound})$ is highly dependent on the assumed gross-direction frequency. The table is not used to select a preferred cross-modal gross-failure rate.

σ_v (km s ⁻¹)		2014-01- π_{gross08}	2026-04- 01	2017-03- 09	2022-07- 28	2009-04- 10
0.55	0.000	0.0000	0.0000	0.0000	0.0000	0.0000
0.55	0.029	0.0000	0.0000	0.0009	0.0008	0.0000
0.55	0.063	0.0000	0.0000	0.0016	0.0019	0.0000
0.55	0.110	0.0000	0.0001	0.0026	0.0034	0.0001
0.77	0.000	0.0000	0.0000	0.0000	0.0000	0.0000
0.77	0.029	0.0000	0.0000	0.0009	0.0008	0.0000
0.77	0.063	0.0000	0.0000	0.0016	0.0020	0.0000

σ_v (km s ⁻¹)	π_{gross08}	2014-01- 01	2026-04- 01	2017-03- 09	2022-07- 28	2009-04- 10
0.77	0.110	0.0000	0.0001	0.0027	0.0035	0.0001
1.29	0.000	0.0000	0.0000	0.0000	0.0000	0.0004
1.29	0.029	0.0000	0.0000	0.0010	0.0010	0.0004
1.29	0.063	0.0000	0.0001	0.0017	0.0022	0.0005
1.29	0.110	0.0000	0.0002	0.0028	0.0040	0.0008
2.00	0.000	0.0000	0.0000	0.0000	0.0001	0.0160
2.00	0.029	0.0000	0.0001	0.0012	0.0013	0.0162
2.00	0.063	0.0000	0.0002	0.0019	0.0027	0.0166
2.00	0.110	0.0000	0.0003	0.0033	0.0050	0.0174

Table 4. Conditional P(bound) for leading events across speed-width and gross-direction sensitivity settings. These are CNEOS-only model checks, not posterior probabilities of interstellar origin.

4.5 Population-count sensitivity

Under the conservative bound-only null used for the +0.9935 km s⁻¹ stage-sensitivity scenario, the calibrated model produces 9.7 apparent unbound states on average (90% range 5-14) against twelve in that sensitivity ladder. The result shows that the number twelve is not, by itself, evidence for an interstellar population. Because the simulation incorporates the same global stage adjustment used to create the twelve-state ladder, it is not a formal null test of the seven-state uncorrected catalogue.

The contamination variant predicts 3.9 false crossings among events that remain bound after the stage adjustment. The gap between the two constructions illustrates the non-identifiability of population counts when the true distance of each event from the parabolic boundary is unknown and the catalogue provides no event-level covariance.




Figure 4. Left: number of nominally unbound states as a function of a global speed adjustment. Right: false-positive count under two null constructions for the same stage-sensitivity scenario. The panel is a sensitivity analysis and is not used to infer an interstellar population fraction.

The population count and the event-level margins answer different questions. The twelve-state stage-sensitivity count is not unusual under the conservative null used here, whereas several nominal events require vector changes larger than those observed in the modern direct calibration. Those event-level margins identify targets for independent measurement; they do not, by themselves, establish interstellar provenance.

4.6 Reproducibility of published candidates

Because every conclusion here depends on a specific catalogue version, we audited the published post-2018 candidates against the current public API. Polar-IM and CNEOS-22 reproduce: our reduction returns heliocentric speeds of 51.73 and 46.98 km s⁻¹ against the published 51.73 and 46.98 km s⁻¹.

CNEOS-25 differs under the current public vector. Cloete & Loeb (2026a) report the 2025 February 12 event (73.4° N, 49.3° E, Barents Sea) at $v_{\odot} = 45.63$ km s⁻¹, exceeding escape by 3.22 ± 0.58 km s⁻¹ with $p_{\text{bound}} < 3 \times 10^{-6}$. The y-component of the velocity vector for that event was flipped in sign from its original value within a day after the preprint by Cloete & Loeb (2026a) was posted online on the arXiv. The revised vector now reads (+9.5, -19.6, -1.9) km s⁻¹, and it yields $v_{\odot} = 37.30$ km s⁻¹ — bound by 5.10 km s⁻¹. Table 5 evaluates all eight sign assignments of the current components: only the two with $v_y = +19.6$ are unbound, and the closest reaches $v_{\odot} = 44.76$ km s⁻¹, still 0.87 km s⁻¹ below the published value.

v_x	v_y	v_z	$v_{\odot}$ (km s ⁻¹)	v_{esc} (km s ⁻¹)	margin (km s ⁻¹)	classification
9.5	19.6	1.9	43.593	42.395	+1.198	unbound
9.5	19.6	-1.9	44.760	42.395	+2.365	unbound
9.5	-19.6	1.9	36.386	42.395	-6.009	bound
9.5	-19.6	-1.9	37.296	42.395	-5.099	bound
-9.5	19.6	1.9	33.379	42.395	-9.016	bound
-9.5	19.6	-1.9	34.250	42.395	-8.146	bound
-9.5	-19.6	1.9	20.931	42.395	-21.464	bound
-9.5	-19.6	-1.9	22.132	42.395	-20.263	bound

Table 5. All sign assignments of the current public CNEOS vector for 2025 February 12. The first row of each pair uses the components exactly as reported in the public CNEOS catalogue snapshot used in this study. Cloete & Loeb (2026a) report $v_{\odot} = 45.63$ km s⁻¹ for this event but the y-component of the posted CNEOS velocity vector was flipped within a day after they posted their analysis, making this trajectory bound to the Solar System.

4.7 Polar-IM: tension with independent stereo measurements

For Polar-IM, the uncorrected CNEOS reduction lies 9.59 km s⁻¹ above local solar escape. The model-free minimum vector change required to reach the parabolic boundary is 9.47 km s⁻¹ (13.7% of the reported speed), and the pure-speed route requires 10.56 km s⁻¹. Under the +0.9935 km s⁻¹ stage-sensitivity scenario those values become still larger. The CNEOS-only state is therefore unusual relative to the modern calibration core, but this fact cannot override an independent measurement of the same event.




Figure 5. Required vector-error magnitudes compared with the direct calibration record. Candidate stars correspond to the atmospheric-stage sensitivity ladder defined by the atmospheric-stage sensitivity analysis described in Section 4.3; the primary uncorrected Polar-IM minimum vector change is 9.47 km s^{-1} . The figure is used as a scale comparison, not as a provenance probability.

Silber et al. (2026) derive two alternative stereo solutions from the same GOES-GLM and MTG-LI observations: 57.3 ± 2.0 and $56.7 \pm 3.0 \text{ km s}^{-1}$, with ECEF vectors $[3.5, -28.3, 49.7]$ and $[3.5, -27.7, 49.4] \text{ km s}^{-1}$, respectively. They report that the velocity directions agree closely with the USG/CNEOS direction while the magnitudes are about 18% lower, and explicitly note that this discrepancy may affect orbital reconstruction. The two solutions are alternative reductions of the same observations and are therefore treated as a model average, not as independent measurements.

A transparent scalar-speed screen illustrates the consequence without claiming a full posterior. For the two published stereo directions, the parabolic speed thresholds are 58.21 and 58.05 km s^{-1} . If the quoted ± 2 and $\pm 3 \text{ km s}^{-1}$ spreads are approximated as Gaussian 1-sigma scales for this purpose, each branch places about one-third of its probability above the parabolic threshold. This approximation is intentionally limited: Silber et al. state that the quoted uncertainties reflect sensitivity to weighting assumptions and navigation errors, and the full trajectory/state covariance is not public. A joint CNEOS-plus-stereo probability would therefore depend on an unvalidated model for the gross CNEOS speed discrepancy and is not reported.

The scalar speeds are also not the whole problem. The stereo directions are part of the orbital solution, and the event lies close enough to the parabolic boundary

that covariance among speed, direction, navigation and deceleration can move appreciable probability across it. The appropriate next calculation is a joint fit to the GLM-LI trajectory geometry with an explicit deceleration model and propagated covariance. Until that analysis exists, Polar-IM should be described as a high-priority unresolved multi-sensor event.

4.8 Calibration-range extrapolation

A limitation more fundamental than any kernel choice constrains the model-based results above. The nineteen direct calibrators span reported geocentric speeds of 6.7 to 38.6 km s⁻¹, with none above 40 km s⁻¹, whereas the two fastest high-priority events lie at 45.2 and 69.1 km s⁻¹. Any tail probability assigned to those events therefore extrapolates the error model beyond the speed range in which it was directly measured.

This matters because the physical arguments run in opposite directions and neither is decided by the data. Faster events are brighter and may be better detected, which can reduce fractional errors; they are also shorter-lived and may present saturation or geometry challenges. Within the calibrated range no speed dependence is detectable, but the sample has little power to exclude one at the speeds that matter. Closing the gap requires common events above 40 km s⁻¹ with independent, uncertainty-bearing trajectories. Polar-IM is one direct example; the 2026 January 30 event would become another if a reproducible ground trajectory and covariance are released.

4.9 Bound-orbit ceiling and direction sensitivity

A simple geometric bound is useful for interpreting the catalogue. A heliocentric orbit bound to the Sun at Earth's distance has $v_{\odot} \leq v_{\text{esc}}$. The geocentric asymptotic speed is maximized when the heliocentric object velocity and Earth's heliocentric velocity are antiparallel, giving $v_{\infty, \oplus} \leq v_{\text{esc}} + v_{\text{Earth}}$, approximately 71.95 km s⁻¹ near 1 au. None of the reconstructed geocentric asymptotic states in the analysed CNEOS sample exceeds this ceiling. Consequently, speed magnitude alone does not establish an unbound orbit for any event in the sample; velocity direction is an essential part of every classification.

The geometric ceiling motivates a direction-channel margin that complements the model-free vector-error statistic. For a fixed geocentric asymptotic speed, the critical angle between the incoming asymptote and Earth's heliocentric velocity can be solved at the parabolic boundary. Table 6 reports this route for the +0.9935 km s⁻¹ stage-sensitivity ladder; it is therefore a geometry sensitivity table, not the primary candidate definition. For the uncorrected Polar-IM state, the corresponding angular change is about 23.0° and the pure-speed change is 10.56 km s⁻¹.

Event (UTC)	$v_{\infty, \oplus}$	fraction of ceiling	θ (°)	θ_{crit} (°)	required radiant error (°)	required speed error (km s ⁻¹)
2014-01-08 17:05:33	44.86	0.617	70.8	114.5	43.69	23.07
2009-04-10 18:42:45	16.91	0.236	18.6	53.2	34.57	3.33
2026-04-01 02:13:14	69.20	0.962	134.8	161.1	26.30	11.55
2017-03-09 04:16:37	36.00	0.498	80.2	100.9	20.74	10.43
2022-07-28 01:36:07	28.60	0.402	68.5	87.7	19.14	6.72
2021-05-06 05:54:26	25.47	0.356	72.7	80.9	8.21	2.84
2022-07-22 00:16:18	14.96	0.210	35.6	40.9	5.23	0.52
2015-02-17 13:19:50	27.57	0.380	80.3	85.6	5.35	2.26
2026-01-30 10:25:37	70.88	0.975	161.2	164.9	3.68	0.94
2013-07-31 03:50:13	14.90	0.209	38.9	40.5	1.55	0.16
2015-03-11 06:18:59	17.70	0.245	55.0	57.2	2.27	0.41
2025-06-26 00:37:20	17.29	0.243	53.1	54.6	1.47	0.23

Table 6. Geometry sensitivity for the +0.9935 km s⁻¹ stage-adjusted ladder. "Fraction of ceiling" is the reconstructed geocentric asymptotic speed divided by $v_{\text{esc}} + v_{\text{Earth}}$. The required radiant and pure-speed changes are separate one-parameter routes to the parabolic boundary. For the primary uncorrected Polar-IM state the corresponding values are approximately 23.0° and 10.56 km s⁻¹.

Under the stage-sensitivity geometry shown in Table 6, Polar-IM lies close to the bound-orbit speed ceiling and can be moved to the parabolic boundary by a smaller relative excursion in direction than by a pure change in speed. In the primary uncorrected state the required radiant rotation is about 23.0°, compared with the 16.31° largest radiant discrepancy in the modern direct calibration, while the required pure-speed change is 10.56 km s⁻¹, compared with a 2.15 km s⁻¹ largest modern speed residual. Both routes therefore extrapolate beyond the directly observed modern calibration. The practical implication is that full three-dimensional uncertainty propagation is essential; a speed-only treatment is

incomplete.




Figure 6. Left: radiant- and speed-route errors required to move each event to the parabolic boundary, compared with the direct calibration record. Right: Jeffreys posteriors for gross discrepancies in the orbit-bearing radiant and GLM position channels; the dashed combined curve is shown only as a cross-modal sensitivity construction and is not adopted by the primary radiant-error model.

4.10 Cross-modal gross-discrepancy diagnostics

The direct orbit-bearing cohort constrains the velocity-direction gross-event frequency only weakly. We therefore examine two larger datasets as external diagnostics of whether several-percent gross solution failures are plausible in other CNEOS observables. The 125 modern CNEOS-GLM common events provide independently measured event positions; at a 100 km threshold, seven are gross position outliers (Jeffreys mean 5.95%, 95% interval 2.54-10.68%).

This position-outlier rate is not a measurement of the radiant-error mixture weight. Likewise, Section 5.5 finds an event-level infrasound bearing-outlier rate near 7%, but back azimuth includes atmospheric propagation and tests source geolocation rather than the velocity radiant. The numerical agreement of these rates is informative as a cross-modal plausibility check, but the Bernoulli events are not exchangeable. We therefore do not pool them into the primary CNEOS velocity-direction kernel.

For transparency, the workbooks retain an alternative pooled 6.45% cross-modal calculation and the GLM threshold grid as labelled sensitivity analyses. Replacing the direct orbit-bearing gross-rate posterior by a several-percent cross-modal value changes the numerical tail probabilities but does not alter the qualitative conclusion for the leading CNEOS-only events. The direct model remains primary because it is the only one calibrated on the same observable used for classification.

4.11 Catalogue version stability

Every result in this paper is conditional on one retrieval of a catalogue that is revised without a changelog, and Section 4.6 had to leave the 2025 February 12 discrepancy unresolved for want of an authenticated earlier copy. One is available. An independently archived scrape of the public CNEOS fireball table, deposited in a public code repository and covering 1988 to 2019 October, preserves 802 events with the same twelve fields, 208 of them carrying velocity components. Comparing it field by field with the raw 2026 API response gives the first empirical measurement of how much this catalogue actually moves.

The comparison is made against the untouched API values, not the analysis fields, so the 2008 TC3 sign correction adopted in Section 3.1 does not enter it; we confirm in passing that the 2026 API still publishes $v_z = +3.8 \text{ km s}^{-1}$ for that event, so the correction of Peña-Asensio et al. (2022) has not been taken up by the catalogue and applying it remains necessary.

Field	Compared	Revised	Rate	Largest change
velocity components v_x, v_y, v_z	154	0	0.00%	—
brightness altitude	294	0	0.00%	—
longitude	516	2	0.39%	345.4° (hemisphere flip)
latitude	516	4	0.78%	5.1°
calculated total	691	175	25.33%	$\times 4.53 / \div 3.42$
impact energy				

Table 7. Field-by-field stability of the public CNEOS catalogue between an archived 2019 snapshot and the 2026 API, over 691 events matched on exact timestamp. Velocity vectors and altitudes are unchanged; impact energies are revised for a quarter of events.

Within the archived pre-2019 comparison subset, none of the 154 comparable velocity vectors and none of the 294 comparable altitudes changed between the 2019 scrape and the 2026 API. This supports frozen-snapshot reproducibility for that historical subset, but it does not establish a zero revision rate for later events. The public API documentation explicitly allows data-format changes, and CNEOS operational notes record updates to existing events. Historical stability should therefore be treated as evidence about the archive, not as proof that a 2025 event could not have been revised.

Second, the calculated impact energies are volatile: 175 of 691 were revised, 128 upward and 47 downward, with a median ratio of 1.073 and extremes of $\times 4.53$ and $\div 3.43$. This matters here because energy is an input, not a bystander. It enters the reporting-selection model of Section 4.1; it defines the 0.45 kt threshold by which Peña-Asensio et al. assign the low-discrepancy regime, and 11 events

(1.6%) cross that threshold between the two versions, so the regime assignment is itself version-dependent for about one event in sixty; and it sets the meteoroid radius that the astrophysical prior of Section 5.2 converts into a number density. Since radius scales as the cube root of energy, the observed revisions imply a version-induced radius factor with a median of 1.02 but extremes of 0.66 to 1.66, and hence a factor of up to 4.5 in the required interstellar number density of Table 9 for the worst-revised events. We record this as a systematic on the provenance calculation rather than propagating it, because the current values are the ones the candidate literature also uses.

Third, timestamps move. Of the 111 events present in the 2019 archive with no exact 2026 counterpart, 107 have a 2026 event within two minutes and are re-timings rather than deletions; only four are genuinely absent, while 134 events were added within the archived span. The re-timings are small — median 1.0 s, 90th percentile 7.8 s, maximum 40 s — and, importantly, none reaches the 60-second window used for every cross-match in this study. The common-event searches of Sections 4.5, 6.2 and 7 are therefore robust to catalogue re-timing, which had been an unquantified risk.




Figure 7. Catalogue stability between the archived 2019 snapshot and the 2026 API. Left: revision rate by field, with counts. Centre: distribution of impact-energy revisions. Right: timestamp shifts for re-timed events, against the 60-second cross-match window used throughout this study.

5. Independent Constraints

The preceding sections exhaust what the orbit-bearing record can supply: nineteen calibrators, none above 38.6 km s^{-1} . Two public datasets that have not entered any previous CNEOS analysis bear directly on what remains. The first tests whether the calibration can be extended by other means. The second supplies the ingredient that converts a likelihood into a statement about origin.

5.1 Shower anchoring as a calibration test

A fireball whose radiant and timing strongly associate it with an established meteor stream has an externally constrained expected pre-atmospheric velocity vector, with a finite stream dispersion. This offers a potential calibration channel without an exact common-event trajectory match, and the major streams span the speed range where the orbit-bearing cohort becomes sparse. We therefore test whether secure stream associations occur often enough in CNEOS to be useful.

We tested this using two public products. The IAU Meteor Data Center shower database supplies 1,736 published solutions across 946 shower codes with solar longitude, radiant, radiant drift and geocentric speed; aggregating to codes carrying at least three independent solutions gives 176 well-established reference streams, whose across-solution scatter is 0.82 km s^{-1} in speed and 2.76° in right ascension. The Global Meteor Network association table supplies 34,711 solar-longitude-resolved nodes for 386 streams with their own association radii (median 3.06°). Neither has been used in a CNEOS calibration.

The design matters. Association uses solar longitude and radiant direction only, never speed; the speed residual is therefore free to take any value and is not truncated by the selection that produced the sample. Chance association is quantified by shifting each event’s solar longitude by a random offset while holding its radiant fixed, which destroys any real membership while preserving both marginal distributions.

No excess above chance is detected under the tested association criteria. Requiring an established stream within 6° of solar longitude and 5° of the drift-corrected radiant associates 3 of the 354 vector-bearing events, compared with 4.82 ± 2.16 expected after randomized solar-longitude shifts ($P = 0.83$). Relaxing to the full 386-stream GMN table associates 8 events compared with 7.6 ± 2.5 expected. Thus the present CNEOS vector sample provides no statistically detectable shower-association excess that could be used to close the high-speed calibration gap.

The result is consistent with the expectation that metre-scale CNEOS impactors are predominantly sporadic and that high-speed shower members at these sizes are rare. The single Southern Taurid association obtained here has a small speed residual, but one association at the chance level cannot calibrate a population. Within the public datasets tested in this study, stream anchoring therefore does not close the high-speed calibration gap; independent trajectories for individual high-speed events remain the more direct route.

5.2 Interstellar flux constraints

Published interstellar-object and meteor-survey constraints provide an important scale against which metre-size candidate rates can be checked. We use the size-distribution families assembled by Peña-Asensio & Seligman (2025) as a

prior-predictive consistency calculation, not as a source of definitive event-level posterior probabilities. The extrapolation from kilometre-scale interstellar objects to metre-scale meteoroids spans orders of magnitude in size and is itself one of the dominant uncertainties. To our knowledge, applying these published size-distribution constraints as a consistency check on the CNEOS candidate set has not previously been carried through in this form.

Peña-Asensio & Seligman (2025) assemble those constraints. Anchoring on 1I/‘Oumuamua and 3I/ATLAS gives a number density of about 10^{-4} AU $^{-3}$ for kilometre-scale interstellar objects. Extrapolating to meteoroid sizes requires a size-distribution slope, and they consider three: $r^{-3} \cdot 0$, connecting spacecraft dust detections to the kilometre-scale objects; $r^{-2} \cdot 7$, the steepest slope compatible with CMOR radar limits, which find five 3-sigma hyperbolic candidates among eleven million orbits; and $r^{-2} \cdot 3$, compatible with optical limits, the Global Meteor Network having found no secure interstellar meteoroid and set an upper limit of about 10^{-6} relative to the interplanetary population.

We convert a candidate’s reported impact energy and speed into a meteoroid mass and radius through Equation (8) at an assumed bulk density $\rho = 1500$ kg m $^{-3}$, evaluate the interstellar number density at that radius under each published slope q , and compute the gravitationally focused Earth encounter rate, Equation (9). Here $n_0 = 10^{-4}$ AU $^{-3}$ at $R_0 = 220$ m is the kilometre-scale anchor, v_∞ the encounter speed and $v_{\text{esc},\oplus}$ Earth’s escape speed. Table 8 gives the result for the sample as a whole.

$$m = \frac{2E}{v^2}, \quad R = \left(\frac{3m}{4\pi\rho} \right)^{1/3} \quad (8)$$

$$n(> R) = n_0 \left(\frac{R_0}{R} \right)^q, \quad \dot{N} = n(> R) v_\infty \pi R_\oplus^2 \left(1 + \frac{v_{\text{esc},\oplus}^2}{v_\infty^2} \right) \quad (9)$$

Slope	n(>R) at R = 0.3 m (AU $^{-3}$)	Expected interstellar events, 2018-2026 CNEOS vector sample
$r^{-2} \cdot 3$	389	6.0e-06
$r^{-2} \cdot 7$	5.45e+03	5.2e-05
$r^{-3} \cdot 0$	3.94e+04	2.6e-04

Table 8. Interstellar meteoroid abundance at the CNEOS size scale under the three published size-distribution slopes, and the resulting expected number of genuine interstellar events in the modern vector-bearing sample. Even the steepest slope predicts fewer than 10^{-3} events.

The resulting expected event counts are far below unity under the published extrapolations. This establishes a strong population-level tension that any interstellar interpretation must address, but it does not by itself decide whether an individual CNEOS state is mismeasured or interstellar.

5.3 Exploratory provenance sensitivity

For completeness, the workbooks retain a speed-only Bayes-factor construction that combines a calibrated CNEOS speed likelihood with alternative bound and interstellar speed priors. We now treat this calculation as exploratory because Section 4.9 shows that CNEOS hyperbolic classification depends on both speed and radiant direction. A one-dimensional speed Bayes factor cannot be interpreted as a full provenance Bayes factor when the decisive likelihood information lies in the three-dimensional state vector.

$$BF = \frac{\int_U L(v_{\text{obs}} | v) \pi_U(v) dv}{\int_B L(v_{\text{obs}} | v) \pi_B(v) dv} \quad (10)$$

Under this exploratory construction, introducing a gross-error component places a finite ceiling on the evidence supplied by an extreme speed residual. The numerical values are useful for sensitivity analysis, but they are not used in the abstract, candidate ranking or conclusions as probabilities of interstellar origin.

The robust inference from this exercise is qualitative: a non-zero gross-error channel prevents arbitrarily large evidence ratios from a single catalogue vector. A defensible event-level Bayes factor will require the joint speed-direction likelihood and, for Polar-IM, the independent stereo geometry.




Figure 8. Exploratory speed-only evidence sensitivity and size-distribution consistency. The vertical sensitivity marker uses the cross-modal gross-outlier

value only to expose model dependence; it is not the primary radiant-error rate, not a full-state Bayes factor, and not a provenance determination.

5.4 Required-density sensitivity

Table 9 is retained as an order-of-magnitude stress test. The most interpretable quantity is the number density required for posterior odds of order unity under the stated speed-only construction, compared with the density supplied by each published size-distribution extrapolation. The columns labelled P(IS) are model-dependent sensitivity outputs, not calibrated event-level posterior probabilities, and no detection claim rests on them.

$$O_{\text{post}} = \frac{N_{\text{IS}}}{N_{\text{ev}}} \cdot BF, \quad n_{\text{req}} = \frac{N_{\text{ev}}}{\bar{N}_1 T \cdot BF} \quad (11)$$

Event (UTC)	R (m)	Bayes fac- tor	P(IS) $r^{-2} \cdot 3$	P(IS) $r^{-2} \cdot 7$	P(IS) $r^{-3} \cdot 0$	n re- quired (AU ⁻³)	n max pub- lished	shortfall
2014-01-08 17:05:33	0.44	2369	0.0014	0.016	0.097	1.18e+05	1.26e+04	9.36
2026-04-01 02:13:14	0.29	2.374e+04	0.0029	0.3	0.76	1.42e+04	4.44e+04	0.32
2017-03-09 04:16:37	1.03	40.89	2.7e-06	2.3e-05	0.00012	8.46e+06	979	8.64e+03
2022-07-28 01:36:07	1.01	8.024	2.5e-07	2.1e-06	1.1e-05	9.63e+07	1.02e+03	9.42e+04
2009-04-10 18:42:45	1.37	1.939	5.7e-08	4.4e-07	2e-06	2.07e+08	412	5.01e+05
2021-05-06 05:54:26	0.69	1.53	1.1e-07	1.1e-06	6.3e-06	5.05e+08	3.19e+03	1.58e+05
2015-02-17 13:19:50	0.61	1.249	2.4e-07	2.5e-06	1.5e-05	3.21e+08	4.72e+03	6.8e+04
2022-07-22 00:16:18	0.96	0.01648	5.7e-10	5e-09	2.5e-08	4.69e+10	1.19e+03	3.94e+07
2026-01-30 10:25:37	0.32	3421	0.0035	0.046	0.26	9.61e+04	3.34e+04	2.88
2015-03-11 06:18:59	0.92	0.01782	1.3e-09	1.2e-08	6.1e-08	2.25e+10	1.38e+03	1.63e+07
2013-07-31 03:50:13	0.98	0.0073954	7e-10	4.1e-09	2.1e-08	5.42e+10	1.13e+03	4.8e+07
2025-06-26 00:37:20	1.11	0.01256	3.1e-10	2.6e-09	1.3e-08	6.15e+10	786	7.82e+07

Table 9. Exploratory provenance sensitivity. Required densities and model-dependent speed-only posterior quantities are shown to expose the astrophysical scale of the assumption. These values are not full-state posterior probabilities because radiant uncertainty and, for Polar-IM, the independent stereo state are not jointly modelled.

Within this exploratory calculation, Polar-IM is the only event whose required metre-scale density overlaps the steepest published size-distribution extrapolation, while the 2026 January 30 and 2014 January 8 events are closer than the remaining candidates. This ranking is useful for follow-up prioritisation, not as evidence that any event is interstellar.

A third hypothesis also deserves a place, because the analysis so far has been framed as a choice between an interstellar origin and a measurement error. Peña-Asensio et al. (2024b) show that Oort-cloud perturbations can deliver Earth impactors on hyperbolic orbits, which is a Solar System origin producing the same nominal signature. Such bodies would arrive with small positive excess speeds rather than the tens of kilometres per second an interstellar population implies, so the hypothesis competes for the marginal candidates rather than for the extreme ones. In the ranking of Table 9 the six events with heliocentric excess speeds below 10 km s^{-1} are precisely the ones it could absorb; Polar-IM, at 31.6 km s^{-1} , and the 2014 January 8 event, at 44.9 km s^{-1} , are not. We do not attempt to quantify this channel, which would require its own flux model, but we record that the weak end of the candidate list has an astrophysical explanation that is neither interstellar nor instrumental.

The size-distribution calculation is dominated by model uncertainty: the steepest slope is itself constrained by meteor-survey limits, the kilometre-scale anchor is based on very few detected objects, impact-energy revisions propagate into inferred radii, and the event-level likelihood is incomplete. We therefore use the calculation only to identify which candidates are astrophysically least implausible under published extrapolations.

the theoretical bearing is a great-circle bearing to the reported position, this residual contains acoustic propagation deviation as well as any CNEOS error, so it bounds the CNEOS contribution from above rather than isolating it.

Taking the median residual per event to avoid counting one bolide several times leaves 121 events, of which eight exceed 10° . The corresponding Jeffreys mean is 6.97% (95% interval 3.18-12.09%). This is numerically similar to the GLM position-outlier rate but it is not an estimate of the CNEOS radiant-error mixture weight: the observable is back azimuth to the acoustic source and the residual also contains atmospheric propagation. The cross-modal pooled value (16/255, 6.45%) is retained in the workbooks only as a labelled cross-modal sensitivity calculation and is not adopted by the primary model.




Figure 10. Left: infrasound observed-minus-CNEOS-derived back-azimuth residuals for 331 detections. Right: separate gross-outlier-rate posteriors for the orbit-bearing radiant, GLM position and acoustic bearing channels. Their similarity is shown as a cross-modal comparison; the primary analysis does not assume a common failure probability.

Four nominal or stage-sensitive events appear in the acoustic database. IM2 was detected by two arrays with substantially different viewing azimuths, and both bearings are consistent with the reported CNEOS source location. This is valuable evidence against a gross geolocation error for that event, but it does not validate the CNEOS velocity radiant.

Event (UTC)	Rank	Detection	Station range (km)	Observed – theoretical bearing ($^\circ$)	Required radiant error ($^\circ$)
2014-01-08 17:05:33	1	2	1750, 2525	+0.35, -2.48	43.7

Event (UTC)	Rank	Detection	Station range (km)	Observed – theoretical bearing (°)	Required radiant error (°)
2009-04-10 18:42:45	5	1	1788	−4.82	34.6
2015-03-11 06:18:59	10	1	1703	−0.76	2.3
2013-07-31 03:50:13	11	1	1511	+4.60	1.6

Table 10. Candidates with independent acoustic bearings. Back azimuth tests the source location, not the velocity direction; the table therefore constrains gross geolocation failure without directly adjudicating hyperbolic provenance.

The inference must be stated precisely. A back azimuth constrains the bearing from an array to the acoustic source and therefore tests source geolocation, not the velocity radiant. For IM2 the two-array agreement is evidence against a gross source-location error, but a velocity-direction error could in principle coexist with an accurate location. For the marginal 2013 and 2015 events, the required radiant changes are smaller than the practical angular precision of this acoustic diagnostic. We therefore use infrasound as an external geolocation check only.

Finally, the infrasound release supplies per-event impactor-property estimates that illustrate the sensitivity of the exploratory size-distribution calculation to assumed physical properties. For IM2 the released density/radius choice changes the required-to-available interstellar density ratio in that exploratory model from 9.36 to 1.53. We do not interpret this as a provenance determination: the density is itself model-derived, the size-distribution extrapolation is uncertain, and the acoustic bearing constrains source location rather than velocity direction. The result is retained only to show that provenance calculations can be as sensitive to assumed impactor properties as to the statistical tail model.

The release additionally carries luminous begin and end altitudes for 86 events, median 40.3 and 25.6 km with a median span of 13.8 km. These are the inputs an altitude-resolved atmospheric-stage model would need, which Section 4.3 identified as the most obvious refinement available; we record their availability rather than building that model here.

5.6 Structure of acoustic residuals

Three hundred and thirty-one paired bearings are enough to ask not only how large the direction error is but what it depends on, and the answers bear on the two limitations this study has been unable to remove by other means. Figure 11 collects the three tests.

The first test concerns speed dependence in the geolocation diagnostic. The acoustic sample carries a CNEOS reported speed for 271 detections, spanning

9.8 to 44.8 km s⁻¹, and the absolute back-azimuth residual shows no detectable dependence on it (Spearman rho = -0.039, p = 0.52). This is useful evidence that source-geolocation agreement does not obviously worsen over the sampled speed range, but it is not a validation of the velocity radiant and it does not close the direct-calibration gap above 38.6 km s⁻¹. Only eight detections lie above 30 km s⁻¹ and none approaches the 69 km s⁻¹ of Polar-IM, so high-speed radiant extrapolation remains an explicit limitation.

The second test concerns what the back-azimuth residual contains. Acoustic propagation through a structured, windy atmosphere can deviate ray paths (Silber 2024), and the residual increases with station range. An illustrative variance decomposition of five equal-count distance bins, $\sigma(D)^2 = \sigma_0^2 + (aD)^2$, gives $\sigma_0 = 3.45^\circ$ and $a = 0.56^\circ$ per 1000 km, with a propagation-scale contribution of 1.58° at the median range of 2820 km. Detections beyond 10° also become more frequent in the farthest bins. We therefore interpret the fitted distance-independent term only as an upper bound on range-independent geolocation/systematic scatter, not as a uniquely identified CNEOS error component. This distance structure is an additional reason not to convert the acoustic outlier fraction into a radiant-error mixture weight.

A third check controls for the geometry effect that motivated this dataset in the first place. Bolide entry angle governs how infrasound is radiated and detected, and steeper entries are preferentially detected (Ronac Giannone & Silber 2026), and the apparent arrival direction has been examined directly (Silber 2025); if the apparent arrival direction also depended systematically on entry angle, part of our residual would be that effect rather than CNEOS error. It does not: the absolute residual is uncorrelated with entry angle over the 271 detections that carry one ($\rho = -0.069$, p = 0.26).

Finally, 86 events carry luminous begin and end altitudes alongside the CNEOS reported altitude. The reported altitude lies at a median 44.5% of the way down the luminous path (interquartile range 30-58%), with median begin, reported and end altitudes of 40.3, 33.3 and 25.6 km. This supports the general relevance of measurement stage, but it still does not identify the epoch represented by the published velocity components, which the CNEOS documentation labels as entry velocities while describing the returned fields as relative to the peak-brightness event. The result therefore strengthens the case for an event-resolved stage model, not for a universal +0.9935 km s⁻¹ correction.




Figure 11. Structure in the acoustic bearing residuals. Left: absolute back-azimuth residual versus reported speed, showing no detectable trend over the sampled range. Centre: robust angular scale versus station range with an illustrative range-independent plus propagation-scale fit. Right: location of the CNEOS reported altitude within the luminous path for 86 events. The right panel concerns altitude geometry and does not determine the epoch represented by the published velocity components.

6. Comparison with Previous CNEOS Calibrations

6.1 Interpretation of published calibration widths

The difference between our calibration and the previously adopted kernel can be traced to how the published summary statistics are interpreted, and this distinction is visible in the source paper’s own table rather than in any disagreement about the underlying data. Table 4 of Peña-Asensio et al. (2025) is captioned “Median and (upper and lower) 1σ standard deviations of the DD values and orbital element, velocity, and geocentric radiant errors for the 18 calibrated fireballs from CNEOS”. For the low-discrepancy group it reports the entries in Table 11 below: a median velocity error of 0.55 km s^{-1} with an asymmetric spread of $+0.37/-0.45$, a median right-ascension error of 1.35° with $+1.12/-0.79$, and a median declination error of 0.84° with $+1.00/-0.79$.

Quantity	PA25 low-discrepancy median	1σ spread	Use in Cloete & Loeb (2026a,b)	Implied σ if the errors are half-normal	This work, measured
velocity error	0.55 km s ⁻¹	+0.37 / -0.45	Gaussian σ	0.82 km s ⁻¹	1.29 km s ⁻¹ (RMS)
right-ascension error	1.35°	+1.12 / -0.79	Gaussian σ	2.00°	—
declination error	0.84°	+1.00 / -0.79	Gaussian σ	1.25°	0.74° per axis (core)

Table 11. The published low-discrepancy uncertainties of Peña-Asensio et al. (2025), their interpretation in recent candidate analyses, and the illustrative effect of treating a median error as a Gaussian scale. The half-normal column divides the median by 0.6745; it is illustrative, not a refit.

These entries are medians of positive per-event error magnitudes accompanied by asymmetric spreads, rather than a fitted zero-mean Gaussian scale parameter. Treating the median error magnitude itself as σ in $N(0, \sigma^2)$ is therefore an additional modelling assumption. As an illustration only, if a positive error magnitude were half-normal, a median of 0.55 km s⁻¹ would correspond to a parent Gaussian sigma of 0.82 km s⁻¹. That value is about 64% of the 1.29 km s⁻¹ modern high-quality RMS measured here. The calculation does not prove a half-normal model; it shows why the numerical value 0.55 km s⁻¹ should not automatically be interpreted as a Gaussian standard deviation.

The asymmetric spreads make the same caution independently: the reported error magnitudes are not described by a single symmetric scale around zero. They motivate sensitivity to non-Gaussian tails, but they do not by themselves identify the two-component gross-error mixture used in this work; that mixture is motivated separately by the observed 16.3° radiant outlier in the modern direct cohort.

This distinction is methodological rather than a criticism of the source calibration or of the analyses that built on it. Peña-Asensio et al. report the measured summary statistics transparently, and using a published central value as a kernel width is a reasonable modelling choice when a full event-level likelihood is unavailable. The distinction becomes important only when the caption and asymmetric spreads are considered together. Peña-Asensio et al. themselves caution that hyperbolic CNEOS candidates should be interpreted carefully, note that CNEOS reports few fast fireballs, and record that more than 22% of CNEOS events show significant inaccuracies in a random sample. Our contribution is to

examine the consequences of alternative kernel interpretations and to propagate those uncertainties explicitly.

6.2 Cross-modal context for gross discrepancies

Previous studies consistently caution that a non-negligible subset of CNEOS solutions can contain large discrepancies, especially in radiant or orbital quantities. Our direct orbit-bearing sample supplies one modern 16.3° radiant outlier in nine comparisons, while the GLM position and infrasound bearing datasets independently show several-percent gross-outlier rates in different observables. These channels support the need for robust-tail sensitivity analyses, but their rates should not be interpreted as repeated measurements of one universal failure probability.

Hajduková et al. (2024) find no evidence for interstellar fireballs in the CNEOS database. At the population level our stage-sensitivity count is compatible with that reading, and the published interstellar size distributions predict far fewer than one genuine event in the modern vector-bearing sample. At the event level, however, several CNEOS states have large required-error margins under the direct calibration. We interpret those states as high-priority anomalies requiring independent measurements, not as events that measurement error has been ruled out for. This distinction makes the population-level null result and the event-level follow-up ranking compatible.

A footnote in the same paper also corroborates Section 4.11 from an unexpected direction: the authors record that during their publication process CNEOS updated the timestamps of half of the hyperbolic events by one second and revised energies. That is precisely the behaviour our archived-snapshot comparison measures, and our median re-timing of 1.0 second matches their observation exactly.

6.3 Relation to recent candidate analyses

With that distinction, our relationship to Peña-Asensio et al. (2025) is narrower and more reproducible. We confirm the post-2018 improvement in CNEOS fidelity, separate reference provenance and observing epoch, distinguish speed from radiant error, and propagate uncertainty in the direct gross-direction rate. We also identify a mean speed offset compatible with atmospheric-stage differences, but treat the corresponding global correction as a sensitivity analysis rather than a primary catalogue modification.

For Cloete & Loeb (2026a, 2026b), we recover CNEOS-22 and Polar-IM from the current catalogue and reproduce their CNEOS-only reductions under the published assumptions. The independent Silber et al. (2026) GLM-LI solution then provides an additional same-event constraint: its directions are similar but its speeds are about 18% lower, placing the event close to the parabolic speed threshold for those directions. Under a simplifying one-dimensional Gaussian

screen of the quoted stereo spreads, roughly one-third of the stereo probability lies on the unbound side. Because the stereo covariance and a validated joint gross-error likelihood are unavailable, we do not report a combined CNEOS-plus-stereo posterior. Taken together, the available public data motivate treating Polar-IM as a high-priority follow-up target whose final dynamical classification remains open pending a full uncertainty-bearing stereo reconstruction.

The comparison with Barghini et al. (2025), who examined hyperbolic-orbit confidence levels in ground catalogues, runs the other way: their point that formal eccentricity errors near the parabolic boundary are unreliable applies with greater force to a catalogue that publishes no errors at all, and is one motivation for the model-free required-error statistic used here.

7. Public-data Controls

We assembled a broad census of publicly accessible meteor, fireball, radar, satellite and infrasound products to determine which sources can constrain each part of the inference. Heterogeneous catalogues are not pooled into a single exposure; each is used only for the observable it actually measures. This section summarises the census, while the companion workbooks retain the source-level audits and provenance notes.

7.1 Quality selection and nominal hyperbolic fractions

Nominal hyperbolic fractions in every large catalogue depend strongly on quality selection, which is the first reason unmatched catalogue counts cannot calibrate CNEOS. In GMN (Vida et al. 2020, 2021), 280,201 of 3,415,540 trajectories are nominally hyperbolic (8.20%), falling to 3.66% under the study quality screen. In SonotaCo (SonotaCo et al. 2021) the fraction falls from 11.06% to 4.98%. In PANSY (Vierinen 2026) it falls from 4.24% to 0.65% across explicit Level 3 selections; in FRIPON (Colas et al. 2020) from 10.12% to 1.82%, with no FRIPON solution more than three reported eccentricity standard deviations above unity. CAMSv3 gives 11.88% falling to 7.12%; SAAMER 10.73% and AMOR 10.00%, neither release carrying the per-record errors needed for a formal test. In the speed-matched band $10 \leq v_g < 30 \text{ km s}^{-1}$ the CNEOS nominal rate is 4/240 (1.667%, Wilson 0.650-4.206%) against 0.336% for GMN and 0.105% for quality-screened PANSY. These ratios are descriptive comparisons of records produced by different instruments, populations and reduction chains; they do not measure the CNEOS false-positive rate.

7.2 Common-event diagnostics

The GLM bolide catalogue (Smith et al. 2021; Ozerov et al. 2024) supplies 130 one-to-one CNEOS matches within 60 s and 300 km, of which 69 have a public CNEOS vector. Twenty-three catalogue-level matches have positions exactly equal to rounded CNEOS coordinates, so catalogue-level independence cannot

be assumed; using native satellite attachment positions instead gives a median event-level horizontal separation of 39.46 km. This is a location diagnostic, not a velocity calibrator, and it is used only as such.

7.3 Product-level propagation limits

Several releases were found to be unusable for the specific purpose they would naturally serve in this analysis, and recording those limitations is part of the source audit. The distributed CAMO/EMCCD covariance arrays (Western Meteor Physics Group 2026) fail the stated numerical admissibility checks for nearly all rows as parsed here, so those covariances are not propagated. The four-pipeline Geminid comparison of Shober et al. (2026) shows that reported velocity uncertainties for the same events can differ substantially between pipelines despite similar trajectory RMS values, cautioning against pooling formal uncertainties across reduction systems. MAARSY/PANSY timing fields (Vierinen 2025, 2026) also require source-specific interpretation. These statements concern the released products as ingested in this study, not the underlying observations or unpublished processing.

7.4 Held-out event checks

Three event summaries are held out of the calibration and used only as checks. For the 14 August 2026 Washington event the reported starting speed is 14.6 km s^{-1} against a CNEOS component norm of 22.899 km s^{-1} , with beginning altitude 77.9 km against a CNEOS brightness altitude of 30 km — a discrepancy whose sign is opposite to ordinary deceleration. For the 17 March 2026 Ohio event the values are 17.5 and 14.915 km s^{-1} , a sign compatible with slowing. The 2023 May 20 Queensland event supplies the nineteenth calibrator with a speed residual of -1.579 km s^{-1} and a radiant separation of 1.295° , both squarely within the calibrated modern core, and is the strongest single piece of evidence that the modern CNEOS vector is accurate at the level claimed here. Florida/Grand Bahama (Hughes et al. 2022) supplies a twentieth speed-channel comparison at -3.351 km s^{-1} but its tabulated radiant equinox is not explicit, so it is not admitted to the frame-harmonised cohort.

8. Discussion

The direct calibration indicates that modern CNEOS vectors are substantially more accurate than the pooled all-epoch residual distribution suggests, but the sample is too small to determine the frequency of large radiant errors. The nine-event modern high-quality subset has a narrow core, while the 16.3° 2019 MO discrepancy requires an explicit gross-error channel. GLM geolocation and infrasound bearing residuals show that percent-level gross discrepancies occur in other observables, but they cannot replace the direct radiant-error rate.

The approximately -1 km s^{-1} mean CNEOS-minus-reference offset is consistent

in sign and magnitude with the independent EFN pre-atmospheric-to-luminous-stage deceleration scale. However, the CNEOS field description, the broad EFN deceleration distribution, and the high reported altitudes of the fastest events do not support a single event-invariant correction. We therefore retain the uncorrected seven-event nominal set and report the five additional $+0.9935 \text{ km s}^{-1}$ crossings only as sensitivity cases.

Polar-IM provides the most informative event-level test because the CNEOS solution and an independent GLM-LI stereo reconstruction disagree primarily in speed. The stereo directions remain close to CNEOS, whereas the reported mean speeds lie near their direction-specific parabolic thresholds. A one-dimensional screen places approximately one-third of the stereo probability on the unbound side if the quoted spreads are treated as Gaussian scales; without the full stereo covariance, the measurement discrepancy is more informative than a refined CNEOS-only tail probability.

The principal observational requirement is a high-speed common event with an independent, uncertainty-bearing three-dimensional trajectory. The 2026 April 1 and 2026 January 30 events are particularly valuable because both lie beyond the 38.6 km s^{-1} upper limit of the direct calibration cohort. Infrasound provides complementary source-geolocation and energetics constraints, but back azimuth alone is not an independent radiant measurement.

The highest-priority follow-up measurements are a covariance-aware reconstruction of the 2026 April 1 GLM-LI stereo event with an explicit deceleration model, recovery of a contemporaneous CNEOS-25 payload, an independent trajectory for the 2026 January 30 New Zealand event, and additional high-speed CNEOS common-event searches. Publication of per-event CNEOS uncertainty or quality metadata would directly reduce the dominant instrumental ambiguity.

Broader implications for impact-energy or flux estimates should await an event-resolved stage model. The present analysis establishes that a measurement-stage contribution is plausible and potentially important; it does not show that every CNEOS speed requires the same $+1 \text{ km s}^{-1}$ adjustment.

9. Summary and Conclusions

The principal results are as follows.

- (1) For the modern high-quality calibration subset ($n = 9$), the speed residual has mean -1.00 km s^{-1} and standard deviation 0.86 km s^{-1} , with a variance-component estimate of 0.77 km s^{-1} for the CNEOS contribution. Radiant residuals consist of a 0.74° per-axis core plus one 16.3° gross event; consequently, the extreme direction-error rate remains weakly constrained.
- (2) The mean speed offset is consistent with the EFN pre-atmospheric-to-luminous-stage deceleration scale, but the available metadata do not justify

a universal $+0.9935 \text{ km s}^{-1}$ correction. Seven nominally hyperbolic states therefore define the primary catalogue result; five additional crossings are retained as atmospheric-stage sensitivity cases.

(3) Polar-IM is the strongest post-2018 CNEOS anomaly but is not independently confirmed as hyperbolic. The two GLM-LI stereo solutions give 57.3 ± 2.0 and $56.7 \pm 3.0 \text{ km s}^{-1}$, close to their direction-specific parabolic thresholds.

(4) GLM geolocation and infrasound back-azimuth records show several-percent gross discrepancies in their own observables, but they are not direct estimates of the CNEOS velocity-radiant tail. The archived 2019 catalogue snapshot also shows field-dependent revision behaviour, making frozen payloads and retrieval metadata essential for reproducibility.

Taken together, the public record can rank high-priority follow-up events and identify the assumptions controlling their classification, but it does not establish an interstellar meteor from a single CNEOS vector. For Polar-IM, the next decisive step is a joint, uncertainty-bearing three-dimensional reconstruction of the independent stereo observations.

Acknowledgments

Facilities and data products: NASA/JPL CNEOS Fireball API; GOES Geostationary Lightning Mapper (GLM); Meteosat Third Generation Lightning Imager (LI); public infrasound-array products described in the cited releases.

Generative AI tools, including OpenAI’s ChatGPT and Anthropic’s Claude Opus, were used during manuscript preparation for language refinement, organizational assistance, code drafting, computational cross-checking, and consistency checking. Source provenance, intermediate results, and relevant computational outputs were retained for verification.

Data Availability

The companion workbooks separate the primary seven-event nominal set from atmospheric-stage sensitivity cases, direct radiant-error calibration from cross-modal diagnostics, and CNEOS-only quantities from the Polar-IM stereo screen. Source provenance, random seeds, required-error statistics, catalogue-version checks, and exploratory prior-sensitivity calculations are retained for reproducibility.

References

Baggaley, W. J. (2000). Advanced Meteor Orbit Radar observations of interstellar meteoroids. *J. Geophys. Res.* 105, 10353.

- Barghini, D., et al. (2025). The Kresáks’ diagram: hyperbolic meteoroid orbits and their confidence level. *A&A* 701, A135.
- Borovička, J. & Charvát, Z. (2009). Meteosat observation of the atmospheric entry of 2008 TC3 over Sudan and the associated dust cloud. *A&A* 507, 1015.
- Borovička, J., Spurný, P., Shrbený, L., Štork, R. & Kotková, L. (2022). Data on 824 fireballs observed by the digital cameras of the European Fireball Network in 2017–2018. *A&A* 667, A157.
- Brown, P. G. & Borovička, J. (2023). On the proposed interstellar origin of the 20140108 fireball. *ApJ* 953, 167.
- Brown, P. G., et al. (2013). A 500-kiloton airburst over Chelyabinsk and an enhanced hazard from small impactors. *Nature* 503, 238.
- Brown, P., Spalding, R. E., ReVelle, D. O., Tagliaferri, E. & Worden, S. P. (2002). The flux of small near-Earth objects colliding with the Earth. *Nature* 420, 294.
- Brown, P., Wiegert, P., Clark, D. & Tagliaferri, E. (2016). Orbital and physical characteristics of metre-scale impactors from airburst observations. *Icarus* 266, 96.
- Cloete, R. & Loeb, A. (2026a). Two robust interstellar meteor candidates in the post-2018 CNEOS fireball database. *arXiv:2602.08956*.
- Cloete, R. & Loeb, A. (2026b). A polar interstellar meteor candidate from 04-01-2026 in CNEOS. *arXiv:2606.04379*.
- Colas, F., et al. (2020). FRIPON: a worldwide network to track incoming meteoroids. *A&A* 644, A53.
- Devillepoix, H. A. R., Bland, P. A., Sansom, E. K., et al. (2019). Observation of metre-scale impactors by the Desert Fireball Network. *MNRAS* 483, 5166.
- Global Meteor Network. Shower-association reference table distributed with the RMS/meteorTools software (gmn_shower_table). <https://github.com/CroatianMeteorNetwork/RMS>
- Guzik, P., et al. (2020). Initial characterization of interstellar comet 2I/Borisov. *Nature Astronomy* 4, 53.
- Hajduková, M. & Kornoš, L. (2020). The influence of meteor measurement errors on the heliocentric orbits of meteoroids. *Planetary and Space Science* 190, 104965.
- Hajduková, M., Stober, G., Barghini, D., Kotten, P., Vaubaillon, J., Sterken, V. J., Ďurišová, S., Jackson, A. & Desch, S. (2024). No evidence for interstellar fireballs in the CNEOS database. *A&A* 691, A8. doi:10.1051/0004-6361/202449569.
- Hughes, A., et al. (2022). Analysis of the April 13, 2021 bolide off the coast of Florida and Grand Bahama Island. *Meteoritics & Planetary Science* 57, 575.
- IAU Meteor Data Center (2026). Meteor databases, version 2026: CAMSv3, SAAMER and AMOR annual archives. <https://ceres.ta3.sk/iaumdcdb/>

- IAU Meteor Data Center. Meteor shower database: established, working and full stream lists (streamfulldata.txt). <https://www.ta3.sk/IAUC22DB/MDC2007/> ; Jopek & Kanuchova (2017), P&SS 143, 3; Jenniskens et al. (2020), P&SS 182, 104821.
- Jenniskens, P., et al. (2009). The impact and recovery of asteroid 2008 TC3. *Nature* 458, 485.
- Meech, K. J., et al. (2017). A brief visit from a red and extremely elongated interstellar asteroid. *Nature* 552, 378.
- NASA/JPL Center for Near Earth Object Studies. Fireball and bolide data; SSD Fireball API. <https://cneos.jpl.nasa.gov/fireballs/>
- NASA/JPL Solar System Dynamics. Horizons system. <https://ssd.jpl.nasa.gov/horizons/>
- Ozerov, A., Smith, J. C., Dotson, J. L. & Longenbaugh, R. S. (2024). GOES GLM, biased bolides, and debiased distributions. *Icarus* 408, 115843. doi:10.1016/j.icarus.2023.115843.
- Peña-Asensio, E. & Seligman, D. Z. (2025). The interstellar flux gap: From dust to kilometre-scale objects. *Astronomy & Astrophysics* 704, L1. doi:10.1051/0004-6361/202557337. arXiv:2511.01957.
- Peña-Asensio, E., Socas-Navarro, H. & Seligman, D. Z. (2025). Error dependencies in the space-based CNEOS fireball database. *A&A* 701, A202. arXiv:2508.01454.
- Peña-Asensio, E., Visuri, J., Trigo-Rodríguez, J. M., Socas-Navarro, H., Gritsevich, M., Siljama, M. & Rimola, A. (2024b). Oort cloud perturbations as a source of hyperbolic Earth impactors. *Icarus* 408, 115844. doi:10.1016/j.icarus.2023.115844.
- Peña-Asensio, E., et al. (2022). Orbital characterization of superbolides observed from space. *AJ* 164, 76.
- Ronac Giannone, M. & Silber, E. A. (2026). The role of source geometry and atmospheric propagation in global bolide infrasound detectability. *Icarus* 458, 117194.
- Seligman, D. Z., Micheli, M., Farnocchia, D., et al. (2025). Discovery and preliminary characterization of a third interstellar object: 3I/ATLAS. *ApJL* 989, L36.
- Shober, P. M., et al. (2026). Comparing the data-reduction pipelines of FRIPON, DFN, WMPL and AMOS: case study of the Geminids. *A&A* 705, A65.
- Silber, E. A. & Sawal, V. (2025). BLADE: a bolide light-curve catalogue. *AJ* 170, 153.
- Silber, E. A. (2024). Perspectives and challenges in bolide infrasound processing and interpretation: a focused review with case studies. *Remote Sensing* 16, 3628.

- Silber, E. A. (2025). Investigating the relationship between bolide entry angle and apparent direction of infrasound signal arrivals. *Pure and Applied Geophysics* 182, 2373–2392. doi:10.1007/s00024-025-03706-1.
- Silber, E. A., et al. (2025). Multiparameter constraints on empirical infrasound period–yield relations for bolides and implications for planetary defense. *AJ* 170, 38. Event-level database: Harvard Dataverse, doi:10.7910/DVN/P6XZ67.
- Silber, E. A., Brown, E., Thompson, A. R. & Sawal, V. (2025). A curated dataset of regional meteor events with simultaneous optical and infrasound observations (2006–2011). *Data* 10, 138. doi:10.3390/data10090138. Dataset: Zenodo doi:10.5281/zenodo.15868512; user guide SAND2025-10874, doi:10.2172/2587514.
- Silber, E. A., Rubbrecht, N., Tillier, C. E., Cano, J. L., Landis, R. R. & Longenbaugh, R. (2026). Multi-sensor observations of a high-velocity fireball over the South Atlantic on 2026 April 1. *Research Notes of the AAS*, 10, 183. doi:10.3847/2515-5172/ae853e.
- Silber, E. A., et al. (2026). The 2023 May 20 Queensland superbolide: multi-modal observations from atmospheric entry to postdisruption dust cloud. *AJ* 172, 178; Zenodo 20739713.
- Smith, J. C., Morris, R. L., Rumpf, C., Longenbaugh, R., et al. (2021). An automated bolide detection pipeline for GOES GLM. *Icarus* 368, 114576.
- SonotaCo, et al. (2021). SonotaCo network simultaneously observed meteor data sets. *WGN* 49, 64.
- Vida, D., et al. (2020, 2021). The Global Meteor Network. *MNRAS* 491, 2688; *MNRAS* 506, 5046.
- Vierinen, J. (2025, 2026). Preliminary meteor head-echo dataset (Zenodo 17139689); PANSY meteor head-echo orbit catalogue v1 (Zenodo 21703650).
- Western Meteor Physics Group (2026). Combined CAMO/EMCCD low-velocity Earth-impacting meteoroids. Zenodo 19615720.