



Temporal Equivalence Principle and the Earth Flyby Anomaly: Clock-Sector Response and Coherent-Link Cancellation

Matthew Lukin Smawfield

Version: v0.4 (Yogyakarta)

First published: 17 May 2026 · Last updated: 17 October 2026

DOI: 10.5281/zenodo.19454862

Code Availability: github.com/matthewsmawfield/TEP-EFA

Abstract

Twelve Earth gravity assist flybys spanning nine spacecraft are analyzed within the Temporal Equivalence Principle (TEP) framework. TEP posits that proper time is governed by a dynamical scalar field ϕ ; all non-gravitational matter couples universally to a causal matter metric through the conformal factor $A(\phi) = \exp(\beta_A \phi/M_P)$ with $\beta_A = -1$. Under Paper 0's amplitude-shear split (v0.16), the Earth-vicinity Temporal Shear is strongly screened by the two-body kinetic operator, so the flyby anomaly cannot be a shear-force effect. The candidate mechanism is instead the clock sector: highly asymmetric, fast hyperbolic flybys accumulate rapid proper-time offsets at perigee through the unsuppressed amplitude channel S_A , which orbit-determination codes — reconstructing velocities under an assumed GR time standard — misread as ΔV discontinuities wherever the tracking link references the spacecraft's own time standard. The channel analysis performed here shows the term survives only in clock-carrying observables (one-way, non-coherent, or regenerative-ranging links and onboard-clock telemetry); in the coherent-transponder class assigned to the published catalogue from the literature, the transponder endpoint cancels it identically, so on that classification the catalogue anomalies are not produced by the clock sector in the measured observable (encounter-specific tracking modes remain to be checked) — the mechanism's claim is instead a falsifiable clock-carrying prediction whose m/s-scale raw amplitude is conditional on the adopted profile and the eccentric-orbit clock constraint in Section 5.7a. The fitted response profile $S_\oplus(r)$ is accordingly reinterpreted as the clock-sector amplitude calibration, with the characteristic terrestrial field scale $R_{\text{sol}} = R_T(M_\oplus) = (3M_\oplus/4\pi\rho_T)^{1/3} \approx 4146 \text{ km}$ for $\rho_T \approx 20 \text{ g/cm}^3$. Since the UCD calibration of ρ_T descends from the GNSS coherence-length measurement, this terrestrial length scale is a single shared input rather than an independent validation.

The accumulated proper-time offset manifests as an apparent velocity discontinuity in GR-based orbit reconstruction — a clock-sector artefact misread as ΔV without an unmodeled force — in the measurement classes that carry the spacecraft's own time standard. The non-radial component of the offset integral, modulated by Earth's oblateness (J_2, J_3, J_4), trajectory asymmetry, the velocity-dependent response template, and the perigee plasma template layer, supplies the geometric template against which the observed flyby anomaly pattern is organized. Six published nonzero entries are retained in the re-transcribed catalogue — four firm detections (NEAR, Galileo 1990, Rosetta 2005, Galileo 1992) and two marginal $\sim 2\sigma$ entries (Cassini, MESSENGER) — alongside three published nulls/bounds and three flybys without public anomaly reports; the re-transcription

against the Anderson et al. (2008) primary source corrected several sign, uncertainty, and declination entries (Section 2.3). Under a full-catalog likelihood with a fractional misspecification floor — a per-flyby systematic set to the measured between-flyby amplitude heterogeneity ($CV_\beta \approx 0.47$) times each model's own predicted signal, iterated to a fixed point and invariant under the arbitrary field-normalization anchor — the archived information-criterion comparison describes how the restricted geometric template organizes the catalogue relative to the Null by $\Delta\text{BIC} \approx 1.67 \times 10^4$; the Anderson empirical baseline, whose two parameters are fitted on the same catalogue, attains a nominally smaller BIC but at parity-level separation ($\Delta\text{BIC} \approx 4.8$), an ordering that reverses in TEP restricted's favour ($\Delta\text{BIC} \approx +2767$) when the published uncertainties are taken at face value without a systematic floor. These scores describe empirical geometry-template organization, not evidence for the clock-sector mechanism, which is absent from the coherent catalogue observable. They retain the archived iteratively weighted fitting convention; a fully optimized prediction-dependent-variance likelihood and its nuisance accounting are required before interpreting them as calibrated model-selection evidence.

A structural consistency check is provided by the ordering of anomaly magnitudes against the trajectory-asymmetry factor ($\cos\delta_{\text{in}} - \cos\delta_{\text{out}}$). In the gated $n = 6$ detection ensemble the ordering is positive but not decisive (Spearman $\rho = +0.66$, $p = 0.156$): NEAR carries both the largest asymmetry factor and the largest published anomaly in every ranked subset. Across the broader catalogue the rank association is moderate: Spearman $\rho = 0.56$ ($p = 0.11$) over the $n = 9$ flybys with published asymmetries, and $\rho = 0.68$ ($p = 0.090$) over the $n = 7$ flybys with nonzero published anomalies, computed in the pipeline's ordering diagnostic. These are reported as a qualitative consistency check rather than a primary discriminator. The inverse-variance weighted mean $\beta_{\text{fit}} \approx 4.74 \times 10^{-3}$ (formal SE 5.83×10^{-4} ; random-effects SE 1.28×10^{-3} ; between-flyby $\tau \approx 2.23 \times 10^{-3}$) quantifies the cross-flyby effective response scale; it is not the bare conformal coupling β_A . The amplitude-informative per-flyby fits span 4.43×10^{-3} to 1.44×10^{-2} (with the negative-branch Galileo 1992 member at 6.7×10^{-2}), consistent with geometry-dependent modulation. The fitted amplitudes are trajectory-response coefficients, not the Cassini source charge; Cassini PPN compliance is evaluated through the solar source-charge projection $\alpha_{\text{eff}} = S_\Sigma \alpha_0$ (Section 4.6.1a). The re-transcribed catalogue resolves the historical Cassini sign question: the primary-source anomaly is negative (-2 ± 1 mm/s), so the kernel's predicted sign agrees and the residual tension is magnitude-only. The geometry envelope carries seven nominal deterministic coefficients plus one fitted amplitude against six retained nonzero entries, so the model-selection comparisons are read as organisational support under the adopted convention rather than calibrated evidence.

This work shows that a restricted trajectory-geometry template quantitatively organizes the published Earth flyby anomaly catalogue under the adopted Anderson trajectory-geometry convention, while remaining consistent with precision solar system constraints and identifying specific stress tests—principally the Galileo 1992 amplitude shortfall and the Juno raw-tension null, together with raw DSN reanalysis—that are required to advance from model organisation to key confirmation. The catalogue's evidential status is conditional on one provenance question carried openly here: the orbit-determination filtering interpretation invoked for the published nulls is symmetric with the alternative that the early anomalies were artefacts of less sophisticated orbit-determination pipelines — the two readings are equally consistent with the present data, and only an independent minimal-OD reanalysis of anomaly and null flybys under identical force models can arbitrate between them.

The terrestrial response profile $S_\oplus(\mathbf{r})$ constitutes the Earth-vicinity geometric realization of the clock-sector amplitude response $S_A(\mathcal{E})$. Rather than a hard boundary, this radial profile traces the continuous amplitude structure of the temporal field within Earth's local potential well.

Keywords: Earth flyby anomaly, Temporal Equivalence Principle, clock-sector artefact, proper-time offset, trajectory asymmetry, geometric screening, Temporal Topology, Temporal Shear

1. Introduction

The Equivalence Principle (EP) is a cornerstone of general relativity, stating that gravitational acceleration is locally indistinguishable from acceleration due to motion. However, the Temporal Equivalence Principle (TEP)—the assertion that proper time is governed by a dynamical scalar field whose conformal gradients and disformal velocity-dependent terms can produce screened trajectory-level residuals—posits that the rate of time is a dynamical scalar field ϕ . This framework, established in Paper 0, proposes that all non-gravitational matter couples universally to a causal matter metric $\tilde{g}_{\mu\nu} = A^2(\phi)g_{\mu\nu} + B(\phi)\nabla_\mu\phi\nabla_\nu\phi$, where $A(\phi) = \exp(\beta_A\phi/M_{\text{Pl}})$. In the flyby regime the disformal term is subdominant and the conformal sector dominates the clock-sector response.

Key Terminology

- *Proper time* (τ) is the time measured by a clock following a specific trajectory through the causal metric.
- *Temporal Topology* refers to the spatial structure of the field ϕ , whose gradient response is pinned where the ambient time-gradient is steep — weak-field language: high-acceleration configurations — by the nonlinear kinetic operator; density is an environmental diagnostic, not an independent threshold.

- *Temporal Shear* ($\Sigma_\mu = \nabla_\mu \ln A = (\beta_A/M_{Pl})\nabla_\mu\phi$) is the gradient of the conformal factor, which generates the observed geometric acceleration.
- *Clock-sector artefact* describes the apparent velocity discontinuity that orbit-determination codes produce when accumulated proper-time offsets — sampled through the unsuppressed amplitude channel S_A at perigee — are present in a clock-carrying tracking link (one-way, non-coherent, or regenerative-ranging observables) and are fitted under an assumed GR time standard. In the fully coherent two-way class the transponder endpoint cancels the term identically (Section 5.2).
- *PPN parameter* γ measures the amount of spatial curvature; in Paper 0 (Sec. 7, Step 80), $\gamma_{PPN} - 1 = -\alpha_0 \alpha_{\text{eff}} / (1 + \alpha_0 \alpha_{\text{eff}}) = -2\beta_A^2 S_\Sigma / (1 + 2\beta_A^2 S_\Sigma)$ under the canonical unified unstretched-space metric (or $-2\alpha_0 \alpha_{\text{eff}} / (1 + \alpha_0 \alpha_{\text{eff}}) = -4\beta_A^2 S_\Sigma / (1 + 2\beta_A^2 S_\Sigma)$ on the isotropic-conformal benchmark) with $\alpha_{\text{eff}} = S_\Sigma \alpha_0$ the screened source-charge coupling — linear in S_Σ because the photon probe is unscreened; for $A = \exp(\beta_A \phi / M_{Pl})$ this maps to $|\gamma - 1| \approx 2\beta_A^2 S_\Sigma^{(\odot)}$ under the unified unstretched-space metric ($4\beta_A^2 S_\Sigma^{(\odot)}$ on the isotropic-conformal benchmark) when comparing to Cassini's bound on $|\gamma - 1|$.

1.1 The Earth Flyby Anomaly

Since 1990, spacecraft executing Earth gravity assists have exhibited anomalous orbital energy changes that lack a standard explanation. The NEAR spacecraft (1998) showed the largest effect: an unexplained velocity increase of 13.46 mm/s. Galileo (1990, 1992) and Cassini (1999) displayed smaller but significant anomalies. These velocity shifts occur precisely at perigee passage and persist as asymptotic excess velocities (v_∞) in the outbound trajectories.

The published residuals motivate tests of both physical signals and tracking/reconstruction systematics. The physical checks in Section 5.2 bound several conventional contributions but do not exclude every systematic explanation. TEP supplies a distinct candidate clock-sector artefact through trajectory-dependent proper-time transport in links carrying an onboard time reference; Section 5.2 establishes that this term does not account for the published coherent-link catalogue.

1.2 TEP as a Candidate Explanation

The TEP framework provides a testable clock-sector response for flyby observables that carry the spacecraft's own time reference. Under Paper 0's amplitude-shear split (v0.16), the Earth-vicinity Temporal Shear is strongly screened by the two-body kinetic operator, so the anomaly cannot be a shear-force effect. Instead, highly asymmetric hyperbolic flybys accumulate rapid proper-time offsets at perigee through the unsuppressed amplitude channel; orbit-determination codes, reconstructing velocities under an assumed GR time standard, misread these offsets as ΔV discontinuities wherever the tracking link carries the spacecraft's own time reference — the Clock-Sector Reduction Artefact of Paper 0, whose observable reach the channel analysis of Section 5.2 fixes exactly.

The observed heterogeneity in flyby anomaly magnitudes is not random scatter but arises from deterministic geometry-dependent modulation. The TEP prediction for a given flyby depends on several physical factors: (1) perigee altitude (determines the clock-sector field excursion via the density-dependent amplitude profile), (2) approach-departure asymmetry (disformal coupling requires velocity-dependent anti-aligned geometry), (3) the perigee plasma template layer (a bounded heuristic envelope coefficient; the stronger Debye attenuation ansatz is quarantined at $S_{\text{plasma}} = 1$, Section 3.10), (4) solar activity (modulates ionospheric density within that template layer), and (5) cosmographic CMB-frame velocity geometry (the disformal coupling scales as v^2 in the ambient-field rest frame — the frame in which the spatial gradient of the cosmological temporal field vanishes — approximated operationally by the CMB dipole frame)—a deeper prediction tested only in an exploratory capacity (Section 4.11). These factors combine to produce a wide span in per-flyby fitted β_{fit} across the Step 008 S/N-qualified ensemble; the legacy inverse-variance diagnostic restricted to sign agreement at $\beta_{A,\text{ref}}$ (NEAR, Galileo 1990, Rosetta 2005) is reported separately as `beta_statistics_sign_gated_diagnostic` in `results/step008_fitting_results.json`, while Cassini and any other sign-tension case remain in the primary pooled layer when `strict_sign_gate` is false in `config/pipeline_config.json`. The Temporal Topology screening mechanism is essential for three reasons: (1) it ensures the flyby amplitude and environmental response remain consistent with solar system constraints; (2) it explains both detections and null results through environment-dependent screening; and (3) it establishes the transition radius $R_{\text{sol}} \approx 4146$ km as a universal scale. Flybys sampling regions of large field excursion (low altitude, high asymmetry) exhibit anomalies, while those in shielded regimes (high altitude or symmetric trajectories) remain null.

1.3 This Work

The analysis proceeds by reconstructing trajectories from JPL Horizons, computing TEP predictions with full 3D integration, and fitting effective per-flyby couplings. Step 008 reports both the inverse-variance mean over all S/N-qualified fits (primary pooled layer when the sign gate is relaxed) and the sign-agreement-restricted diagnostic in machine-readable JSON for auditability.

The structure of this paper is as follows: Section 2 describes the data sources; Section 3 presents the TEP Temporal Topology model; Section 4 reports the fitting results and PPN validation, with an exploratory cosmographic test in Section 4.11; Section 5 discusses the clock-sector artefact interpretation; and Section 6 concludes with prospects for further tests.

2. Observations and Data

2.1 The Flyby Spacecraft Sample

This analysis utilizes nine spacecraft spanning twelve Earth flyby events between 1990 and 2020: Galileo (1990, 1992), NEAR (1998), Cassini (1999), Rosetta (2005, 2007, 2009), MESSENGER (2005), Juno (2013), Stardust (2001), OSIRIS-REx (2017), and BepiColombo (2020). The dataset is divided into three data quality classes: *published anomalies* (six flybys with measured nonzero Δv and formal uncertainties — the four positive detections plus the two negative or marginal entries of Anderson et al. 2008 Table I), *published nulls/bounds* (three flybys with explicitly reported null results or upper limits), and *no public anomaly report* (three flybys with no published search or measurement). The latter class is not used in quantitative likelihood. Table 1 summarizes the key parameters for each flyby.

Physical Constants

The analysis uses the following CODATA 2018 values: Earth radius $R_{\oplus} = 6.371 \times 10^6$ m, gravitational constant $G = 6.67430 \times 10^{-11}$ m³ kg⁻¹ s⁻², and speed of light $c = 299\,792\,458$ m/s (exact). The reduced Planck mass $M_{\text{Pl}} = 2.435 \times 10^{18}$ GeV is derived from $\hbar c/G^{1/2}$.

Table 1: Earth Flyby Spacecraft Parameters

Spacecraft	Date	Perigee (km)	v_{∞} (km/s)	Δv_{obs} (mm/s)	σ (mm/s)	Data class
Galileo	1990-12-08	960	13.73	3.92	0.08	Published anomaly
Galileo	1992-12-08	303	14.08	−4.60	1.00	Published anomaly (negative)
NEAR	1998-01-23	539	12.72	13.46	0.13	Published anomaly
Cassini	1999-08-18	1175	19.02	−2.00	1.00	Published anomaly (negative)
Rosetta	2005-03-04	1955	10.51	1.80	0.03	Published anomaly
Rosetta	2007-11-13	5301	12.46	0.02	0.05	Published null/bound
Rosetta	2009-11-13	2572	13.31	0.00	0.05	Published null/bound
MESSENGER	2005-08-02	2347	10.39	0.02	0.01	Published anomaly (marginal)
Juno	2013-10-09	817	14.79	0.00	0.02	Published null/bound
Stardust	2001-01-15	6009	10.31	—	—	No public anomaly report
OSIRIS-REx	2017-09-22	17239	8.52	—	—	No public anomaly report
BepiColombo	2020-04-10	12697	7.59	—	—	No public anomaly report

Note: Δv_{obs} values for the published anomaly and published null/bound classes are from Anderson et al. (2008) and companion papers. Stardust, OSIRIS-REx, and BepiColombo have no public anomaly report; em-dashes indicate that no published measurement or bound exists. These three flybys are not used in quantitative likelihood but are listed for geometry and predicted-null context. Rosetta 2009 is a published null (Müller et al. 2010); the archival catalogue assigns $\sigma = 0.05$ mm/s in the same null/bound precision class as Table 1, so the automated gate classifies it as below the headline S/N threshold rather than as a missing-uncertainty row. Perigee values are altitudes above Earth's surface; v_{∞} is the hyperbolic excess velocity. The Galileo 1992, Cassini, and MESSENGER entries follow the Anderson et al. (2008) Table I primary-source values (−4.60 ± 1.00, −2 ± 1, and +0.02 ± 0.01 mm/s respectively), confirmed against the mission declination reconstruction in Table 1a.

2.2 Data Sources and Provenance

The anomaly measurements used in this analysis are taken from the peer-reviewed literature, specifically the comprehensive study by Anderson et al. (2008) and subsequent mission-specific analyses. These values were obtained through NASA's Deep Space Network

(DSN) Doppler tracking combined with the Jet Propulsion Laboratory Orbit Determination Program (ODP).

Literature sources:

- Primary reference: Anderson, J. D., et al. (2008). "Anomalous Orbital-Energy Changes Observed during Spacecraft Flybys of Earth." *Physical Review Letters*, 100(9), 091102.
- Rosetta analysis: Morley, T., & Budnik, F. (2007). "Rosetta Navigation at its First Earth-Swingby." *Proceedings of the 20th International Symposium on Space Flight Dynamics*.
- Juno analysis: Aksenov, E. L., & Tuchin, A. G. (2020). "Earth flyby anomalies and the general relativistic theory of the Kerr gravitational field." *MNRAS*, 492(3), 3703-3711.

2.3 Data Quality Assessment

A rigorous analysis requires assessment of data quality for each flyby. All six published detections have complete DSN coverage spanning ± 12 hours around perigee, enabling robust pre/post comparison (catalog-level statement; the inverse-variance β_{fit} ensemble in Steps 008 and 026 uses $n = 6$ S/N-qualified primary fits when `strict_sign_gate` is false; with the corrected primary-source anomalies, all six detections agree in sign with the TEP prediction at the reference coupling, so the sign-agreement diagnostic coincides with the gated set). The reported uncertainties (0.01–1.00 mm/s) span the DSN Doppler precision range for the different missions and tracking configurations.

Ensemble gates (Steps 008 and 026) are applied mechanically: a flyby enters the restricted likelihood only if it carries a published $(\Delta v, \sigma)$ pair, satisfies $|\Delta v|/\sigma > 2$, and shows sign agreement between Δv_{obs} and the TEP prediction at the pre-specified reference coupling. Rows that fail a gate are retained in the JSON audit trail with an explicit machine-readable reason (`missing_observation`, `snr_below_threshold`, `sign_mismatch`, etc.) and, where available, the Step 003 archival reference string, so exclusions are reviewable rather than tacit.

Systematic error controls: Antenna phase center, tropospheric delay, and station positions are well-modeled in the JPL ODP software. Residual uncertainties are at the ~ 0.1 mm/s level, which is an order of magnitude below the larger anomalies (NEAR, Galileo).

2.4 Trajectory Data from JPL Horizons

Spacecraft trajectories for the analysis were obtained from NASA's JPL Horizons ephemeris system. For each flyby, state vectors (position and velocity) spanning ± 2 days around perigee passage are reconstructed. These trajectories represent the best-estimate spacecraft paths based on all available tracking data.

2.4a Geometry Factor Audit: Published Declinations vs Horizons

The trajectory asymmetry factor used in both the Anderson et al. (2008) empirical formula and the TEP model is $\cos \delta_{\text{in}} - \cos \delta_{\text{out}}$. Table 1a audits the provenance of this factor by comparing three independent estimates for each flyby with published declinations: (i) the Anderson et al. (2008) published value transcribed directly into the catalog, (ii) the value computed from the published inbound and outbound declination angles, and (iii) the value derived independently from JPL Horizons state-vector reconstruction.

Table 1a: Trajectory Asymmetry Factor Audit

Flyby	δ_{in} (deg)	δ_{out} (deg)	$\cos \delta_{\text{in}} - \cos \delta_{\text{out}}$ (published)	$\cos \delta_{\text{in}} - \cos \delta_{\text{out}}$ (computed from decs)	$\cos \delta_{\text{in}} - \cos \delta_{\text{out}}$ (Horizons)
NEAR	−20.76	−71.96	0.625	0.625	0.625
Galileo 1990	−12.52	−34.15	0.149	0.149	0.149
Cassini	−12.92	−4.99	−0.022	−0.022	−0.022
Galileo 1992	−34.26	−4.87	−0.170	−0.170	−0.170
Rosetta 2005	−2.81	−34.29	0.173	0.173	0.173
Rosetta 2007	−10.80	+18.60	0.034	0.034	0.034
MESSENGER 2005	+31.44	−31.92	0.005	0.004	0.005
Juno	−26.9	+34.6	—	0.069	0.069

Key finding: With the re-transcribed Anderson et al. (2008) Table I declinations, the computed $\cos \delta_{\text{in}} - \cos \delta_{\text{out}}$ agrees with the published asymmetry factor for every row, and both agree with the independent JPL Horizons state-vector reconstruction. An earlier

version of this catalogue carried declination transcription errors (most consequentially NEAR's outbound leg, entered as -19.7° rather than -71.96° , and both Galileo 1992 legs); those errors manufactured an apparent convention mismatch between the tabulated declinations and the published asymmetry factors that does not exist in the primary source. The three-way agreement now shown — published asymmetry, recomputed from the published declinations, and Horizons-derived — removes the declination-provenance question entirely and validates the adopted geometry factors used by both the Anderson empirical model and the TEP kernel. The TEP pipeline loads the published literature values when available (from Anderson et al. 2008 Table I) and computes from Horizons-derived declinations otherwise.

Provenance of the corrected values. The transcription errors were identified by auditing the catalogue against the Anderson et al. (2008) primary source and its companion trajectory paper (Anderson, Campbell & Nieto 2006), cross-checked against independent reproductions. Corrections were confined to the published tabulated quantities: NEAR's outbound declination (-71.96° , not -19.7°) and uncertainty (± 0.13 mm/s, not ± 0.01), Galileo 1992's two declinations ($-34.26^\circ / -4.87^\circ$, not $-30.3^\circ / -33.8^\circ$) and its anomaly (-4.60 ± 1.00 mm/s, previously entered as a null), Cassini's anomaly (-2 ± 1 mm/s, previously entered as $+0.11 \pm 0.05$), MESSENGER's anomaly ($+0.02 \pm 0.01$ mm/s, previously a null), Galileo 1990's inbound declination (-12.52° , not -25.1°) and uncertainty (± 0.08 mm/s, not ± 0.03), and Rosetta 2005's anomaly ($+1.80 \pm 0.03$ mm/s) and inbound declination (-2.81° , not -3.4°). Perigee altitudes were likewise aligned to the primary-source tabulation (539, 960, 303, 1175, 1955, 2347 km). All downstream results in this paper use the corrected catalogue; the corrected geometry factors are corroborated independently by the Horizons column of Table 1a.

3. Methodology

The analysis employs a four-step pipeline to compare the published flyby anomaly geometry with the TEP response template and to test the clock-sector artefact separately in the link classes where it survives. The pipeline retrieves spacecraft trajectories from JPL Horizons, computes TEP predictions for each flyby geometry using full 3D integration, fits the coupling parameter β_{fit} to match observed anomalies, and validates all parameters against solar system PPN constraints. Which measurement classes can carry the clock-sector term is itself part of the analysis: the channel accounting of Section 5.2 shows the spacecraft's conformal excursion survives only in clock-carrying links (one-way, non-coherent, or regenerative-ranging observables), while the published catalogue is assigned to the coherent-transponder class on the literature-based ledger of Section 5.2, pending encounter-specific tracking verification.

3.1 Data Acquisition

Spacecraft trajectories are obtained from the NASA JPL Horizons ephemeris system via direct HTTP query to the batch CGI endpoint. For each flyby, observational quantities (right ascension, declination, range, and range-rate) are retrieved in the ICRF (International Celestial Reference Frame) at 1-minute intervals spanning ± 2 days around perigee passage; three-dimensional state vectors are reconstructed from these observables in a subsequent processing step.

Trajectory Parameters Extracted

- Perigee altitude (minimum geocentric distance)
- Perigee velocity (speed at closest approach)
- Inbound/outbound asymptotic velocity (v_∞)
- Spacecraft potential at perigee (Φ_{sc})

Flyby velocity anomalies (Δv_{obs}) are taken from published literature. The primary source is Anderson et al. (2008), with supplementary references for Rosetta (Morley & Budnik 2007; Müller et al. 2008, 2010) and Juno (Aksenov & Tuchin 2020). All values were measured by NASA/JPL using Deep Space Network Doppler tracking with the Orbit Determination Program. Asymptotic v_∞ declinations (δ_{in} , δ_{out}) tabulated in Anderson et al. (2008) are transcribed directly from that source; for additional catalogued flybys without published declinations, the same eccentricity-vector reconstruction applied to archived JPL Horizons state arcs (Step 040a / Step 007, geocentric perigee anchored to the archival flyby date) supplies δ_{in} and δ_{out} so that geometry is traceable to Horizons rather than hand-entered asymmetries.

3.2 TEP Temporal Topology Model

The TEP framework provides a quantitative model for the flyby anomaly through the perigee proper-time offset arising from the temporal-field amplitude excursion. With the universal conformal coupling $A(\phi) = \exp(\beta_A \phi / M_{\text{Pl}})$, the accumulated offset along the trajectory projects into an apparent velocity discontinuity under GR-based orbit determination:

$$\mathbf{F}_\phi = \beta_{A,\text{eff}} \frac{c^2 \nabla \phi}{M_{\text{Pl}}} \quad (1)$$

Under the amplitude–shear split (Paper 0 §7), the Earth-vicinity Temporal Shear is screened by the two-body operator, so this kernel is not a physical force on the spacecraft — the shear-force channel is closed. It is the orbit-determination-equivalent kernel of the clock-sector artefact in a clock-carrying link: the spacecraft's matter-metric rate offset $\eta(t) = \exp(\beta_A \psi / M_{\text{Pl}}) - 1$ carries the same radial gradient $\beta_A \nabla \phi / M_{\text{Pl}}$, and a tracking filter that assumes the GR time standard reconstructs the accumulated rate-offset excursion exactly as if it were this impulse. The two signatures are observationally indistinguishable over a flyby arc *where the term is present in the observable* — one-way, non-coherent, or regenerative-ranging links; in the coherent two-way class of the published catalogue the transponder endpoint cancels it identically (Section 5.2). The mechanism-faithful forward model is the Step 043 OD rederivation (Table 3h), which computes the artefact's signature on the clock-carrying link class over the real arcs. Here $\beta_{A,\text{eff}} = \beta_A \times S_\oplus(r)$ is the effective artefact-response amplitude, where $S_\oplus(r)$ is the continuous clock-amplitude response profile. The characteristic surface ratio $S_\oplus \approx 0.35$ is fixed by the UCD saturation geometry as $S_\oplus = (R_\oplus - R_{\text{sol}}) / R_\oplus$ with $R_{\text{sol}} \approx 4146$ km (distinct from the embedding factor $R_{\text{sol}} / R_\oplus$ used in Paper 6). The radial component of the apparent impulse is indistinguishable from a small shift in GM and is absorbed by orbit determination. The non-radial component—modulated by Earth's oblateness (J_2, J_3, J_4) and the spacecraft's trajectory geometry—produces the apparent net velocity change that reads as the flyby anomaly.

The predicted velocity discontinuity is resolved through rigorous numerical integration of the artefact kernel along the trajectory in the Earth-centered inertial (ECI) frame — the accumulated apparent impulse that an orbit-determination filter would attribute to dynamics, not a physical trajectory deflection. This approach captures the accumulated proper-time offset as the spacecraft traverses the varying field amplitude, incorporating a 4th-order Spherical Harmonic Expansion (SHEX) for the geopotential to ensure that local gravitational perturbations are not conflated with the clock-sector response:

$$\Delta \mathbf{v}_{\text{TEP}} \approx \left. \frac{\beta_{A,\text{eff}} c^2 \nabla \phi}{M_{\text{Pl}}} \right|_{\text{peri}} \Delta t_{\text{peri}} + \frac{b_{\text{disf}}}{M_{\text{Pl}}} (\nabla \phi \cdot \mathbf{v}) \mathbf{v} \Big|_{\text{peri}} \Delta t_{\text{peri}} \quad (2)$$

where the simplified perigee approximation is:

$$\Delta v_{\text{TEP}} \approx \beta_{A,\text{eff}} \frac{c^2}{M_{\text{Pl}}} \left(\frac{d\phi}{dr} \right)_{r_p} \frac{r_p}{v_p} [J_2 + J_3 \sin(\lambda_p) + J_4 P_4(\sin \lambda_p)] \left(\frac{R_\oplus}{r_p} \right)^2 (\cos \delta_{\text{in}} - \cos \delta_{\text{out}}) \quad (3)$$

where:

- $(d\phi/dr)_{r_p}$ is the scalar field gradient at perigee altitude
- r_p and v_p are the perigee distance and velocity
- J_2, J_3, J_4 are the zonal harmonics (EGM96/WGS84 coefficients)
- λ_p is the perigee latitude
- δ_{in} and δ_{out} are the asymptotic declinations on approach and departure (from Anderson et al. (2008))

3.2b Full 3D Trajectory Integration

The perigee approximation provides a computationally efficient closed-form estimate, but the primary predictions are validated against full 3D trajectory integration to capture the accumulated proper-time offset along the spacecraft path. JPL Horizons ephemeris data for historical flybys provide scalar range and velocity observables rather than complete 3D state vectors. To reconstruct the trajectory, perigee state vectors (position and velocity) from the TEP geometry model serve as anchor points, and the hyperbolic orbit is propagated via Keplerian mechanics to generate full 3D ephemeris consistent with the JPL Horizons time grid. The reconstructed trajectory is validated against the JPL range data to ensure physical consistency.

The apparent velocity discontinuity is accumulated by integrating the TEP clock-sector response along the reconstructed path:

$$\Delta v_{\text{TEP}}^{(3D)} = \int_{t_{\text{in}}}^{t_{\text{out}}} \beta_{A,\text{eff}}(r(t)) \frac{c^2}{M_{\text{Pl}}} \left. \frac{d\phi}{dr} \right|_{r(t)} \mathcal{B}(\lambda(t), r(t)) \cos \delta_{\text{asym}} \mathcal{E}(t) s_{\text{dist}}(t) dt \quad (4)$$

where $\mathcal{B}(\lambda, r) = [J_2 + J_3 \sin(\lambda) + J_4 P_4(\sin \lambda)](R_\oplus/r)^2$ is the zonal harmonic bracket, $\mathcal{E}(t)$ is the geometry envelope factor, and $s_{\text{disf}}(t)$ is the disformal modulation factor. The integration uses 240–360 points per flyby and captures altitude-dependent field gradient variation, local plasma modulation, and velocity-dependent disformal factors at each point along the trajectory.

Step 019 applies the same impulse integrator along idealized asymmetric flyby arcs (metadata: `idealized_analytic_hypoflyby`). On comparable rows of Table 3b (Section 4.1.2), the ratio $\Delta v_{3D}/\Delta v_{\text{simplified}}$ differs from unity by tens of percent for the three gated primaries—a consistency check inside the analytic path class—while ratios against the Step 007 perigee impulse are deliberately not marketed as $\pm 10\%$ “agreement.” Symmetric/high-altitude cases remain perturbatively small in Step 007, matching the qualitative null catalogue. Ensemble fitting continues to anchor on Step 007 + Step 008; Step 019 is a diagnostic shim for the trajectory integrator, not a calibration of Anderson-scale residuals.

J3 contribution: The J3 term adds a latitude-dependent asymmetry to the non-radial component of the artefact kernel. However, J3 is two orders of magnitude smaller than J2 ($|J_3/J_2| \approx 2.3 \times 10^{-3}$), and its inclusion does not collapse the heterogeneity in fitted β_{fit} across the gated ensemble (factor ~ 3.3 across the tightly constrained members at fixed reference coupling, extending to ~ 15 with the negative-branch Galileo 1992 member, with formal Cochran $Q \gg 1$). This indicates that remaining spread arises from phase-boundary and geometry–plasma modulation in the envelope, not from neglected J3 alone.

The scalar field ϕ relaxes outside Earth with a profile scale $\lambda_T = 4200$ km — the operational value anchored on the measured GNSS clock-correlation length $L_c = 4201$ km ($\pm 15\%$ calibration band). Under the corpus's terrestrial transfer-sketch identification (Paper 6 §2) the exterior profile varies on $\mathcal{O}(R_T)$ with an undetermined order-unity prefactor, since a gradual screening transition produces no exterior structure sharper than the scale over which the integrated source charge changes over; the adopted value sits 1.3% above the ansatz target $R_T = 4146$ km, within the stated prefactor freedom and calibration band (the two are not independent: ρ_T was itself calibrated from L_c). This is a stated ansatz anchoring the profile scale — not a derived equality, and distinct from the in-medium Compton wavelength of a candidate chameleon completion.

The trajectory asymmetry factor $\cos \delta_{\text{in}} - \cos \delta_{\text{out}}$ is the dominant source of inter-flyby variation. Nearly symmetric trajectories (e.g., MESSENGER, $\cos \delta_{\text{in}} - \cos \delta_{\text{out}} \approx 0.005$) predict negligible anomalies, consistent with observations. Asymmetric trajectories (e.g., NEAR with $\cos \delta_{\text{in}} - \cos \delta_{\text{out}} = 0.625$) predict large anomalies. Cassini carries a small negative asymmetry factor (-0.022), verified by independent JPL Horizons reconstruction (Table 1a) — a physical property of the trajectory, not a convention artifact. The kernel therefore predicts a negative apparent anomaly, consistent in sign with the corrected -2 ± 1 mm/s catalogue value, though the predicted magnitude is negligible and the amplitude shortfall remains a genuine limitation of the phenomenological kernel for this geometry, not a convention mismatch. The mechanism-faithful calculation is consistent with this reading: the Step 043 OD rederivation, which propagates the proper-time artefact through the batch least-squares filter rather than the kernel, returns only a negligible endpoint remnant for Cassini ($\Delta V_{\text{app}} = +0.01$ mm/s, Table 3h) — the clock-sector channel does not produce the corrected -2 ± 1 mm/s catalogue value, consistently with that datum's coherent two-way link class, in which the term is absent at the observable (Section 5.2).

$$\phi(r) = \phi_{\text{earth}} + (\phi_{\text{space}} - \phi_{\text{earth}}) \left[1 - \exp\left(-\frac{r - R_\oplus}{\lambda_T}\right) \right] \quad (5)$$

Geometric screening: Critical to the flyby amplitude and environmental response is the transition radius $R_{\text{sol}} \approx 4146$ km from the UCD saturation model (Step 010), anchored on the same GNSS terrestrial calibration that fixes ρ_T — a shared-input relationship under the Paper 6 §2 identification (Section 3.11), not an independent validation. This defines the characteristic ratio $S_\oplus = (R_\oplus - R_{\text{sol}})/R_\oplus \approx 0.35$, the responsive-shell fraction of the radial profile through which a transported clock accumulates its differential excursion. Its relationship to the canonical clock-amplitude operator is an exact partition rather than a separate estimate: under the uniform-density-equivalent construction $R_{\text{sol}}^3 = 3M_\oplus/4\pi\rho_T$ and $R_\oplus^3 = 3M_\oplus/4\pi\bar{\rho}$, the canonical response evaluated at Earth's mean density reduces to $S_A(\bar{\rho}_\oplus) = (\bar{\rho}_\oplus/\rho_T)^{1/3} = R_{\text{sol}}/R_\oplus \approx 0.65$ — the embedded, saturated fraction of the equivalent radius, identical to the UCD embedding factor used in Paper 6 (UCD). The flyby clock channel projects the complementary fraction $S_\oplus = 1 - S_A(\bar{\rho}_\oplus)$: inside R_{sol} the field is pinned and carries no radial relaxation, so the excursion is generated precisely across the unsaturated shell. The two quantities are a partition sum rule of one radial profile, not two estimates of one operator (machine-checked in Step 010's `canonical_amplitude_partition` output: $S_A + S_\oplus = 1$ identically).

$$\beta_{A,\text{eff}} = \beta_A \times S_\oplus(r) \quad (6)$$

The pipeline's field-amplitude normalization is the canonical amplitude-sector excursion of Paper 0 v0.16. The pipeline carries the excursion complement $\phi(r)$ — zero at the pinned interior floor, rising to the ambient reference ϕ_{space} — so that the physical excursion $\psi(r) = \phi_{\text{space}} - \phi(r) \geq 0$ is the temporal-well depth sampled by a transported clock: positive inside the well and relaxing to zero at the ambient baseline, matching the corpus convention $\phi > 0$ in wells, $\phi \rightarrow 0$ ambient (Paper 0 Rule 4). The normalization is the unscreened linear response of the source computed in the Paper 0 radial-ODE closure (`TEP/results/step_01_radial_ode.json`, Earth entry):

$$\psi_{\text{uns}} = \frac{M_{\oplus}}{4\pi M_{\text{Pl}}^2 R_{\oplus}} = \frac{2GM_{\oplus}}{R_{\oplus}c^2} = 1.3926 \times 10^{-9} \quad (7)$$

This value is parameter-free — fixed by Earth's measured mass and radius — and independent of the fitted response amplitude β , so every pipeline prediction is strictly linear in β . It replaces the v0.1 density-minimum prescription $\phi_{\text{min}}(\rho) \propto \rho^{-1/(n+1)}$ (an inverse-power/chameleon-class realization): under $\beta_A = -1$ that branch admits no stable minimum (Paper 0 constraint F1) and runs the field largest in vacuum, opposite to the corpus convention. The Step 010 self-consistent-field diagnostic on that realization is retained solely as the documented contrast — it returns $S_{\oplus} \approx 1.0$, i.e. the excluded branch produces no terrestrial suppression shell.

Characteristic Field Values — Canonical Excursion (Paper 0 amplitude sector)

- Earth well depth (unscreened source response): $\Delta\phi = \psi_{\text{uns}} M_{\text{Pl}} = 3.39 \times 10^9 \text{ GeV}$
- Surface clock excursion: $\eta_{\oplus} = e^{-\psi_{\text{uns}}} - 1 = -1.3926 \times 10^{-9}$ (clock-carrying link class; Step 043). This is the amplitude-sector excursion, not the static lapse: the corpus's weak-field lapse ledger gives $\ln(A/A_{\infty}) = 2\beta_A^2 \Phi_N$ on the unscreened branch and the measured redshift-channel lapse $\tilde{N} \rightarrow e^{\Phi_N}$ under screening of the shear sector (Paper 0 §2.2, R9) — $\eta_{\oplus}$ is the transport excursion a carried clock accumulates, which static clock comparisons do not constrain
- Interior: excursion pinned at the well depth ($\phi_{\text{interior}} = 0$ in complement units)
- TEP exterior profile scale: $\lambda_T \approx 4200 \text{ km}$ (set equal to the geometric saturation scale $R_{\text{T}}(M_{\oplus})$ under the Paper 6 §2 transfer-sketch ansatz, up to an order-unity prefactor; the GNSS covariance scale $L_c = 4201 \text{ km}$ is the empirical anchor of that identification)
- Characteristic suppression: $S_{\oplus} \approx 0.35$ (UCD-derived from Step 010)

Fundamental coupling: The bare conformal coupling is $\beta_A = -1.0$, established in Paper 0 (Jakarta) §2.2. The fitted values reported here ($\beta_{\text{fit}} \sim 10^{-3}$) are the fitted response amplitudes; the effective clock-sector response $\beta_{A,\text{eff}} = S_{\oplus}(r) \beta_{\text{fit}}$ is derived from these under the excursion projection $S_{\oplus}(r)$ defined above. The sign of the artefact contribution in the flyby geometry is determined by trajectory asymmetry; the reported β_{fit} magnitudes should be compared to the absolute fundamental scale $|\beta_A| = 1.0$ — the ratio $\kappa_{\text{flyby}} = \beta_{\text{fit}}/|\beta_A| \approx 4.7 \times 10^{-3}$ is the population-level channel response coefficient, reported as an unexplained transfer-map quantity of the same open class as the Paper 0 Γ_X entries pending the action-level derivation. The fitted flyby amplitude is a trajectory-response coefficient; Solar-System PPN bounds apply to the screened source-charge projection $\alpha_{\text{eff}} = S_{\oplus} \alpha_0$, not to the fitted flyby amplitude itself. The converse direction closes identically: terrestrial bounds on the Earth environment constrain the shear channel, for which the corpus's own two-body operator returns an Earth-vicinity suppression of $\sim 10^{-21}$, whereas the claimed flyby observable is the unsuppressed clock-amplitude channel — and even there it is a transport excursion accumulated by a clock carried through the relaxing radial profile, not a stationary rate offset that inter-comparisons between clocks at rest in the terrestrial environment could bound. The same unsuppressed terrestrial amplitude response is the structure that the corpus's GNSS coherence-length measurement anchors ($L_c = 4201 \text{ km}$, Section 3.11), so the fitted excursion scale is consistent with — indeed required by — the corpus's own terrestrial calibration rather than bounded away by it.

Ambient field value: Under the canonical excursion normalization, the ambient reference carries the full well depth $\phi_{\text{space}} = \Delta\phi = 3.39 \times 10^9 \text{ GeV}$ in complement units, while the physical excursion $\psi(r)$ decays to the interplanetary baseline over λ_T — no ad-hoc density cutoff enters, since the profile is geometric rather than density-minimum-derived. The normalization is β -independent, so the reference amplitude $\beta = 10^{-4}$ enters only as a linear scaling of the response, not of the field.

Geometry modulation factors: Per-flyby fitted β_{fit} in the six-member sign-consistent ensemble spans roughly a factor of three across the tightly constrained members (extending to a factor of ≈ 15 across the four amplitude-informative fits) — formally divergent for the near-zero-prediction Cassini and MESSENGER members — reflecting substantial geometry-, plasma-, and velocity-dependent modulation encoded in the Step 007 envelope. Four physical mechanisms are tracked in the analysis: (1) *inclination-dependent screening*—higher latitude trajectories sample less equatorial bulge; (2) *J2 oblateness*—altitude-dependent screening from Earth's shape; (3) *plasma environment*—ionospheric density modulates local screening (IRI at perigee, Step 033); and (4) *velocity effects*—disformal coupling in the high-velocity regime. These factors are incorporated into the clock-sector response calculation.

3.2c Component-Level Geometry Factor Analysis

To address the extreme heterogeneity in fitted β_{fit} values across flybys (see Section 4.8), a component-level analysis extracts the effective geometry factor for each flyby independently. The geometry factor isolates trajectory-dependent modulation of the path-resolved scalar response, revealing the physical origin of the observed variation.

The effective geometry factor is defined as the ratio of observed anomaly to the gradient prediction at the reference coupling:

$$G_{i,\text{eff}} = \frac{\Delta v_{i,\text{obs}}}{\Delta v_{\text{grad},i}(\beta_{\text{fit},0} = 10^{-4})} \quad (8)$$

This factor absorbs all geometry-dependent modulation—altitude, J2 oblateness, trajectory asymmetry, velocity-dependent disformal coupling, plasma screening, and OD absorption—into a single observable per flyby. The implied reference-scale amplitude is then $\beta_{0,\text{implied}} = 10^{-4} \times G_{i,\text{eff}}$.

Correlation analysis tests whether $G_{i,\text{eff}}$ varies systematically with trajectory parameters:

- Altitude: higher perigee $\rightarrow$ lower G_{eff} (weaker field gradient)
- Velocity: higher $v_p \rightarrow$ lower G_{eff} (shorter field exposure time)
- Asymmetry: positive $\cos \delta_{\text{in}} - \cos \delta_{\text{out}} \rightarrow$ higher G_{eff} (stronger disformal enhancement)

A multiple linear regression in log space quantifies the combined contribution:

$$\log_{10} |G_{\text{eff}}| = c_0 + c_1 \tilde{h} + c_2 \tilde{v} + c_3 \tilde{a} \quad (9)$$

where $\tilde{h}$, $\tilde{v}$, $\tilde{a}$ are normalized altitude, velocity, and asymmetry. Non-zero coefficients confirm geometry-dependent TEP coupling; R^2 near unity indicates that the three trajectory parameters explain most of the observed heterogeneity.

This approach is complemented by a four-parameter hierarchical Bayesian model ($\beta_{\text{fit},0}$, b_{disf} , σ , α_{res}) sampled via MCMC. The pre-computed gradient and disformal components from the TEP clock-sector response model contain the full perigee physics; the likelihood scales these channels through population-level hyperparameters, with any residual unmodeled modulation captured by α_{res} . Posterior predictive checks validate the model against per-flyby observations.

3.2d Parameter budget: fitted coupling versus envelope heuristics

The restricted Step 008 tier fits a *single* amplitude parameter (shared β_{fit} rescaling of the reference prediction at $\beta_{A,\text{ref}} = 10^{-4}$) on the sign-gated high-S/N ensemble. Separately, the Step 007 geometry envelope carries seven *nominal* deterministic coefficients (inclination scale, J2/altitude modulation, plasma-density power-law parameters, velocity-screening exponent, and near-symmetry coherence threshold) documented in `scripts/utils/tep_geometry_envelope.py`. These coefficients are *not* varied by the nonlinear least-squares map in Step 008; they are nevertheless physical knobs whose $\pm 50\%$ stress band must be machine-checked because skeptical readers will count them toward an “epicycle” budget.

The pipeline therefore emits an explicit parameter audit in `results/step008_fitting_results.json` (`parameter_budget_audit`), including a defensibility score $\text{defensibility} = k_{\text{fit}} / (k_{\text{fit}} + k_{\text{heur}})$ for the restricted tier, and a conservative BIC/AIC companion row that adds k_{heur} to the effective parameter count as a stress test (`TEP_heuristic_pessimistic` inside `enhanced_validation.model_comparison`). Independent $\pm 50\%$ Monte Carlo sweeps over all seven coefficients are archived in `results/step041_envelope_heuristic_sensitivity.json` (Step 041), including one-at-a-time leverage rankings for the inverse-variance pooled β_{fit} .

Ionospheric screening uses a bounded phenomenological ansatz family compared under identical IRI perigee densities in `results/step017_plasma_modulation.json` (`plasma_ansatz_robustness`). A first-principles scalar-plasma coupling derived from the TEP action remains the correct long-term replacement for these ansätze.

3.3 Ensemble Selection Protocol

The gated ensemble for inverse-variance β_{fit} fitting and Bayesian model comparison is constructed via two *a priori* criteria applied *before* any fitting. These criteria are encoded in `scripts/utils/flyby_ensemble.py` and contain no mission-specific exclusions.

A Priori Ensemble Gates

1. Signal-to-noise gate: The published anomaly must satisfy $S/N = |\Delta v_{\text{obs}}| / \sigma > 2$. This ensures that the measurement is statistically distinguishable from zero.
2. Sign-agreement gate (configurable): When `parameters.analysis.tep_physics.strict_sign_gate` is true, the observed anomaly must have the same sign as the TEP prediction at $\beta_{A,\text{ref}} = 10^{-4}$. The pipeline default is `false`: opposite-sign rows still receive an amplitude-only reference fit and remain in the primary pool; the sign-agreement-

restricted subset is reported as `recommended_beta_sign_gated_diagnostic` and `sign_agreement_model_comparison` in Step 026.

Flybys failing either gate are excluded from the *fitted ensemble* but remain in the full catalog for hierarchical diagnostics, fixed- β_{ref} cross-catalog stress tests, and literature comparison. The criteria are applied automatically by the pipeline; no human decision enters after the criteria are specified.

Application to the current dataset:

- *NEAR (1998)*: $S/N = 103.5 > 2$; sign agreement (+13.46 mm/s vs +0.30 mm/s predicted at $\beta_{A,\text{ref}}$). Passes both gates.
- *Galileo 1990*: $S/N = 49 > 2$; sign agreement (+3.92 mm/s vs +0.03 mm/s predicted at $\beta_{A,\text{ref}}$). Passes both gates.
- *Rosetta 2005*: $S/N = 60 > 2$; sign agreement (+1.80 mm/s vs +0.03 mm/s predicted at $\beta_{A,\text{ref}}$). Passes both gates.
- *Cassini (1999)*: $S/N = 2.0 \geq 2$; passes sign agreement at $\beta_{A,\text{ref}}$ on the re-transcribed catalogue (−2.00 mm/s observed vs a negative, near-zero predicted total, sign product > 0). The negative predicted total arises because Cassini's high perigee velocity (19.02 km/s $> v_{\text{trans}} \approx 16.8$ km/s) places it on the template's mixed-regime branch, where the conformal-gradient term (-1.1×10^{-5} mm/s) and disformal term ($+5 \times 10^{-7}$ mm/s) nearly cancel. Because the reference prediction is numerically negligible, the closed-form β_{fit} rescaling is formally divergent ($\beta_{\text{fitted}} \approx 19$ with an uncertainty of comparable size); Cassini enters the Step 008 pool and the Step 026 gated likelihood but contributes negligible weight to the pooled amplitude.
- *Galileo 1992*: $S/N = 4.6 > 2$ on the re-transcribed catalogue (-4.60 ± 1.00 mm/s vs -0.007 mm/s predicted at $\beta_{A,\text{ref}}$, sign agreement). Passes both gates; the prediction is negative but underpredicts the observed magnitude by an order, so the fit contributes a positive residual rather than an amplitude anchor.
- *MESSENGER (2005)*: $S/N = 2.0 \geq 2$ ($+0.02 \pm 0.01$ mm/s vs $+6 \times 10^{-8}$ mm/s predicted at $\beta_{A,\text{ref}}$, sign agreement). Passes both gates; the reference prediction is numerically zero, so its closed-form rescaling is formally divergent and contributes negligible weight.
- *Rosetta 2007*: $S/N = 0.4 < 2$ ($+0.02 \pm 0.05$ mm/s). Fails the S/N gate. Excluded.
- *Rosetta 2009*: $S/N = 0$ ($\Delta v = 0.00$ mm/s bound). Fails the S/N gate. Excluded.
- *Juno (2013)*: $\Delta v = 0.00$ mm/s. Fails S/N gate ($S/N = 0$). Excluded.

Distinction between caveat and exclusion. The manuscript acknowledges that Galileo 1990's high-gain antenna failure and spin-rate changes introduce additional spacecraft-specific uncertainty (Section 5.13). This is a *caveat*—a flag for cautious interpretation—not an *exclusion criterion*. Galileo 1990 remains in the fitted ensemble because it satisfies the pre-specified objective gates. Conflating a stated caveat with an operational exclusion would be a methodological error.

3.4 Deterministic Factor Computation

Deterministic Factors

- Trajectory geometry (G_{traj}): $G_{\text{traj}} = \exp(-(h - 300 \text{ km})/2000 \text{ km}) \times (1 + |\cos \delta_{\text{asym}}|)$
- Clock-amplitude response factor ($S_{\oplus}$): $S_{\oplus} = (R_{\oplus} - R_{\text{sol}})/R_{\oplus}$ with $R_{\text{sol}} \approx 4146$ km (UCD saturation radius; matches `CHARACTERISTIC_SUPPRESSION` in `scripts/utils/physics.py`)
- OD absorption (F_{OD}): Mission-specific fraction of injected TEP signal surviving standard OD processing. Step 039 withholds post-OD columns until Step 021 supplies defensible mission OD configuration data.
- Plasma factor (F_{plasma}): bounded heuristic envelope coefficient on IRI-derived densities, modulated by solar activity indices (F10.7 flux, Kp index); the stronger Debye attenuation ansatz is quarantined at $S_{\text{plasma}} = 1$ (Section 3.10)
- Disformal factor (F_{disf}): velocity-activated sign reversal of the response template for $v > 16.8$ km/s with negative asymmetry (phenomenological ansatz, §3.5)

3.4a Deterministic Geometry Modulation Analysis (Step 009)

With $n = 6$ sign-consistent detections, formal variance decomposition into structural, observational, environmental, and residual components remains statistically underpowered: the standard error on a sample variance estimate with $\text{ddof} = 1$ is a large fraction of the estimate at this sample size, rendering any reported percentage unreliable. Step 009 therefore does not produce an ANOVA-style partition. Instead, it reports three complementary analyses that are statistically defensible at small n :

1. Beta scatter statistics. The raw span and log-standard deviation of fitted β_{fit} are reported as descriptive statistics, with an explicit note that the Step 007 prediction already includes the geometry envelope, so residual scatter reflects genuine coupling heterogeneity or model incompleteness rather than unmodeled trajectory geometry.

2. Full-catalog detection pattern. Across all $n = 12$ catalogued flybys, the Step 007 deterministic prediction is classified against the published observation as true positive (predicted and observed detection), true negative (predicted and observed null), false negative, or false positive. This yields a classification accuracy that tests whether the model correctly identifies which flybys should show anomalies independent of the small fitted- β_{fit} sample.
3. Rank correlation. A Spearman rank correlation between predicted and observed anomaly magnitudes across the full catalog tests whether the model captures the relative ordering of anomaly sizes, again independent of the $n = 6$ β -fitting sample.

The legacy four-stage chained-heuristic output (`variance_decomposition.stages`) is retained in the JSON for backward compatibility but is explicitly marked `DEPRECATED` and must not be used for manuscript inference.

3.5 Disformal Transition Criterion

A disformal transition criterion Ξ is defined to classify flybys into conformal-dominated, mixed, or disformal-dominated regimes of the fitted response template. This provides a formal regime label for Cassini's velocity-activated branch (the status of the underlying scale is assessed below).

$$\Xi_i = \left(\frac{v_i}{v_{\text{trans}}} \right)^p \times |\cos \delta_{\text{in}} - \cos \delta_{\text{out}}| \times \left(\frac{|\nabla \phi_i|}{|\nabla \phi_{\oplus}|} \right)^q \times \text{sgn}(\cos \delta_{\text{in}} - \cos \delta_{\text{out}}) \quad (10)$$

where:

- $v_{\text{trans}} \approx 16.8$ km/s is the transition velocity (an empirical scale of the response template, see below)
- v_i is the flyby perigee velocity
- $p = 2$ is the velocity exponent
- $q = 1$ is the gradient exponent
- $\nabla \phi_{\oplus}$ is the field-gradient excursion at Earth's surface
- $\nabla \phi_i$ is the field-gradient excursion at flyby altitude
- sgn indicates aligned (positive) vs anti-aligned (negative) disformal response

Status of the Transition Scale v_{trans}

The transition velocity enters the model as a fixed scale of the velocity-dependence response template — a pipeline constant carrying an assigned $\pm 20\%$ uncertainty — not as an output of the TEP field equations. The distinction matters and is stated explicitly here, because an earlier draft presented v_{trans} as a derived quantity. The disformal term of the TEP matter metric,

$$ds^2 = A^2(\phi)c^2dt^2 - A^2(\phi)d\mathbf{x}^2 + B(\phi)\partial_\mu\phi\partial_\nu\phi dx^\mu dx^\nu, \quad (11)$$

would produce a genuine velocity-dependent response at the scale where the disformal contribution becomes comparable to the conformal perturbation, $B(\phi)(\mathbf{v} \cdot \nabla \phi)^2/c^2 \sim A^2(\phi) - 1$. Evaluated on the solved Earth profile under the canonical envelope $B(\phi) \propto B_0\phi^2$, however, that balance is not realized at flyby velocities: the microscopic coefficient required is $B_0 \sim 1.6 \times 10^{18}$, exceeding both the calibrated Paper-0/28 envelope normalization by ~ 21 orders of magnitude and the projected 10^{-19} -fractional holonomy bound ($|B_0| \lesssim 5 \times 10^{-8}$) by ~ 26 orders (Step 011 audit). Under the canonical action, therefore, no kilometre-per-second disformal transition exists, and flyby phenomenology at these velocities is conformal-sector — consistent with the fitted disformal amplitude of the hierarchical layer, $b_{\text{disf}} = 0.026 \pm 0.043$, which is consistent with zero and is the expected outcome on the admissible branch (Section 4.2.2).

The coefficient $B_{\text{eff}} = 4|\beta_A|R_{\oplus}/\lambda_T \approx 6.1$ quoted in earlier drafts is accordingly disclosed for what it is: a bookkeeping identity of the template rather than a derived coefficient. Writing $B_{\text{eff}} \equiv (A^2 - 1)/(v_{\text{trans}}/c)^2$ defines it through the balance condition itself, so inserting it back into the balance returns the assumed scale; it is not an action-level or shell-averaged property of the solved field. For the same reason the closed form

$$v_{\text{trans}} = \frac{c}{\sqrt{2}} \left(\frac{\lambda_T}{R_{\oplus}} \right)^{1/2} \left(\frac{|\nabla \phi_{\oplus}| \lambda_T}{M_{\text{Pl}} \epsilon_\phi} \right)^{1/2}, \quad (12)$$

where $\epsilon_\phi \equiv |\nabla \phi_{\oplus}| \lambda_T / M_{\text{Pl}}$ is the dimensionless field excursion across one relaxation length of the screened profile (the Helmholtz solution of $\nabla^2 \phi - \lambda_T^{-2} \phi = -(\beta_A/M_{\text{Pl}})\rho$ for the uniform-sphere source), is retained here only as a dimensional anchor. Evaluated

at the UCD-pinned surface combination $\epsilon_\phi = |\beta_A| \rho_T \lambda_T^2 / M_{Pl}^2 \approx 6.6 \times 10^{-9}$ (the field pinned at the saturation-density source inside R_{sol} ; the documented $\pm 40\%$ GNSS uncertainty on ρ_T spans $\epsilon_\phi \approx (4.0\text{--}9.2) \times 10^{-9}$), the combination evaluates to 14--16.5 km/s, numerically bracketing the 16.8 km/s template value used below (Step 011 audit). This proximity — the geometric combination of the clock-offset velocity $c\sqrt{2\epsilon_\phi}$ diluted by the layer aspect ratio $\sqrt{\lambda_T/4R_\oplus}$ — is recorded as a consistency observation that motivates the template scale; it is not a derivation, and no claim of a field-theoretic transition is made. The UCD-derived characteristic suppression $S_\oplus = (R_\oplus - R_{sol})/R_\oplus \approx 0.35$ is connected to the same field solution through the surface-to-transition geometry and continues to supply the amplitude prior (Step 010).

Within the template, v_{trans} therefore marks the empirical scale separating the velocity-suppressed and velocity-unsaturated response regimes of the fitted envelope. Cassini's high perigee velocity (19.02 km/s $> v_{trans}$) places it on the velocity-activated branch of the template; the physical anomaly mechanism claimed in this paper remains the clock-sector channel (Section 4.6.3a), and the regime labels below are phenomenological descriptors of the fitted response, not statements about a propagating disformal sector.

Classification (by $|\Xi|$):

- $|\Xi| < 0.05$: Conformal-dominated
- $0.05 \leq |\Xi| \leq 0.10$: Mixed
- $|\Xi| > 0.10$: Disformal-dominated

The sign of Ξ encodes the orientation of the template's disformal-term response: positive for aligned trajectories and negative for anti-aligned trajectories. Cassini, with its high perigee velocity (19.02 km/s) and negative asymmetry, falls into the mixed regime with a negative sign — on the template reading, it occupies the anti-aligned branch where the conformal-gradient and disformal terms partially cancel. As established above, this classification is a phenomenological descriptor of the fitted envelope rather than a derived field-theoretic regime assignment.

Velocity shift formula: The predicted velocity anomaly combines four physical effects:

$$\Delta v_{TEP} = \underbrace{\frac{\beta_{A,eff} c^2}{M_{Pl}}}_{\text{field gradient}} \cdot \underbrace{\left. \frac{d\phi}{dr} \right|_{r_p}}_{\text{perigee time}} \cdot \underbrace{\frac{r_p}{v_p}}_{\text{non-radial fraction}} \cdot \underbrace{J_2 \left(\frac{R_\oplus}{r_p} \right)^2}_{\text{trajectory asymmetry}} \cdot (\cos \delta_{in} - \cos \delta_{out}) \quad (13)$$

Each factor has a distinct physical origin:

1. Field amplitude excursion $d\phi/dr = (\Delta\phi/\lambda_T) e^{-h/\lambda_T}$: the clock-sector response strength at perigee altitude h , decaying exponentially with the GNSS-established relaxation length. Lower flybys experience stronger field excursions.
2. Perigee dwell time r_p/v_p : the effective duration of the close encounter. Slower, lower flybys accumulate larger proper-time offsets.
3. J_2 oblateness $J_2(R_\oplus/r_p)^2$: the non-radial component of the offset arising from Earth's oblateness. The radial component is absorbed into the orbit determination program's estimate of GM ; only the non-radial residual produces an apparent net velocity change.
4. Trajectory asymmetry $\cos \delta_{in} - \cos \delta_{out}$: the difference in approach and departure v_∞ declinations (from Anderson et al. (2008)). This factor determines how asymmetrically the spacecraft samples the oblate field. For symmetric trajectories ($\delta_{in} \approx \delta_{out}$), the non-radial impulse cancels and the predicted anomaly vanishes—correctly predicting null results for flybys such as Galileo 1992 and MESSENGER.

3.6 Robust Bayesian Fitting

The primary Step 008 coupling estimate is an inverse-variance scaling fit on the sign-gated detections, and the Step 026 model-comparison layer uses Gaussian weighted least-squares likelihoods. A Student's t-distribution likelihood with degrees of freedom $\nu = 3$ is used only in the auxiliary robust Bayesian/hierarchical checks, where it tests whether the conclusions are sensitive to outlier treatment in the small sample.

Primary fit (Step 008). After the pre-specified S/N and sign gates, each retained flyby maps the Step 007 prediction at $\beta_{ref} = 10^{-4}$ to a single effective amplitude using the linear rescaling $\beta_{fit} = \beta_{ref} (\Delta v_{obs}/\Delta v_{TEP})$ implied by the canonical β -independent excursion normalization (Section 3.3). This is a deterministic per-flyby consistency check, not an iterative maximisation of a heavy-tailed likelihood; it replaces the v0.1 $(\Delta v)^{4/3}$ map, which traced to the β -dependence of the excluded inverse-power density minimum rather than to the canonical action.

Hierarchical population model (Step 015). The four-parameter MCMC layer uses independent Gaussian residuals for each flyby with variance $\sigma_{pop}^2 + \sigma_{obs,i}^2$ in the log-likelihood (see `scripts/steps/step_015_hierarchical_bayesian.py`). Sampling targets the log-posterior $\ln p(\theta | \text{data})$ under the stated priors; reported medians and credible intervals are standard MCMC summaries.

$$\ln \mathcal{L}_i = -\frac{1}{2} \left(\frac{\Delta v_{\text{obs},i} - \Delta v_{\text{pred},i}(\theta)}{\sigma_i} \right)^2 - \frac{1}{2} \ln(2\pi\sigma_i^2), \quad \sigma_i^2 = \sigma_{\text{pop}}^2 + \sigma_{\text{obs},i}^2 \quad (14)$$

Auxiliary tail diagnostics. Cochran's Q , I^2 , and Student- t critical values are used only when summarising scatter among the per-flyby $\beta_{A,i}$ and when constructing heterogeneity-aware error bars in Step 008 (`analyze_fit_quality`); they do not replace the primary scaling law above.

PPN constraint validation is reported in Section 3.9 and Section 4.6.

3.7 Statistical Analysis

The weighted mean $\bar{\beta}_{\text{fit}}$ across all detections is:

$$\bar{\beta}_{\text{fit}} = \frac{\sum_i w_i \beta_{\text{fit},i}}{\sum_i w_i}, \quad w_i = \frac{1}{\sigma_{\beta_{\text{fit},i}}^2} \quad (15)$$

with inverse-variance weights derived from propagated measurement uncertainties. The weighted standard error is:

$$\sigma_{\bar{\beta}_{\text{fit}}} = \left(\sum_i w_i \right)^{-1/2} \quad (16)$$

The NEAR detection dominates the pooled estimate through its fitted- β uncertainty ($\sigma_{\beta} \approx 2.3 \times 10^{-5}$, versus 4.9×10^{-5} for Rosetta 2005, 2.4×10^{-4} for Galileo 1990, and formally divergent values for the near-zero-prediction members).

Heterogeneity assessment: Following meta-analysis conventions (Higgins & Thompson, 2002), heterogeneity is quantified using:

$$Q = \sum_i w_i (\beta_{\text{fit},i} - \bar{\beta}_{\text{fit}})^2 \quad (\text{Cochran's } Q) \quad (17)$$

$$I^2 = \max\left(0, \frac{Q - (n-1)}{Q}\right) \times 100\% \quad (\text{percentage variance due to heterogeneity; Higgins 2002}) \quad (18)$$

An $I^2 > 75\%$ indicates extreme heterogeneity, justifying uncertainty inflation by $\sqrt{Q/(n-1)}$ to account for model scatter beyond measurement error.

Robustness verification: Step 008 parametric bootstrap (10^4 draws) yields median $\beta \approx 4.77 \times 10^{-3}$ with 95% interval $[4.42 \times 10^{-3}, 1.47 \times 10^{-2}]$, and leave-one-out recomputations 5.88×10^{-3} (without NEAR), 4.58×10^{-3} (without Galileo 1990), and 4.64×10^{-3} (without Rosetta 2005). The stability coefficient ≈ 0.092 is below the 0.5 robustness guideline, indicating moderate leave-one-out stability on the gated trio.

1. *Smooth bootstrap* ($n = 10\,000$): Resampling with replacement combined with Gaussian noise injection validates the weighted mean distribution and provides confidence intervals.
2. *Leave-one-out cross-validation*: Systematically excluding each detection verifies that no single flyby dominates the conclusion. Stability coefficient < 0.5 indicates robustness.

3.8 Orbit Determination Filtering Mechanism (Hypothesis)

Modern orbit determination (OD) employs a multi-stage processing pipeline that may inadvertently filter TEP-like signals. Understanding this potential mechanism is relevant for interpreting why some flybys show null results despite TEP predictions. This remains a hypothesis requiring independent verification through raw DSN data analysis.

Standard OD processing chain:

1. Raw Doppler measurements: Two-way/3-way Doppler tracking from DSN stations, typically at X-band (8.4 GHz) or Ka-band (32 GHz), with sampling rates of 1-60 Hz.
2. Cycle-slip detection and correction: Automated algorithms detect discontinuities in phase measurements and correct them to maintain phase continuity.
3. Outlier rejection: Measurements deviating by more than 3σ from the expected trajectory are flagged and removed as erroneous data points.
4. Smoothing and averaging: Raw measurements are typically averaged over 10-60 second intervals to reduce noise and computational load.
5. Bias estimation and removal: Systematic biases (e.g., station clock offsets, media delays) are estimated and subtracted from the measurements.
6. Empirical acceleration estimation: To absorb unmodeled forces, OD fits empirical accelerations (constant, once-per-revolution, stochastic) that absorb any residual systematic errors.
7. Residual analysis: Final residuals are examined; large residuals trigger additional data editing or model refinement.

Hypothesized filtering of TEP signals: TEP produces a sudden velocity shift precisely at perigee passage (± 2 hours), characterized by:

- Sharp temporal structure (not gradual acceleration)
- Correlation with gravitational potential gradient
- Consistent amplitude across multiple spacecraft geometries
- Occurrence at a predictable location (perigee)

These characteristics could cause TEP signals to be treated as systematic errors in the OD pipeline:

1. Outlier rejection: The sharp perigee anomaly could appear as an outlier in the Doppler residuals and be removed by the 3σ threshold.
2. Empirical acceleration absorption: The sudden velocity shift could be absorbed by empirical acceleration terms, effectively modeling it as a force rather than a clock rate effect.
3. Smoothing: Averaging over 10-60 second intervals could dilute the sharp perigee signal, reducing its amplitude.
4. Bias estimation: The perigee anomaly could be partially absorbed into station bias estimates.

Proposed minimal OD approach for validation: To test whether TEP signals can be recovered from raw data, a minimal OD pipeline is recommended:

- Use reduced gravity field (10×10 instead of 50×50 or higher)
- Disable empirical acceleration estimation
- Disable outlier rejection (or use relaxed threshold)
- Use raw Doppler without smoothing
- Fit only initial state and solar radiation pressure coefficient

This minimal approach would preserve TEP signals while still providing adequate orbit determination for anomaly extraction. Where the pipeline reports perigee-window statistics from sequential pairwise Doppler differences (Steps 006 and 030), those quantities are explicitly *not* commensurate with published post-OD Δv anomalies unless a batch OD residual chain is bound; model comparison for the flyby ensemble (Step 026) uses a fractional misspecification floor as the headline likelihood. Each signal-bearing model carries a per-flyby systematic uncertainty $\sigma_{\text{sys},i} = CV_\beta |\hat{y}_i|$, where $CV_\beta = \tau/\bar{\beta} \approx 0.47$ is the measured between-flyby amplitude heterogeneity (DerSimonian–Laird $\tau \approx 2.23 \times 10^{-3}$ on the Step 008 scale factors; $\bar{\beta} \approx 4.74 \times 10^{-3}$) and $\hat{y}_i$ is the model's own fitted prediction, iterated to a fixed point. The floor is dimensionless and therefore invariant under the arbitrary reference-amplitude anchor β_{ref} — a property the earlier constant geometry-spread floor lacked, since the spread of predictions evaluated at β_{ref} rescales with that anchor and can alter the model ordering without any change to the data. The Null model predicts zero signal and accordingly receives no floor: its published uncertainties are taken at face value, so the large anomaly detections penalize it directly. Per-flyby floors under the fitted models span $\approx 0\text{--}14$ mm/s, concentrated where a model claims a large signal. The former constant geometry-spread floor and a published-uncertainties-only block ($\sigma_{\text{sys}} = 0$) are retained as labelled sensitivity analyses together with a constant-floor sweep, not the primary evidence claim.

3.9 PPN Constraints and Cassini Solar Conjunction

For scalar-tensor theories with conformal coupling, the PPN parameter deviation is bounded by

$$\alpha_{\text{eff}}^{(\odot)} = S_{\Sigma}^{(\odot)} \alpha_0, \quad \gamma_{\text{PPN}} - 1 = -\frac{2\beta_A^2 S_{\Sigma}^{(\odot)}}{1 + 2\beta_A^2 S_{\Sigma}^{(\odot)}} \simeq -2\beta_A^2 S_{\Sigma}^{(\odot)} \quad \left(\text{isotropic benchmark: } -\frac{2\alpha_0 \alpha_{\text{eff}}^{(\odot)}}{1 + \alpha_0 \alpha_{\text{eff}}^{(\odot)}} \simeq -4\alpha_0 \right)$$

where γ is the PPN parameter, $\alpha_0 = \sqrt{2}\beta_A = -\sqrt{2}$ is the DEF-normalized coupling ($\beta_A = -1$ frozen; the deviation is linear in the screened source charge because the photon probe is unscreened), and $S_{\Sigma}^{(\odot)}$ is the Solar-System source-charge screening factor. The Cassini solar conjunction experiment provides the tightest bound on the post-Newtonian light-propagation sector. It measured the gravitationally induced frequency shift of radio photons exchanged with the spacecraft and obtained $\gamma = 1 + (2.1 \pm 2.3) \times 10^{-5}$. The Earth-vicinity flyby response geometry $S_{\oplus}(r)$ is a distinct projection of the same environment-dependent scalar configuration; it governs the flyby impulse but is not the Cassini source-charge parameter.

Cassini constrains the reciprocity-even radio light-time observable in the screened solar-system environment. In the TEP decomposition, this constrains three specific sectors:

A. Gravitational/light-propagation sector (directly constrained): Cassini requires that any unscreened solar scalar charge, any long-range conformal/disformal coupling affecting the radio link, or any deviation in the solar-system Shapiro sector be smaller than roughly the measured γ uncertainty: $|\gamma - 1| \lesssim 2.3 \times 10^{-5}$.

B. Conformal clock-sector structure (not directly tested): A purely conformal transformation $\tilde{g}_{\mu\nu} = A^2(\phi)g_{\mu\nu}$ preserves null cones. Therefore, a conformal clock-sector field can evade a direct Cassini light-cone constraint only if it does not create an observable solar-system γ shift or anomalous clock/redshift signature.

C. Screening sector (boundary condition): For the solved high-acceleration solar configuration, where the kinetic response suppresses Temporal Shear, then Cassini becomes a boundary condition: $\Sigma_{\mu} = \nabla_{\mu} \ln A \approx 0$ in the solar-system Shapiro regime. This is not a weakness but exactly how the theory must be formulated.

Therefore Cassini should be treated not as irrelevant to TEP, but as a stringent boundary condition: a viable TEP model must reduce to the GR PPN light-propagation limit near the Sun while reserving its discriminating predictions for observables outside the Cassini measurement class (spatial clock covariance, one-way residual shear, low-density temporal-shear recovery).

The deep potential well of the Sun suppresses Temporal Shear toward zero, providing screening in the solar environment. The UCD-derived characteristic ratio $S_{\oplus} \approx 0.35$ at Earth's surface calibrates the clock-sector response profile along flyby trajectories, while the solar-screening calculation (Section 4.6.1a) shows that the effective coupling along the Cassini radio path also remains below the Cassini bound.

3.10 Plasma Modulation

The Cassini flyby exhibits a near-cancellation regime where the conformal-gradient term is negative (-5.6×10^{-5} mm/s) and the disformal term is positive ($+2.5 \times 10^{-6}$ mm/s), yielding a numerically negligible negative total (-5.3×10^{-5} mm/s) at the reference coupling. The historical sign mismatch was a catalogue transcription artifact: the primary-source anomaly is itself negative (-2 ± 1 mm/s, Anderson et al. 2008), so the predicted sign agrees and the residual gap is one of magnitude — corrected in the present re-transcribed catalogue. Plasma-dependent attenuation is treated as a secondary modulation effect.

The plasma density along the flyby trajectory is computed using:

$$n_{\text{plasma}}(h) = n_{\text{iono}}(h) + n_{\text{mag}}(h) \quad (20)$$

where the ionospheric component is obtained from the International Reference Ionosphere (IRI) empirical model (Step 033), which provides continuous electron density profiles along spacecraft trajectories using historical F10.7 solar flux data. The IRI model replaces the Chapman layer approximation with real ionospheric data, improving accuracy for plasma environment reconstruction (Step 020). For theoretical reference, the Chapman layer model is:

$$n_{\text{iono}}(h) = n_{\text{max}} \exp \left[0.5 \left(1 - \frac{h - h_{\text{max}}}{H_{\text{scale}}} - e^{-(h - h_{\text{max}})/H_{\text{scale}}} \right) \right] \quad (21)$$

with $h_{\text{max}} = 300$ km, $H_{\text{scale}} = 50$ km, and $n_{\text{max}} = 10^6$ cm $^{-3}$ (solar maximum). The magnetospheric component scales with L-shell as $n_{\text{mag}} \propto L^{-4}$.

A Debye-like plasma attenuation ansatz is used as a phenomenological proxy for ionospheric screening:

$$S_{\text{plasma}} = \exp\left(-\frac{n_e}{n_{\text{ref}}}\right) \quad (22)$$

where n_e is the electron density in cm^{-3} and $n_{\text{ref}} = 10^4 \text{ cm}^{-3}$ is a reference density. Plasma attenuation is retained only as a heuristic sensitivity layer; the primary TEP prediction uses $S_{\text{plasma}} = 1$, so the headline result does not depend on the phenomenological plasma ansatz.

Plasma attenuation does not cause sign reversal—it only modulates the magnitude of the scalar field. The primary mechanism for sign reversal is disformal coupling (Section 3.5), which produces velocity-dependent effects for high-velocity anti-aligned trajectories.

Solar activity data for plasma density estimation are obtained from documented historical records: F10.7 solar flux from NOAA/SWPC and the Kp geomagnetic index from the GFZ German Research Centre for Geosciences. The current implementation uses continuous International Reference Ionosphere (IRI) model electron density data fetched for the exact historical trajectories of each flyby (Step 033) and ingested by the plasma environment reconstruction step (Step 020). The IRI model is a well-validated empirical model based on decades of ionospheric measurements. Step 033 is configured so its mission keys mirror the cached JPL Horizons trees under `data/raw/jpl_horizons/<mission>`, and the Step 017 lookup table (`iri_mission_map`) maps catalogue names to those keys; together this prevents a silent plasma-path gap for a flyby that already has a reconstructed Horizons arc.

The IRI electron density and computed phenomenological screening factor at perigee show that the ansatz predicts stronger attenuation at lower altitudes (higher plasma density), which is physically intuitive. Rosetta 2007 at 5430 km has the weakest attenuation ($S_{\text{plasma}} \approx 0.96$) because it samples the most tenuous plasma environment, while NEAR at 539 km has the strongest ($S_{\text{plasma}} \approx 1.3 \times 10^{-6}$) due to the dense F-region ionosphere.

Because the Debye ansatz is phenomenological and not derived from the TEP action, the primary predictions in Table 3 use $S_{\text{plasma}} = 1$ (no plasma suppression): the ansatz is computed and recorded for every flyby but is never applied to a prediction. Applying it literally would suppress NEAR’s predicted anomaly to $\sim 5 \times 10^{-6}$ mm/s, far below the observed 13.46 mm/s and inconsistent with the detection. A weaker, bounded power-law plasma factor $f_{\text{plasma}} = (1 + \rho/\rho_{\text{crit}})^{-0.3}$ on IRI-derived perigee densities is distinct from the quarantined ansatz: it remains an active nominal coefficient inside the Step 007 deterministic geometry envelope — one of the seven documented heuristic knobs carrying the machine-checked $\pm 50\%$ stress band (Steps 007, 041) — contributing modulation factors of 0.19–1.00 across the catalogue and 0.18–0.91 on the gated ensemble. It is carried explicitly as a template-layer coefficient, not as an action-derived scalar–plasma coupling, whose derivation remains an open refinement.

3.11 UCD-Motivated Temporal Topology Derivation

To eliminate systematic bias from phenomenological suppression factors, R_{sol} and the characteristic suppression $S_{\oplus}$ are derived from the Temporal Topology Saturation Scale (UCD) saturation model. The saturation radius is calculated from the UCD ansatz using Earth’s total mass and the saturation scale $\rho_T \approx 20 \text{ g/cm}^3$, a scale established across astrophysical systems from dwarf-galaxy cores to galaxy-cluster halos (Paper 6; Schive et al. 2014; Mocz et al. 2018). GNSS atomic-clock correlation analysis (Step 016) yields a covariance decay length $L_c \approx 4201 \text{ km}$. Under the corpus’s terrestrial calibration ansatz — the Paper 6 §2 transfer sketch, which identifies L_c with the network-projected scale of the exterior clock-transport field at $\mathcal{O}(1) \times R_T(M_{\oplus})$ — inverting the saturation formula returns $\rho_T \approx 19.2 \text{ g/cm}^3$, within 3.9% of the adopted 20 g/cm^3 . Because the adopted ρ_T is itself anchored on L_c , this agreement is a shared-input consistency check (the global anchor recovered at its own terrestrial calibration point), not an independent validation; what it does establish is that the density scale carries no tuning to Earth flyby data.

$$R_{\text{sol}} = \left(\frac{3M_{\oplus}}{4\pi\rho_T}\right)^{1/3} \approx 4146 \text{ km} \quad (23)$$

This yields the UCD saturation estimate anchored on the terrestrial GNSS scale ($L_c = 4201 \text{ km} \rightarrow \Delta R/R = 0.34$, 2% agreement with the terrestrial calibration anchor):

$$\frac{\Delta R}{R} = \frac{R_{\oplus} - R_{\text{sol}}}{R_{\oplus}} = 0.349 \approx 0.35 \quad (24)$$

The systematic uncertainty on $\rho_T = 20 \pm 7 \text{ g/cm}^3$ (35%) from Paper 6 (UCD) — a formal scaling constant of the $R_T(M)$ mass relation, not a literal interior density — propagates to $\Delta R_{\text{sol}} \approx \pm 540 \text{ km}$ ($\sim 13\%$) and $\Delta S_{\oplus} \approx \pm 0.09$ ($\sim 25\%$). The GNSS comparison ($L_c = 4201 \pm 1967 \text{ km}$, Step 016) verifies that the adopted global value remains consistent with its own terrestrial

calibration anchor under the Paper 6 §2 identification; it is not an independent determination. The independent empirical content entering the prior is instead supplied by the cross-scale astrophysical consistency of ρ_T (Paper 6) and the flyby altitude-threshold check (Step 010, approach 5). Together, these constraints establish $S_{\oplus} = 0.35^{+0.09}_{-0.09}$ as a rigorously derived prior.

3.12 Cosmographic Temporal Shear Modulation Analysis

A deeper TEP prediction is that the disformal coupling term depends on the total velocity in the ambient-field rest frame — the frame in which the spatial gradient of the cosmological temporal field vanishes. If the CMB dipole frame approximates this frame, the Solar System's ~ 370 km/s bulk motion toward (RA, Dec) = (167.94°, −6.93°) provides a cosmographic modulation of the effective coupling strength. Step 040 tests this using full 3D spacecraft state vectors from JPL Horizons, computing heliocentric distance, CMB dipole projection, and disformal enhancement proxies. With only $n = 9$ flybys carrying both reconstructed 3D state vectors and a catalogued anomaly value or published bound (the three modern passes lack a published anomaly), all cosmographic tests remain exploratory; the full extraction protocol is documented in the Step 040 pipeline output.

4. Results

4.1 Individual Flyby Fits

The TEP clock-sector response model with J2/J3/J4 multipole contributions, *disformal coupling*, perigee plasma factors (Step 017 / IRI, Step 033), and *Temporal Topology screening* quantitatively reproduces the gated primary detections as apparent velocity discontinuities at a reference coupling $\beta_{A,\text{ref}} = 10^{-4}$, then rescales to a universal β_{fit} by per-flyby fitting subject to the pre-fit $S/N \geq 2$ admission gate (in the active configuration `strict_sign_gate = false`; sign agreement at $\beta_{A,\text{ref}}$ is tracked as a companion diagnostic rather than an exclusion filter). Evaluated on the ungated primary-source catalogue of Table 1 prior to any S/N or sign filtering, all 7 of 7 published nonzero anomalies (NEAR, Galileo 1990, Cassini, Galileo 1992, Rosetta 2005, Rosetta 2007, and MESSENGER; exact one-sided binomial $p = 0.0078$, `ungated_sign_concordance` in Step 008) — including the corrected Cassini entry (-2 ± 1 mm/s) and the Galileo 1992 detection (-4.60 ± 1.00 mm/s) — match the sign of their TEP reference predictions, so the six $S/N \geq 2$ detections and the sign-concordant subset coincide identically without circular selection. Cassini and MESSENGER have reference predictions of order 10^{-4} – 10^{-6} mm/s at $\beta_{A,\text{ref}}$; their per-flyby rescalings are therefore formally unconstrained (fractional uncertainty $\gtrsim 0.7$) and contribute negligible weight to the pooled amplitude, while remaining in the gated ensemble, the Step 039 classification, and the hierarchical diagnostics (Step 015). Table 3 lists per-flyby predictions at $\beta_{A,\text{ref}}$ and fitted β_{fit} for the ensemble members.

Table 3: Step 008 fitted β_{fit} and fitted $\Delta v_{\text{TEP,fit}}$ for the six-member sign-gated ensemble (reference-coupling predictions at $\beta_{A,\text{ref}} = 10^{-4}$ are shown in Table 3a). Rows marked † have reference predictions ≈ 0 at $\beta_{A,\text{ref}}$; their per-flyby rescalings are unconstrained and carry negligible weight in the pooled mean.

Spacecraft	Date	$\Delta v_{\text{TEP,fit}}$ (mm/s)	Δv_{obs} (mm/s)	β_{fit}	σ_{β}
NEAR	1998-01-23	13.46	13.46	4.43×10^{-3}	4.28×10^{-5}
Galileo 1990	1990-12-08	3.92	3.92	1.44×10^{-2}	2.93×10^{-4}
Cassini †	1999-08-18	−2.00	−2.00	1.89×10^1	9.44
Galileo 1992	1992-12-08	−4.60	−4.60	6.65×10^{-2}	1.44×10^{-2}
Rosetta 2005	2005-03-04	1.80	1.80	5.15×10^{-3}	8.58×10^{-5}
MESSENGER †	2005-08-02	0.02	0.02	3.40×10^1	1.70×10^1

The four amplitude-informative gated fits span 4.43×10^{-3} to 6.65×10^{-2} , with the three tightly constrained members (NEAR, Galileo 1990, Rosetta 2005) clustering within a factor of about 3.3 (4.43×10^{-3} to 1.44×10^{-2}), consistent with geometry-dependent modulation of the effective coupling (the plasma term is a bounded heuristic envelope coefficient; the stronger Debye attenuation ansatz is quarantined at $S_{\text{plasma}} = 1$, Section 3.10). The inverse-variance weighted mean is $\beta_{\text{fit}} = 4.74 \times 10^{-3} \pm 5.83 \times 10^{-4}$ (formal uncertainty from Step 008); the random-effects uncertainty is 1.28×10^{-3} . Cross-validation indicates moderate robustness (stability coefficient $0.092 < 0.5$). Residuals on the six-member gated set are consistent with normality (Shapiro–Wilk $p \approx 0.95$).

Table 3a shows the component-level breakdown for the six gated detections.

Table 3a: Component-Level TEP Predictions at $\beta_{A,\text{ref}} = 10^{-4}$ for the Gated Ensemble

Flyby	Δv_{grad} (mm/s)	Δv_{disf} (mm/s)	Δv_{total} (mm/s)	Δv_{obs} (mm/s)	Regime
NEAR	+0.18	+0.13	+0.30	13.46	Gradient-dominated
Galileo 1990	+0.02	0.00	+0.03	3.92	Mixed gradient-disformal
Rosetta 2005	+0.03	0.00	+0.03	1.80	Gradient-dominated
Cassini	-0.00	0.00	-0.00	-2.00	Sign-consistent, sub-threshold amplitude
Galileo 1992	-0.08	+0.04	-0.03	-4.60	Sign-consistent negative branch
MESSENGER 2005	+0.00	0.00	+0.00	0.02	Negligible-amplitude consistency

4.1.2 Full 3D Trajectory Integration Validation

The perigee approximation used in the primary analysis (Table 3) is cross-checked against full 3D trajectory integration reconstructed from JPL Horizons scalar data via Keplerian orbit propagation (Section 3.2b). Table 3b compares the perigee estimate Δv_{peri} with the integrated velocity shift Δv_{int} for the analyzed flybys.

Table 3b: 3D Trajectory Integration vs Perigee Approximation (Step 011)

Spacecraft	Δv_{peri} (mm/s)	Δv_{int} (mm/s)	Ratio $\Delta v_{\text{int}}/\Delta v_{\text{peri}}$	n_{points}	Path length (km)		
NEAR	+0.304	+6.933	+8.704	1.26	28.67	diagnostic	diagnostic
Galileo 1990	+0.027	+6.376	+8.396	1.32	307.69	diagnostic	diagnostic
Rosetta 2005	+0.035	+5.308	+7.749	1.46	221.68	diagnostic	diagnostic
Cassini	-0.000	+6.118	+8.247	1.35	-778764.53	diagnostic	diagnostic
Galileo 1992	-0.007	+7.280	+8.887	1.22	-1283.79	diagnostic	diagnostic
Rosetta 2007	+0.000	+3.201	+6.153	1.92	465373.66	diagnostic	diagnostic
MESSENGER_2005	+0.000	+4.963	+7.520	1.52	127682306.18	diagnostic	diagnostic
Stardust	-0.000	+2.922	+5.897	2.02	-88513.81	diagnostic	diagnostic
OSIRIS-REx	+0.000	+1.019	+3.547	3.48	8313217.05	diagnostic	diagnostic
BepiColombo	-0.000	+1.460	+4.229	2.90	-61194535.62	diagnostic	diagnostic

† Ratios are not reported where the perigee estimate is numerically zero; they carry no information there.

For the three tightly constrained detections (NEAR, Galileo 1990, Rosetta 2005), the 3D integrated velocity shift exceeds the perigee estimate by factors of 6.2–18.6 at the pooled $\beta_{\text{fit}} = 4.74 \times 10^{-3}$ (a factor of ~ 47.4 in coupling relative to $\beta_{A,\text{ref}}$), with trajectory-curvature and altitude-dependent modulation partially suppressing the signal along the path. NEAR shows the largest ratio (18.6), Galileo 1990 is intermediate (14.6), and Rosetta 2005 shows the smallest ratio (6.2), consistent with its higher perigee altitude (1955 km) where the field gradient varies more significantly along the trajectory. The perigee approximation therefore captures the dominant physics at $\beta_{A,\text{ref}}$ but should not be naively rescaled to the fitted β_{fit} without the geometry envelope. Ratios are not reported where the perigee estimate is numerically zero, since they carry no information there.

Rosetta 2007, Rosetta 2009, and MESSENGER 2005 predict negligible anomalies in both methods, consistent with their published sub-threshold results; Rosetta 2007 is a high-altitude suppression case (5301 km perigee). Galileo 1992 is the informative negative-asymmetry case: the perigee estimate reproduces the observed sign (-0.007 mm/s versus the published -4.60 ± 1.00 mm/s), while the integrated 3D shift changes sign along the extended trajectory ($+0.30$ mm/s at the pooled coupling), so the two layers bound the response differently and both underpredict the observed magnitude. BepiColombo's 12,697 km pass integrates stably to a numerically negligible shift ($\Delta v_{\text{int}} = 7.8 \times 10^{-9}$ mm/s), consistent with the suppression expected at its altitude.

The overall conclusion is that the perigee approximation captures the dominant physics at the reference coupling $\beta_{A,\text{ref}} = 10^{-4}$, while full 3D integration at the pooled $\beta_{\text{fit}} = 4.74 \times 10^{-3}$ provides a cross-check that incorporates trajectory curvature, altitude-dependent field-gradient variation, and the complete geometry envelope. The ensemble fitting retains the perigee approximation as the primary computational tool, with 3D integration serving as a validation layer where numerically stable.

4.2 Hierarchical Bayesian Model Results

4.2.1 Per-Flyby Geometry Factor Extraction

The component-level analysis extracts the effective geometry factor $G_{i,\text{eff}} = \Delta v_{\text{obs}} / \Delta v_{\text{grad}}(\beta_{\text{fit},0} = 10^{-4})$ for each flyby. This factor represents the multiplicative scaling between the observed anomaly and the gradient prediction at the reference coupling, isolating geometry-dependent modulation from the universal coupling strength.

Table 3c: Per-Flyby Effective Geometry Factors

Flyby	G_{eff}	Altitude (km)	Velocity (km/s)	Asymmetry	$\beta_{0,\text{implied}}$
NEAR	75.8	539	12.7	+0.625	7.58×10^{-3}
Galileo 1990	168.3	960	13.7	+0.149	1.68×10^{-2}
Cassini	$1.8 \times 10^5 \dagger$	1175	19.0	−0.022	unconstrained
Galileo 1992	307.8	303	14.1	−0.170	3.08×10^{-2}
Rosetta 2005	51.8	1955	10.5	+0.173	5.18×10^{-3}
Rosetta 2007	$1.6 \times 10^3 \dagger$	5301	12.5	+0.034	unconstrained

The geometry factor spans a wide dynamic range across the gated ensemble. Rows marked $\dagger$ divide by a reference prediction of order 10^{-4} mm/s or smaller, so their G_{eff} is a numerical artefact of the near-zero denominator rather than a constrained amplitude. Among the four amplitude-informative members, the implied couplings span $\beta_{0,\text{implied}} \sim 5.2 \times 10^{-3}$ – 3.1×10^{-2} , bracketing both the Step 008 weighted mean (4.74×10^{-3}) and the hierarchical Bayesian median (5.27×10^{-3}) within the ensemble spread (the G_{eff} normalization divides by the gradient component only, so it sits above the pooled amplitude). Correlations against the trajectory parameters are reported in the Step 015 output: at $n = 6$ the only approach to significance is against velocity ($r = 0.90$, $p = 0.011$); altitude ($r = -0.13$) and asymmetry ($r = -0.28$) show none.

4.2.2 MCMC Hierarchical Inference

The four-parameter hierarchical Bayesian model ($\beta_{\text{fit},0}$, b_{disf} , σ , α_{res}) is sampled via MCMC. The $\log\beta_{\text{fit},0}$ prior is centered on the Step 008 weighted mean with a heterogeneity-aware width ($\beta_{\text{fit},0} \sim \text{LogNormal}(\ln(4.74 \times 10^{-3}), 1.08)$, sourced from the random-effects scatter so the prior stays diffuse on the scale of the observed flyby-to-flyby scatter); the layer therefore functions as an independent population-level estimate rather than an echo of the pooled fit. The pre-computed gradient and disformal components from Step 007 contain the perigee physics; the likelihood scales these components by inferred universal couplings, with residual modulation captured by α_{res} .

Posterior parameter estimates (Step 015, current pipeline):

- $\beta_{\text{fit},0} = 5.27 \times 10^{-3} \pm 2.55 \times 10^{-3}$ (16th–84th: 3.57 – 7.57×10^{-3})
- $b_{\text{disf}} = 0.026 \pm 0.043$
- $\sigma = 1.65 \pm 0.62$ mm/s (flyby-to-flyby scatter)
- $\alpha_{\text{res}} = +0.022 \pm 0.286$ (consistent with zero)

The hierarchical posterior median $\beta_{\text{fit},0} = 5.27 \times 10^{-3}$ is consistent with the Step 008 inverse-variance weighted mean (4.74×10^{-3}) within the heterogeneity width, and the agreement is not prior-imposed: the stored likelihood profile over $\beta_{\text{fit},0}$ — evaluated with the remaining parameters at their posterior medians and no prior weighting — peaks at 5.57×10^{-3} , with the Step 008 centre within $\Delta \ln \mathcal{L} \approx 0.75$ of the maximum. The two estimation layers thereby reconcile on a common $\mathcal{O}(10^{-3})$ amplitude. Posterior predictive checks track NEAR and Rosetta 2005 well; Galileo 1990, Galileo 1992, and Cassini retain residuals of 2.5, 4.0, and 2.0 mm/s respectively in this hierarchical layer, motivating continued envelope and OD work.

Table 3d: Posterior Predictive Checks

Flyby	Δv_{obs} (mm/s)	Δv_{pred} (mm/s)	Residual (mm/s)
NEAR	13.46	13.54 ± 1.85	−0.08
Galileo 1990	3.92	1.39 ± 0.29	+2.53
Cassini	−2.00	$−0.00 \pm 0.00$	−2.00
Galileo 1992	−4.60	$−0.58 \pm 0.33$	−4.02
Rosetta 2005	1.80	1.90 ± 0.50	−0.10

Flyby	Δv_{obs} (mm/s)	Δv_{pred} (mm/s)	Residual (mm/s)
Rosetta 2007	0.02	0.00 ± 0.00	+0.02
MESSENGER 2005	0.02	0.00 ± 0.00	+0.02

The posterior median $\beta_{\text{fit},0} = 5.27 \times 10^{-3}$ agrees with the Step 008 inverse-variance weighted mean (4.74×10^{-3}) within the heterogeneity width; the width is driven by the population scatter $\sigma \approx 1.65$ mm/s absorbed by the hierarchical layer rather than by any residual offset between the two estimators. The factor-of-order-unity rescaling from the reference coupling ($\beta_{A,\text{ref}} = 10^{-4}$) is a computational anchor only: because the canonical excursion normalization is β -independent, the response model is strictly linear ($\Delta v_{\text{TEP}} \propto \beta$), so the closed-form inversion $\beta_{\text{fit}} = \beta_{A,\text{ref}} (\Delta v_{\text{obs}} / \Delta v_{\text{TEP}})$ is algebraically independent of $\beta_{A,\text{ref}}$ — the fitted amplitude is set by the data, and re-evaluating the same flyby at $\beta_{A,\text{ref}} = 10^{-5}$ or 10^{-3} returns the identical fitted value. The residual bookkeeping is therefore a single channel quantity: the population response coefficient $\kappa_{\text{flyby}} = \beta_{\text{fit}} / |\beta_A| \approx 4.7 \times 10^{-3}$ — an unexplained transfer-map coefficient of the same open class as κ_{Cep} and the Paper 0 Γ_X entries, reported pending the action-level derivation rather than absorbed into the theory. The Step 043 clock-channel forward model carries the same bookkeeping in its own normalization: the OD-level rescaling $C_A \approx 4.2$ between reconstructed apparent shifts and published anomalies is the clock-channel analogue of κ_{flyby} and is reported as such.

4.3 Deterministic Geometry Modulation Analysis

With only $n = 6$ sign-gated detections, formal variance decomposition into structural, observational, environmental, and residual components remains underpowered: the standard error on a sample variance estimate with $\text{ddof} = 1$ exceeds 100% of the estimate for $n < 5$ and stays near 50% at $n = 6$. Step 009 therefore does not report ANOVA-style percentages. Instead, three complementary analyses that are defensible at small n are presented.

4.3.1 Beta Scatter Statistics

Across the six gated fits, fitted β spans 4.43×10^{-3} to 3.40×10^{-1} (log STD = 1.77 dex, $n = 6$), dominated by the two unconstrained near-zero-prediction rows; the four amplitude-informative members span 4.43×10^{-3} to 6.65×10^{-2} . Because the Step 007 prediction already includes the geometry envelope, the residual scatter among the constrained members reflects genuine coupling heterogeneity or model incompleteness, not unmodeled trajectory geometry. With $n = 6$, no formal partitioning of this scatter into structural, observational, or environmental buckets is attempted.

4.3.2 Full-Catalog Detection Pattern

Across all $n = 12$ catalogued flybys (Step 009), evaluating the deterministic geometry envelope at the Step 008 pooled coupling $\beta_{\text{pooled}} = 4.74 \times 10^{-3}$ correctly classifies 11/12 flybys at the 0.5 mm/s threshold (accuracy 91.7%: 3 true positives [NEAR +14.39 mm/s predicted vs +13.46 mm/s observed; Galileo 1990 +1.29 mm/s vs +3.92 mm/s; Rosetta 2005 +1.66 mm/s vs +1.80 mm/s], 8 true negatives, 1 false negative [300 km perigee drag-dominated Galileo 1992, predicted at -0.33 mm/s], and 0 false positives), and achieves 12/12 (100% accuracy: 4 true positives, 8 true negatives, 0 false negatives, 0 false positives) at the 0.3 mm/s threshold. (By contrast, at the unfitted placeholder reference $\beta_{A,\text{ref}} = 10^{-4}$ used prior to Step 008, magnitudes sit a factor of 47.4 lower so all four detections fall below 0.5 mm/s while retaining the exact observed signs.) Cassini's corrected -2 mm/s entry sits at the adopted detection boundary (2σ) and is classified as a true negative; the sign of its prediction nevertheless agrees with observation. No false positives are present: no published null or sub-threshold row is predicted as a detection.

4.3.3 Rank Correlation

A Spearman rank correlation between predicted and observed anomaly magnitudes across the full catalog ($n = 12$) yields $\rho = 0.56$ ($p = 0.059$). This indicates moderate rank agreement: the model captures the relative ordering of anomaly sizes, with the largest predicted anomalies corresponding to the largest observed anomalies. The companion ordering diagnostic between $|\text{trajectory asymmetry}|$ and $|\text{observed anomaly magnitude}|$ gives $\rho = 0.56$ ($p = 0.11$, $n = 9$ flybys with published asymmetries) and $\rho = 0.68$ ($p = 0.090$, $n = 7$ nonzero anomalies); within the gated $n = 6$ subset the ordering is positive but not decisive ($\rho = +0.66$, $p = 0.156$): NEAR carries both the largest asymmetry and the largest anomaly in every subset.

The legacy four-stage variance-decomposition output in `results/step009_variance_analysis.json` is retained for backward compatibility but is explicitly marked `DEPRECATED` and must not be used for manuscript inference. The dominant formal heterogeneity statistic remains Cochran Q / I^2 on the six gated β fits (Step 008).

4.4 Disformal Transition Criterion Results

The disformal transition criterion Ξ classifies flybys into conformal-dominated, mixed, or disformal-dominated regimes of the fitted response template, using velocity, asymmetry, and altitude. The definition $\Xi = (v/v_{\text{trans}})^2 \times |\text{asym}| \times (|\nabla\phi|/|\nabla\phi_{\ominus}|) \times \text{sgn}(\text{asym})$ uses the empirical template scale $v_{\text{trans}} \approx 16.8$ km/s (status: phenomenological ansatz, Section 3.5 — under the canonical $B(\phi)$ envelope no kilometre-per-second disformal transition exists, the required microscopic coefficient being excluded by ~ 26 orders of magnitude); the analyzed flybys span multiple template regimes.

Cassini, with its high perigee velocity (19.02 km/s) and weakly negative asymmetry ($\cos_{\text{asymmetry}} = -0.022$), lies in the template's mixed regime where gradient and disformal contributions partially cancel at $\beta_{A,\text{ref}}$, yielding a small negative total prediction (-5.3×10^{-5} mm/s). With the corrected catalogue entry the published anomaly is also negative (-2 ± 1 mm/s), so Cassini passes the sign gate; its residual tension is in amplitude only, since the reference prediction is far below the observation and the unconstrained per-flyby rescaling is absorbed into the negligible-weight fit reported in Table 3. Consistent with the canonical branch, the hierarchical layer's fitted disformal amplitude is null ($b_{\text{disf}} = 0.026 \pm 0.043$), so the velocity-template terms act as phenomenological descriptors rather than a detected propagating sector.

4.4a Full-Catalog Raw Stress Test

Before restricting attention to the sign-gated detections, the universal- β_{fit} prediction is tested against all published rows with explicit raw TEP predictions in Step 039. This raw-layer stress test includes NEAR, Galileo 1990, Rosetta 2005, Cassini, Galileo 1992, MESSENGER, Rosetta 2009, Juno, and Rosetta 2007 ($n = 9$). Flybys with no public anomaly report are also excluded. The uncertainty is $\sigma_{\text{tot}}^2 = \sigma_{\text{obs}}^2 + \sigma_{\text{raw}}^2$, combining the published measurement uncertainty with the propagated universal- β_{fit} prediction uncertainty.

Table 3e: Full-Catalog Raw Stress-Test Likelihood (Step 039)

Quantity	Null	Raw TEP universal- β_{fit}
Included rows	$n = 9$ published rows with explicit raw predictions	
Log likelihood	−81.79	−38.76
χ^2	175.61	89.54
Improvement	$\Delta \log L_{\text{TEP-null}} \approx +43.03$; $\Delta \chi_{\text{null-TEP}}^2 \approx 86.07$	

This table is the headline full-catalog stress test: once the random-effects prediction uncertainty is included, the raw universal- β_{fit} model improves over the null by $\Delta \log L \approx 43.03$. The improvement reflects the large NEAR, Galileo 1990, and Rosetta 2005 signals, while still exposing the remaining stress cases: Galileo 1992 is classified as a raw surplus (correct sign, underpredicted magnitude; the uncertainty-aware scheme flags it as an observed detection with an uncertain prediction), and Juno remains the explicit raw-tension null (predicted $+0.12$ mm/s at the pooled β_{fit} , a $\sim 3\sigma$ tension under the propagated uncertainty). These results are therefore stronger than a pure three-row fit, but they are not a post-OD mission likelihood because the F_{OD} columns remain withheld until real OD configuration data are available.

4.5 Bayesian Model Comparison

The four-tier model comparison below evaluates Null, Anderson, TEP restricted, and TEP flexible models on the full catalog of nine flybys with published observations ($n = 9$), using Gaussian likelihoods and systematic uncertainty from the Step 026 heterogeneity budget. The TEP restricted amplitude β_{fit} is fitted on the six-member sign-gated ensemble (all published detections now pass the sign gate on the corrected catalogue), then applied to the full $n = 9$ catalog for the model-comparison likelihood. This tests whether a coupling inferred from the detections also compresses the full catalog.

4.5.1 Model Definitions and Parameter Status

Model	Fitted Parameters	Pre-specified Quantities	Description
Null (M_0)	0	—	Predicts $\Delta v = 0$ for all flybys.
Anderson Empirical (M_A)	2 (A, B)	Geometry (declinations) from JPL Horizons	$\Delta v = A(\cos \delta_{\text{in}} - \cos \delta_{\text{out}}) + B$. Captures the core empirical correlation identified by Anderson et al. (2008). Perigee latitude is omitted because it is not catalogued.
TEP Restricted (M_T^{res})	1 (β_{fit})	$\lambda_T = 4200$ km (GNSS Step 016); $S_{\oplus} \approx 0.35$ (UCD Step 010); $v_{\text{trans}} \approx 16.8$ km/s	$\Delta v = dv_{\text{pred}}^{\text{base}}(\beta_{\text{fit}}/\beta_{\text{ref}})$. All physics except the coupling amplitude is pre-specified from independent data or stated response templates; the excursion

Model	Fitted Parameters	Pre-specified Quantities	Description
		(empirical template scale); geometry from JPL Horizons	normalization is β -independent under the canonical Paper 0 amplitude sector, so the response is linear.
TEP Flexible (M_T^{flex})	3 (β_{fit} , b_{disf} , offset)	Same pre-specified quantities as restricted	$\Delta v = (\beta_{\text{fit}}/\beta_{\text{ref}})(dv_{\text{grad}} + b_{\text{disf}} dv_{\text{disf}}) + \text{offset}$. Allows disformal amplitude and residual modulation (plasma, OD) to vary freely.

4.5.2 Log-likelihoods and Information Criteria

The archived Step 026 models are trajectory-geometry templates, not distinct clock-sector predictions for the coherent catalogue. Their iteratively weighted least-squares scores are retained below for reproducibility. Because the fractional variance depends on the fitted prediction, this fixed-point procedure is not a certificate of maximum likelihood for the full heteroscedastic model; the fitted scatter and template choices also require explicit nuisance accounting. The following log-likelihood, AIC and BIC values are therefore archived convention-dependent scores, not calibrated evidence for the clock mechanism:

- Null: log L = -8360.96, AIC = 16721.9, BIC = 16721.9
- Anderson: log L = -11.59, AIC = 27.2, BIC = 27.6
- TEP restricted: log L = -15.09, AIC = 32.2, BIC = 32.4
- TEP flexible: log L = -30.88, AIC = 67.8, BIC = 68.4
- TEP restricted (pessimistic): $k_{\text{eff}} = 8$ (1 fitted β_{fit} + 7 heuristic envelope coefficients), BIC = 47.8

Sign-gated ensemble ($n = 6$) log-likelihoods under the same fractional floor: Null log L = -8368.02, Anderson log L = -6.55, TEP restricted log L = -16.97, TEP flexible log L = -35.42 (BICs 16736.0, 16.7, 35.7, 76.2 respectively). The gated sensitivity at $\sigma_{\text{sys}} = 0$ — published uncertainties taken at face value — favours TEP restricted over Anderson by $\Delta\text{BIC} \approx +848$ (log L -582.0 versus -1005.2), the same direction as the full-catalog $\sigma_{\text{sys}} = 0$ result (+2767); the nominal Anderson preference exists only under a nonzero systematic floor whose scale is a modelling convention rather than a measurement.

4.5.3 Information Criteria and Model Comparison

Information criteria for the full $n = 9$ catalog:

- Anderson vs Null: $\Delta\text{BIC} \approx 16694.3$
- TEP restricted vs Null: $\Delta\text{BIC} \approx 16689.5$
- TEP flexible vs Null: $\Delta\text{BIC} \approx 16653.6$
- TEP restricted vs Anderson: $\Delta\text{BIC} \approx -4.8$ (parity; Anderson nominally first)

Interpretation. The three geometric templates organize the catalogue with archived score differences $\Delta\text{BIC} \gtrsim 1.66 \times 10^4$ against the zero-signal Null under the adopted uncertainty convention. These scores quantify geometric compression of the trajectory-asymmetry pattern; the link-class condition under which the clock-sector mechanism populates this template (requiring an independent onboard clock reference rather than coherent two-way turnaround) is established in Section 5.2, and because the Null is assigned no additional fractional scatter, the comparison does not by itself exclude geometry-correlated tracking systematics. On the direct head-to-head, the Anderson empirical model attains the smallest BIC (27.6 versus 32.4 for TEP restricted) — a $\Delta\text{BIC} \approx 4.8$ gap that lies in the weak-evidence regime and constitutes statistical parity rather than a decisive preference. Two context points matter. First, the Anderson tier is fitted on the very catalogue it is scored against (amplitude A and offset B), whereas TEP restricted carries a single amplitude on top of geometry pre-specified from independent measurements. Second, the ordering between the two tiers is systematic-convention dependent: under the published-uncertainties-only sensitivity block ($\sigma_{\text{sys}} = 0$), TEP restricted leads Anderson by $\Delta\text{BIC} \approx +2767$, because its geometry tracks the dominant NEAR amplitude while the linear asymmetry law does not; the former constant geometry-spread floor placed Anderson ahead by ≈ 98 only through its dependence on the arbitrary β_{ref} anchor, and the constant-floor sweep in Step 026 shows the ordering crossing zero near $\sigma_{\text{sys}} \approx 5$ mm/s. The headline fractional floor sits inside this range and yields parity. The TEP flexible tier (BIC = 68.4) does not improve on the restricted tier under the canonical linear response. The BIC-approximated Bayes factor ($BF \approx e^{\Delta\text{BIC}/2}$) is a large-sample result. For $n < 10$ and extreme signal-to-noise it is unreliable; when $\log_{10}(BF) > 100$ the value is driven by formal uncertainties and should not be read as a literal probability ratio (Step 026 flags all comparisons in this regime). The displayed scores are arithmetic functions of the archived objective values and parameter counts. Their probabilistic interpretation additionally requires a verified likelihood optimum and treatment of all data-estimated nuisance quantities; those checks are distinct from the channel theorem. Even under the pessimistic parameter count ($k_{\text{eff}} = 8$), the TEP restricted BIC = 47.8 still beats the Null by $\Delta\text{BIC} \approx 16674.1$, confirming that the catalog compression is not an artifact of aggressive parameter counting; under that penalty the Anderson tier's margin over TEP restricted widens to ≈ 20 BIC units.

Akaike weights (Step 026, full $n = 9$ catalog, AICc): Anderson ≈ 0.86 , TEP restricted ≈ 0.14 , TEP flexible ≈ 0 , Null ≈ 0 . The weight distribution quantifies the same parity-level ordering as the BIC comparison — neither tier dominates under the fractional floor.

The restricted model remains the scientifically important tier because every quantity except β_{fit} is pre-specified from independent measurements or stated response templates. Its large ΔBIC separation on the full $n = 9$ catalog therefore reflects predictive compression relative to the Null with a catalog-independent geometry, while the Anderson tier measures how much additional compression is purchasable with a second catalog-fitted parameter.

Artefact-Kernel Impulse Consistency Verification

The fitted β_{fit} values provide a direct probe of the clock-sector artefact structure through the kernel impulse diagnostic. The artefact-kernel impulse $\mathcal{I} = \int_{\text{path}} \mathbf{F}_\phi \cdot d\mathbf{r}$ measures the net accumulation of the clock-sector response along the flyby trajectory. In the TEP framework, the predicted apparent velocity shift relates to the impulse via $\Delta v_{\text{TEP}} \propto \beta_{A,\text{eff}} \cdot \mathcal{I}$, modulated by trajectory geometry and disformal coupling. The consistent mapping between fitted β_{fit} values and the geometric impulse computed from each flyby's 3D trajectory (using JPL Horizons ephemerides) supports the conclusion that the clock-sector response model respects the fundamental field structure of the TEP equations. On the corrected catalogue the impulse magnitude is not positively correlated with the fitted β_{fit} ($r \approx -0.4$ across the six-member ensemble): the flybys with near-zero reference impulses (Cassini, MESSENGER) return divergent rescalings that carry no amplitude information, while the well-conditioned members (NEAR, Galileo 1990, Rosetta 2005) span the tightly constrained range $\beta_{\text{fit}} \sim (4.4\text{--}14.4) \times 10^{-3}$. The diagnostic therefore confirms that the per-flyby rescaling is well-posed where the reference impulse is non-negligible, and that the ensemble amplitude is set by the members with genuine response leverage.

The fitted β_{fit} values are flyby trajectory-response amplitudes, not PPN γ predictions. Cassini PPN compliance is evaluated through the solar source-charge projection $\alpha_{\text{eff}}^{(\odot)} = S_\Sigma^{(\odot)} \alpha_0$ (Section 4.6.1a). The ensemble weighted mean yields $\beta_{\text{fit}} = 4.74 \times 10^{-3} \pm 5.83 \times 10^{-4}$. The solar source-charge projection in Section 4.6.1a adopts the Paper 0 v0.16 convention $S_\Sigma^{(\odot)} \sim 2 \times 10^{-6}$, inside the Cassini requirement $S_\Sigma^{(\odot)} \lesssim 1.15 \times 10^{-5}$ under the unified unstretched-space metric (5.8×10^{-6} on the isotropic-conformal benchmark); the two-branch kinetic-operator evaluation on the solar profile gives $y_{\text{tb}}(1.6 R_\odot) = 2.16 \times 10^{-8}$ ($|\gamma_{\text{PPN}} - 1| = 4.32 \times 10^{-8} \ll 2.3 \times 10^{-5}$) and $S_\Sigma \sim 10^{-15}$ at the 1 AU path. This supports Temporal Topology screening in the solar environment without overstating the margin.

4.6 PPN Constraints and Validation

4.6.1 PPN Constraint Derivation

The PPN (Parametrized Post-Newtonian) formalism characterizes deviations from General Relativity. For scalar-tensor theories with conformal coupling $A(\phi) = \exp(\beta_A \phi / M_{\text{Pl}})$, the PPN parameter γ relates to the coupling strength under the unified unstretched-space metric $\tilde{g}_{\mu\nu} = -A_{\text{loc}}^2(u_{\text{grav}}(\mathbf{x}))U_\mu U_\nu + A_{\text{sp}}^2(u_{\text{drift}}(t, \mathbf{x}))\Pi_{\mu\nu} + B(\phi)\Pi_\mu^\alpha \Pi_\nu^\beta \nabla_\alpha \phi \nabla_\beta \phi$ (Paper 0 §2.2, Step 80; with $4\beta_A^2 S_\Sigma^{(\odot)}$ on the isotropic-conformal benchmark):

$$|\gamma - 1| \approx 2\beta_A^2 S_\Sigma^{(\odot)} \quad \left(\text{isotropic benchmark: } 4\beta_A^2 S_\Sigma^{(\odot)}; \text{ for small } S_\Sigma \right) \quad (25)$$

Derivation (Paper 0, Sec. 7, Step 80): in the screened-source limit the photon probe is unscreened while the source charge is suppressed, so under the unified unstretched-space metric $\gamma_{\text{PPN}} - 1 = -2\beta_A^2 S_\Sigma / (1 + 2\beta_A^2 S_\Sigma)$ (or $-2\alpha_0 \alpha_{\text{eff}} / (1 + \alpha_0 \alpha_{\text{eff}}) = -4\beta_A^2 S_\Sigma / (1 + 2\beta_A^2 S_\Sigma)$ on the isotropic-conformal benchmark), linear in S_Σ , with $\alpha_{\text{eff}} = S_\Sigma \alpha_0$ and $\alpha_0 = \sqrt{2} \beta_A$ the DEF-normalized microscopic coupling. The Cassini source-charge parameter is $\alpha_{\text{eff}}^{(\odot)} = S_\Sigma^{(\odot)} \alpha_0$, giving $|\gamma - 1| \approx 2\beta_A^2 S_\Sigma^{(\odot)}$ ($4\beta_A^2 S_\Sigma^{(\odot)}$ isotropic) for magnitude comparisons to Cassini (the measured $\gamma - 1$ is negative in both conventions). The Earth-vicinity flyby response geometry $S_\oplus$ is a distinct projection and is not the Cassini source-charge parameter.

The fitted β_{fit} values are flyby trajectory-response amplitudes for Earth-vicinity geometry, distinct from the Cassini source-charge parameter $\alpha_{\text{eff}}^{(\odot)} = S_\Sigma^{(\odot)} \alpha_0$. The Earth-screened flyby response amplitudes $\beta_{\text{fit}} \times S_\oplus$ are:

- NEAR: $\beta_{\text{fit}} \times S_\oplus = 1.55 \times 10^{-3}$
- Galileo 1990: $\beta_{\text{fit}} \times S_\oplus = 5.03 \times 10^{-3}$
- Galileo 1992: $\beta_{\text{fit}} \times S_\oplus = 2.33 \times 10^{-2}$
- Rosetta 2005: $\beta_{\text{fit}} \times S_\oplus = 1.80 \times 10^{-3}$
- Cassini and MESSENGER: formally gated but unconstrained (reference prediction ≈ 0); Cassini PPN compliance is evaluated through the solar source-charge projection $\alpha_{\text{eff}}^{(\odot)} = S_\Sigma^{(\odot)} \alpha_0$ (Section 4.6.1a), not through the flyby amplitudes.

These are trajectory-response amplitudes, not PPN γ predictions. Cassini PPN compliance is evaluated through the solar source-charge projection $\alpha_{\text{eff}}^{(\odot)} = S_\Sigma^{(\odot)} \alpha_0$ (Section 4.6.1a), not through the fitted flyby amplitudes.

4.6.1a Solar-Screening PPN Check for Cassini

The Cassini Shapiro-delay measurement constrains the scalar field along the radio path during solar conjunction, not at Earth's surface. Cassini PPN compliance is evaluated through the solar source-charge projection, not through the fitted flyby amplitudes. Following the Jakarta mapping (Paper 0, §7):

$$\alpha_{\text{eff}}^{(\odot)} = S_{\Sigma}^{(\odot)} \alpha_0, \quad \gamma_{\text{PPN}} - 1 = -\frac{2\beta_A^2 S_{\Sigma}^{(\odot)}}{1 + 2\beta_A^2 S_{\Sigma}^{(\odot)}} \simeq -2\beta_A^2 S_{\Sigma}^{(\odot)} \quad \left(\text{isotropic benchmark: } -4\beta_A^2 S_{\Sigma}^{(\odot)}\right), \quad (26)$$

where $\alpha_0 = \sqrt{2}\beta_A$ is the DEF-normalized bare microscopic coupling (linear in the source charge because the photon probe is unscreened) and $S_{\Sigma}^{(\odot)}$ is the solar Temporal-Topology screening factor. With $\beta_A = -1$ (Jakarta §2.2), the Cassini bound $|\gamma - 1| < 2.3 \times 10^{-5}$ requires (Paper 0, Jakarta v0.16 §7, Step 80):

$$S_{\Sigma}^{(\odot)} \lesssim \frac{2.3 \times 10^{-5}}{2\beta_A^2} \approx 1.15 \times 10^{-5} \quad \left(\text{isotropic benchmark: } \frac{2.3 \times 10^{-5}}{4\beta_A^2} \approx 5.8 \times 10^{-6}\right). \quad (27)$$

Cassini constrains the solar exterior source-charge projection $S_{\Sigma}^{(\odot)}$, not the Earth-vicinity flyby response $S_{\oplus}$. Paper 0 v0.16 adopts the Solar-System source-charge convention $S_{\Sigma}^{(\odot)} \sim 2 \times 10^{-6}$, which satisfies the Cassini requirement $S_{\Sigma}^{(\odot)} \lesssim 1.15 \times 10^{-5}$ (5.8×10^{-6} isotropic) comfortably; the two-branch kinetic-operator evaluation on the solar field returns $y_{\text{tb}}(1.6 R_{\odot}) = 2.16 \times 10^{-8}$ ($|\gamma_{\text{PPN}} - 1| = 4.32 \times 10^{-8} \ll 2.3 \times 10^{-5}$) and $S_{\Sigma} \sim 1.6 \times 10^{-15}$ at the 1 AU path (Paper 0 Appendix E, static-branch benchmarks), with the microscopic uniqueness of that charge within the wider $K(X)$ class remaining open. The flyby response and the solar PPN source charge are therefore treated as distinct environmental projections of the same scalar configuration.

4.6.2 Sensitivity Analysis

To assess robustness, the TEP model is tested against variations in key parameters. Table 3f shows how results change when parameters are varied within physically plausible ranges:

Table 3f: Sensitivity Analysis - Parameter Variations

Parameter	Nominal Value	Tested Range	Fit stable?	Impact on β
Geometric suppression factor ($S_{\oplus}$)	0.35	0.30 – 0.50	✓ Yes (all values)	±6%
Relaxation length (λ_T)	4200 km	2000 – 8000 km	✓ Yes (within range)	≤50% (sensitivity coefficient 0.36)
J2 coefficient	1.08263×10^{-3}	±0.1%	✓ Yes	<1%
J3 coefficient	-2.54×10^{-6}	±10%	✓ Yes	negligible
Trajectory uncertainty	declination ±0.5°	±1°	✓ Yes	±5%

Robustness conclusion: The TEP model maintains fit stability across a broad range of parameter values. The phase-boundary factor can vary by ±32% (0.25 to 0.45) and all fitted β_{fit} values remain within the fitted range. This suggests that the fit stability is not fine-tuned but is a feature of the screening mechanism. The relaxation length has moderate impact on predicted Δv but does not affect fit stability because the fitted β values adjust to compensate.

4.6.3 OD Filter Simulation: Suppression Hypothesis Validation

Step 021 now withholds all mission-specific OD survival factors until real mission OD configuration data are available. The earlier synthetic Step 012 batch least-squares experiment is retained only as a diagnostic development artifact: its empirical-acceleration implementation is numerically unstable in the current 3D form, and the generated result is not valid for computing F_{OD} or for supporting quantitative claims about modern OD suppression.

Current OD Evidence Status

- Mission F_{OD} values: not computed; mission OD configuration files are required.
- Step 012 synthetic OD run: quarantined as `synthetic_diagnostic_not_for_manuscript_inference`.

- Manuscript policy: no era-based or synthetic OD survival factors are used in the Step 039 classification table.

Interpretation: The OD-suppression mechanism remains a falsifiable hypothesis rather than a calibrated correction. It is physically plausible that empirical acceleration states and residual editing can absorb unmodeled perigee-local forces, but this paper does not assign numerical survival fractions without mission-specific OD settings.

Table 3g: OD Survival-Factor Status

Quantity	Status	Use in likelihood?
Mission-specific F_{OD}	Not computable without OD configuration files	No
Step 012 synthetic OD diagnostic	Quarantined; not valid for manuscript inference	No

Connection to observations: Step 039 classifies fixed-amplitude raw predictions (Step 008 pooled β_{fit}) with the Step 007 geometry envelope (3 deterministic true positives, 4 deterministic true nulls, 1 sign-consistent raw surplus, and 1 deterministic fixed-amplitude warning case across the 9 published rows). After propagating Step 008 random-effects amplitude scatter, the uncertainty-aware raw layer retains 1 raw-tension case (Juno) and moves Galileo 1992 to "observed detection, prediction uncertain". Post-OD survival factors are withheld until mission OD configuration yields defensible F_{OD} estimates.

Juno tension: The Juno non-detection (universal- β_{fit} prediction $+0.12 \pm 0.03$ mm/s, observed 0.00 ± 0.02 mm/s) is the most serious raw-tension case and motivates independent raw DSN re-analysis. In the current pipeline run, the DSN ingestion layer emits mission discovery and request artifacts but does not yet provide indexed raw Doppler products for most missions (ingest status `no_indexed_products`), so this re-analysis remains a defined falsification pathway rather than a completed reprocessing result.

4.6.3a Clock-Sector Artefact: End-to-End Orbit-Determination Rederivation

The clock-sector mechanism has now been propagated through the full reduction chain rather than asserted (Step 043). For each of the twelve flybys with reconstructed JPL Horizons state-vector arcs, the canonical field excursion $\psi(r) = \phi_{space} - \phi(r)$ of the pipeline's screened field solution is evaluated along the real trajectory; the matter-metric rate offset $\eta(t) = \exp(\beta_A \psi / M_{Pl}) - 1$ ($\beta_A = -1$; the unsuppressed amplitude channel S_A , not the screened shear channel) accumulates the proper-time offset $\delta\tau(t) = \int \eta dt$ and injects the apparent range-rate offset $c\eta(t)$ that the excursion produces in a clock-carrying link — a spacecraft-local oscillator term entering once wherever the downlink is referenced to an onboard time standard (one-way, non-coherent, or regenerative-ranging observables). In the fully coherent two-way class of the published catalogue the term is absent at the observable: the transponder multiplies the received carrier by a fixed ratio, and the endpoint conformal factors cancel identically (Section 5.2). The step therefore computes what the clock excursion reconstructs as on a clock-carrying link over the real arcs — the mechanism's falsifiable prediction — together with the postfit-residual signature that quantifies how much of it survives reduction into the orbit solve. The identical batch least-squares orbit-determination filter used in the OD-suppression study (six-state estimator, point-mass dynamics, assumed GR time standard) is then run on clean and corrupted synthetic Doppler, and the differential of the fitted end-minus-start speed change isolates the clock-induced apparent ΔV . Truth dynamics are pure point-mass gravity — consistent with the screened shear channel contributing nothing to the trajectory — while $\eta(t)$ is sampled on the real arcs, preserving the physical in/out asymmetry.

Table 3h: Clock-sector OD rederivation — per-flyby proper-time signature, reconstructed apparent velocity shift, and postfit residual signature on a clock-carrying link (Step 043; noiseless differential estimator, three-complex DSN network, ± 12 h tracking window at 180 s cadence). ΔV_{app} is an absorption remnant — it shrinks as the arc lengthens (NEAR: $+0.68/+0.18/+0.05$ mm/s for $\pm 6/12/24$ h) — so the discriminating observable is the unabsorbed postfit residual column. The injected term is absent from the catalogue's coherent two-way class (Section 5.2); the comparison column therefore evaluates sign and scale, not attribution.

Flyby	η at perigee	$\delta\tau$ total (μ s)	ΔV_{app} (mm/s)	Postfit resid. RMS (peak) mm/s	Observed (mm/s)
NEAR	-1.22×10^{-9}	-2.0	+0.18	5.8 (35)	13.46 ± 0.13
Galileo 1990	-1.11×10^{-9}	-1.6	+0.07	7.7 (83)	3.92 ± 0.08
Galileo 1992	-1.29×10^{-9}	-1.8	+0.06	4.6 (42)	-4.60 ± 1.00
Cassini	-1.05×10^{-9}	-1.0	+0.01	1.3 (13)	-2.00 ± 1.00
Rosetta 2005	-8.73×10^{-10}	-2.0	+0.50	4.7 (27)	1.80 ± 0.03
Rosetta 2007	-3.93×10^{-10}	-0.7	+0.01	0.7 (5)	0.02 ± 0.05
Rosetta 2009	-7.70×10^{-10}	-1.2	+0.05	3.0 (26)	0.00 ± 0.05
MESSENGER	-7.97×10^{-10}	-1.9	+0.49	4.7 (31)	0.02 ± 0.01
Juno	-1.22×10^{-9}	-1.6	+0.04	1.8 (15)	0.0 ± 0.02

Flyby	η at perigee	$\delta\tau$ total (μ s)	ΔV_{app} (mm/s)	Postfit resid. RMS (peak) mm/s	Observed (mm/s)
Stardust	-3.33×10^{-10}	-0.8	+0.04	0.5 (3)	—
OSIRIS-REx	-2.31×10^{-11}	-0.1	+0.02	0.3 (1)	—
BepiColombo	-6.78×10^{-11}	-0.3	+0.05	0.2 (0.7)	—

Three findings follow, all conditional on the link class. First, the amplitude regime is correct: the canonical field excursion ($\Delta\phi = 3.4 \times 10^9$ GeV, $\eta_{\text{surface}} = -1.39 \times 10^{-9}$) produces accumulated proper-time offsets of 0.1–2 μ s; the raw signature reaches $c\eta \sim$ m/s at perigee, so the excursion would be an unmissable feature of any one-way, non-coherent, or onboard-referenced flyby measurement. Second, the discriminating observable is the postfit residual, not the reconstructed endpoint ΔV : the six-state fit absorbs the bulk of the perigee-concentrated offset into the estimated state, but a structured remnant always survives — 0.2–7.7 mm/s RMS across the catalogue (peaks up to ~ 83 mm/s), i.e. 4–25% of the injected signature leaks through orbit determination at ~ 180 s cadence and remains detectable at DSN noise floors. The endpoint ΔV_{app} column is reported for completeness but is an absorption remnant rather than a signal amplitude — it shrinks by a factor of ~ 15 as the tracking window widens from ± 6 to ± 24 h (NEAR: $+0.68 \rightarrow +0.05$ mm/s) — while the residual signature is window-stable (the linearized filter response reproduces the nonlinear result within 10–40%). Third, the recovered artefact is controlled by the projection of the $\eta(t)$ profile through the tracking geometry into the fit basis, not by the leg-integrated proper-time asymmetry: across the nine published rows the reconstructed apparent shift correlates only moderately with the catalogue (Pearson $r = 0.17$, Spearman $\rho = 0.58$), so the artefact is intrinsically geometry-projection-dependent, consistent with the observed anomaly's documented resistance to simple energy-change parameterizations.

Amplitude calibration: evaluated against the detection subset, the raw canonical normalization under-produces the detected anomalies, and the global clock-response coefficient — the clock-channel analogue of the Step 008 response amplitude β_{fit} — returns $C_A = 4.2$ on the sign-inclusive $> 3\sigma$ detection subset (NEAR, Galileo 1990, Galileo 1992, Rosetta 2005; $n = 4$, Pearson $r = 0.08$ at this normalization) with $\chi^2 = 1.17 \times 10^4$ on 3 dof: no single global clock-channel scaling describes the detections, so the coefficient is reported rather than adopted. The per-flyby ordering of ΔV_{app} does not, by itself, reproduce the catalog pattern: the same geometry envelope and response-amplitude structure fitted in §4.3–4.5 remains load-bearing, and the clock-sector rederivation supplies the physical channel through which that response acts in a clock-carrying link rather than a replacement fit. Because the catalogued anomalies were recorded in the coherent two-way class (Section 5.2), the coefficients calibrated here characterize the forward model's response scale for clock-carrying observables, not an attribution of the published values. Limitations are explicit: a generic three-complex network rather than each mission's historical tracking schedule, point-mass truth dynamics, and single-channel (clock-only) injection. The per-mission tracking-schedule reconstruction and joint clock-plus-environment injection are identified refinements.

4.6.4 Leave-One-Out Cross-Validation

To verify that the weighted mean β is not dominated by any single detection, the analysis is repeated excluding each flyby successively:

Table 3i: Leave-One-Out Cross-Validation Results

Excluded Flyby	β without this flyby	Change from full sample
None (full sample)	4.74×10^{-3}	—
NEAR (1998)	5.88×10^{-3}	+24%
Galileo (1990)	4.58×10^{-3}	-3.5%
Rosetta (2005)	4.64×10^{-3}	-2.1%
Cassini (1999)	4.74×10^{-3}	<0.1%
Galileo (1992)	4.74×10^{-3}	<0.1%
MESSENGER (2005)	4.74×10^{-3}	<0.1%

The stability coefficient (relative standard deviation of LOO estimates divided by their mean) is 0.092, indicating moderate robustness (values < 0.5 are considered robust). Even when the high-S/N NEAR detection is excluded, the remaining five flybys yield $\beta = 5.88 \times 10^{-3}$, within the 95% bootstrap interval. This indicates that the TEP conclusion does not depend on any single detection.

4.6.5 Enhanced Statistical Validation

Artefact-kernel impulse consistency: The clock-sector response model's apparent velocity predictions integrate the field excursion along 3D trajectories while preserving the TEP metric structure. For each flyby, the predicted Δv_{TEP} is computed via path

integration of the artefact kernel $\mathbf{F}_\phi = \beta_{A,\text{eff}} c^2 \nabla \phi / M_{\text{Pl}}$ — the OD-equivalent impulse of the proper-time rate-offset excursion, not a physical force — along the actual spacecraft trajectory from JPL Horizons ephemeris. The open-path impulse $\mathcal{I} = \int \mathbf{F}_\phi \cdot d\mathbf{r}$ is consistently mapped to observable velocity shifts. This geometric consistency check distinguishes TEP from phenomenological force laws that lack field-theoretic structure.

Effect size analysis: Cohen's d compares each detection to the null-result population mean, using the pooled standard deviation of the two groups. The null population comprises the three published sub-threshold flybys (Rosetta 2007, Rosetta 2009, Juno) with mean $\Delta v = 0.007 \pm 0.012$ mm/s. The detection population ($n = 6$) has mean $\Delta v = 2.10 \pm 6.30$ mm/s. The pooled standard deviation is $\sigma_{\text{pooled}} = 5.33$ mm/s. Cohen's d for each detection vs. the null population:

- NEAR: $d = +2.53$ — large effect ($d \gg 0.8$)
- Galileo 1990: $d = +0.73$ — medium effect
- Galileo 1992: $d = -0.87$ — large effect on the negative branch
- Cassini: $d = -0.38$ — small effect on the negative branch
- Rosetta 2005: $d = +0.34$ — small effect
- MESSENGER 2005: $d = +0.003$ — negligible effect

NEAR shows the largest effect, and the corrected catalogue places Galileo 1992 on the negative branch with a large $|d|$ — consistent with the negative sign the TEP kernel predicts for both it and Cassini. The two strongest detections (NEAR and Galileo 1990) provide the bulk of the positive-branch statistical separation.

Bayesian model comparison: Stable four-tier model comparison (Step 026) on the full $n = 9$ catalog favors TEP restricted over Null ($\Delta\text{BIC} \approx 16689.5$) under the fractional misspecification floor, while the Anderson empirical model — whose two parameters are fitted on this same catalogue — attains the nominally smallest BIC at parity-level separation (TEP restricted minus Anderson $\Delta\text{BIC} \approx -4.8$; the ordering reverses in TEP restricted's favour under the published-uncertainties-only sensitivity). See Section 4.5 for tier definitions. The physics-based single-amplitude model therefore compresses the full catalog far beyond the null, with the residual gap to the empirical tier quantifying the amplitude-level heterogeneity documented in Section 4.2.

Prediction accuracy: On the Step 008 primary comparison, $R^2 \approx 0.850$, Pearson $r \approx 0.930$, MAE ≈ 1.67 mm/s, RMSE ≈ 2.24 mm/s, and MAPE $\approx 62.5\%$. NEAR dominates the variance fraction; small-anomaly rows inflate percentage errors.

Residual analysis: Shapiro–Wilk normality on the gated prediction residuals gives $p \approx 0.95$ (consistent with Gaussian tails at $n = 6$).

4.6.6 Characteristic Suppression from UCD Saturation Model

The characteristic suppression $S_\oplus \approx 0.35$ —critical to the flyby amplitude and the magnitude of the flyby anomaly—is derived from the UCD saturation model in Step 010. The derivation uses Earth's total mass and the saturation scale $\rho_T = 20$ g/cm³, yielding a transition radius $R_{\text{sol}} \approx 4146$ km and suppression factor $S_\oplus = (R_\oplus - R_{\text{sol}})/R_\oplus \approx 0.35$. Under the corpus's terrestrial identification $L_c \sim \mathcal{O}(1) R_T(M_\oplus)$ (Paper 6 §2 transfer sketch), the GNSS covariance scale returns $S_\oplus = 0.34$ — a shared-input consistency check at the 2% level, since the same measurement anchors ρ_T , rather than an independent validation. The independent empirical checks — the flyby altitude threshold and the astrophysical cross-scale consistency of ρ_T across dwarf-galaxy cores and cluster halos (Paper 6) — converge on $S_\oplus \in [0.34, 0.39]$. See Step 010 for the complete derivation and cross-scale consistency arguments.

Distinction from the canonical amplitude operator: The EFA uses $S_\oplus = (R_\oplus - R_{\text{sol}})/R_\oplus$ as the field-excursion response ratio at the surface. Its relation to the corpus's canonical clock-amplitude operator $S_A(\mathcal{E}) = \min[1, (\bar{\rho}/\rho_T)^{1/3}]$ (Paper 0/Paper 26) is an exact identity, not a second estimate: under the uniform-density-equivalent construction, $S_A(\bar{\rho}_\oplus) = (\bar{\rho}_\oplus/\rho_T)^{1/3} = R_{\text{sol}}/R_\oplus \approx 0.65$ — the saturated (embedded) fraction of the equivalent radius, coinciding with the UCD embedding factor used in Paper 6 (UCD). The flyby clock channel projects the complementary responsive-shell fraction $S_\oplus = 1 - S_A \approx 0.35$, because a transported clock's differential excursion relative to the ground accumulates only across the unsaturated shell where the field relaxes; the pinned interior contributes no radial excursion. The Step 010 output records the machine-checked partition sum $S_A + S_\oplus = 1$ (`canonical_amplitude_partition`). For strongly centrally condensed objects the enclosed mean density exceeds ρ_T , in which case S_A saturates at unity and the excursion projection clamps at $S_\oplus \rightarrow 0$ (a fully pinned profile) rather than going negative — the uniform-density-equivalent partition applies to bodies whose enclosed mean density remains below ρ_T .

4.6.7 Systematic Uncertainty Budget

A comprehensive uncertainty budget quantifies the contribution of each uncertainty source to the fitted β_{fit} parameters. The corrected uncertainty analysis (Step 025) distinguishes between variance contributions and total relative uncertainty:

Variance Contributions:

- Statistical: 0.8%
- Systematic: 3.4%

- Heterogeneity: 95.8%

Total Relative Uncertainty:

- Statistical: 14.5%
- Systematic: 29.2%
- Heterogeneity: 155.4%
- Total: 158.8%

Systematic Breakdown:

- Trajectory reconstruction: 1.0%
- Characteristic suppression (UCD): 25.0% (from $\rho_T = 20 \pm 7 \text{ g/cm}^3$, Paper 6) ← DOMINANT
- Multipole coefficients: 0.1%
- Relaxation length (UCD): 15.0% (SCF theoretical prior)

Interpretation: The total relative uncertainty of 152.7% is dominated by heterogeneity (149.4%), which reflects genuine geometry-dependent physical variation in the effective coupling across flybys. This is expected in the TEP framework where $\beta_{A,\text{eff}}$ varies with altitude, latitude, velocity, plasma environment, and trajectory asymmetry (the plasma term is a bounded heuristic envelope coefficient, Section 3.10). The systematic uncertainty (29.2%) is dominated by characteristic suppression uncertainty (25.0%) from the UCD saturation model ($\rho_T = 20 \pm 7 \text{ g/cm}^3$, Paper 6), with relaxation length uncertainty (15.0%) from the SCF theoretical prior as the second-largest source.

4.7 Model Predictions for All Flybys

Table 4 presents the full prediction set evaluated at the universal weighted-mean coupling constant ($\beta_{\text{fit}} = 4.74 \times 10^{-3}$), scaled from the reference predictions ($\beta_{\text{fit},0} = 10^{-4}$) via the linear response established in Step 008 (the canonical excursion normalization is β -independent, so $\Delta v \propto \beta_{\text{fit}}$). Each row reports the raw TEP prediction at that universal coupling, the universal- β_{fit} residual $\Delta v_{\text{obs}} - \Delta v_{\text{TEP}}^{\text{raw}}$, and the raw classification from Step 039. Mission-specific OD survival factors F_{OD} and post-OD predictions are emitted only when Step 021 supplies defensible mission OD configuration data; otherwise those columns are withheld. The raw classification uses a 3σ detection threshold relative to the published uncertainty.

Table 4: Per-Flyby Universal- β_{fit} Predictions and Classification (Step 039)

Spacecraft	Data class	Alt. (km)	$\cos \delta_{\text{in}} - \cos \delta_{\text{out}}$	Δv_{obs} (mm/s)	$\Delta v_{\text{TEP}}^{\text{raw}}$ (mm/s)	Residual (mm/s)	F_{OD}	$\Delta v_{\text{TEP}}^{\text{post-OD}}$ (mm/s)	Classification
NEAR	Published anomaly	539.0	+0.625	$+13.46 \pm 0.13$	$+14.394 \pm 3.901$	−0.934	—	—	1
Galileo 1990	Published anomaly	960.0	+0.149	$+3.92 \pm 0.08$	$+1.294 \pm 0.351$	+2.626	—	—	1
Rosetta 2005	Published anomaly	1955.0	+0.173	$+1.80 \pm 0.03$	$+1.657 \pm 0.449$	+0.143	—	—	1
Cassini	Published anomaly	1175.0	−0.022	$−2.00 \pm 1.00$	$−0.001 \pm 0.000$	−1.999	—	—	1
Galileo 1992	Published anomaly	303.0	−0.170	$−4.60 \pm 1.00$	$−0.328 \pm 0.089$	−4.272	—	—	F
MESSENGER	Published anomaly	2347.0	≈ 0	$+0.02 \pm 0.01$	0.000 ± 0.000	+0.020	—	—	1
Rosetta 2009	Published null/bound	2572.0	+0.037	0.00 ± 0.05	$+0.011 \pm 0.003$	−0.011	—	—	1
Juno	Published null/bound	817.4	+0.069	0.00 ± 0.02	$+0.124 \pm 0.034$	−0.124	—	—	F
Rosetta 2007	Published null/bound	5301.0	+0.034	$+0.02 \pm 0.05$	$+0.001 \pm 0.000$	+0.019	—	—	1
Stardust	No public anomaly report	6008.9	−0.053	—	$−0.003 \pm 0.001$	—	—	—	U

Spacecraft	Data class	Alt. (km)	$\cos \delta_{\text{in}} - \cos \delta_{\text{out}}$	Δv_{obs} (mm/s)	$\Delta v_{\text{TEP}}^{\text{raw}}$ (mm/s)	Residual (mm/s)	F_{OD}	$\Delta v_{\text{TEP}}^{\text{post-OD}}$ (mm/s)	
OSIRIS-REx	No public anomaly report	17239.1	+0.150	—	0.000 ± 0.000	—	—	—	
BepiColombo	No public anomaly report	12697.3	−0.034	—	$−0.000 \pm 0.000$	—	—	—	

Summary: At the universal weighted-mean β_{fit} , Step 039 classifies three published anomalies as raw true positives (NEAR, Galileo 1990, Rosetta 2005). Cassini is a raw true null (corrected -2 ± 1 mm/s entry versus a near-zero prediction), Galileo 1992 is a raw surplus (correct sign, underpredicted magnitude at fixed universal amplitude), MESSENGER and the three published null/bound cases are consistent with universal- β_{fit} predictions under the Step 007 geometry envelope, and one raw-tension case remains (Juno). Stardust, OSIRIS-REx, and BepiColombo lack public anomaly reports and are not used in quantitative likelihood.

Falsifiability criterion: Raw-tension cases define model stress tests independent of era-based OD survival factors. Juno remains the sole deterministic fixed-amplitude warning case at the refit weighted-mean β_{fit} ($+0.12$ mm/s raw prediction vs. 0.00 ± 0.02 mm/s observed) with random-effects prediction uncertainty ± 0.03 mm/s. Even after propagating Step 008 random-effects scatter, the uncertainty-aware Step 039 layer retains Juno as a raw-tension case ($\sim 3\sigma$), while Galileo 1992 moves to the uncertainty-aware class "observed detection, prediction uncertain" — the model assigns the correct sign but insufficient magnitude at the pooled amplitude. Post-OD false-negative counts are reported only when mission-specific F_{OD} values are available from Step 021; with current OD configuration data, those columns are withheld.

Honest assessment: The clock-sector response model reproduces the three largest published anomalies at universal β_{fit} with sub-millimetre to millimetre residuals, while high-altitude or symmetric trajectories remain consistent with null predictions. Galileo 1992 — now correctly catalogued as a -4.60 ± 1.00 mm/s published detection — and Juno define the priority raw DSN reanalysis targets. Cassini is no longer a sign-tension case: the corrected -2 ± 1 mm/s entry carries the same sign as the TEP prediction, leaving only an amplitude shortfall that requires independent DSN OD or envelope refinements, not only a universal rescaling of β_{fit} .

4.8 Heterogeneity and Robustness Analysis

Heterogeneity assessment: The six gated fitted β_{fit} values span 4.43×10^{-3} to 3.40×10^1 (Step 009), the upper end driven by the two unconstrained near-zero-prediction rows; the four amplitude-informative members span 4.43×10^{-3} to 6.65×10^{-2} . Step 009 does not perform a formal variance decomposition because $n = 6$ remains underpowered for ANOVA-style partitioning. Instead, Step 009 reports three complementary deterministic geometry modulation analyses: (1) beta scatter statistics (log STD = 1.77 dex across all 6 rows, or a factor of 3.24 across the three clean detections NEAR, Galileo 1990, and Rosetta 2005); (2) full-catalog detection-pattern classification at the Step 008 pooled coupling $\beta_{\text{pooled}} = 4.74 \times 10^{-3}$ (91.7% accuracy, 11/12 at the 0.5 mm/s threshold: 3 true positives, 8 true negatives, 1 false negative, 0 false positives; and 100% accuracy, 12/12: 4 true positives, 8 true negatives at the 0.3 mm/s threshold); and (3) Spearman rank correlation between predicted and observed anomaly magnitudes ($\rho = 0.56$, $p = 0.059$, $n = 12$) together with the asymmetry-ordering diagnostic ($\rho = 0.56$, $p = 0.11$, $n = 9$). The formal heterogeneity statistics on the gated ensemble are dominated by sub-percent measurement precision:

Table 5: Heterogeneity Statistics

Statistic	Value	Interpretation
Cochran's Q	$\approx 1.18 \times 10^3$	Large (expected: ~ 5 for 5 d.o.f.)
Degrees of freedom	5	$n - 1$ for $n = 6$ gated fits
Reduced χ^2	≈ 235.7	$\gg 1$ (scatter exceeds measurement noise)
I^2	$\approx 99.58\%$	Formally extreme ($I^2 > 75\%$)
β_{fit} range (gated)	$4.43 \times 10^{-3} - 3.40 \times 10^1$	Includes two unconstrained near-zero-prediction rows; the four amplitude-informative members span a factor ≈ 15
Random-effects scatter	$\tau \approx 2.23 \times 10^{-3}$	Geometry- and environment-dependent modulation across the ensemble

The elevated I^2 on the six gated β_{fit} fits reflects formal tension between sub-percent per-flyby uncertainties and a multiplicative spread in fitted coupling, amplified by the two formally unconstrained rows. The metric is designed for meta-analyses where true effects are identical; here, geometry, velocity, and heuristic envelope structure intentionally modulate $\beta_{A,\text{eff}}$, so large I^2 is expected until additional physics is folded into a single generative curve or n grows.

Bootstrap resampling: To assess uncertainty given the small sample ($n = 6$ gated detections), parametric bootstrap resampling with 10 000 iterations is performed in Step 008:

- *Bootstrap median*: $\beta_{\text{fit}} \approx 4.77 \times 10^{-3}$ (tracks the inverse-variance weighted mean)
- *68% interval*: $[4.54 \times 10^{-3}, 5.93 \times 10^{-3}]$
- *95% confidence interval*: $[4.42 \times 10^{-3}, 1.47 \times 10^{-2}]$

The bootstrap median reproduces the weighted mean while the interval is widened by Galileo 1990's higher per-flyby β_{fit} ; the occasional resampled divergent row inflates the bootstrap mean, so the median and percentile intervals are the robust summaries.

Leave-one-out (Step 008): Inverse-variance β_{fit} is recomputed excluding each gated detection:

- *Exclude NEAR*: $\beta_{\text{fit}} \approx 5.88 \times 10^{-3}$ (+24%)
- *Exclude Galileo 1990*: $\beta_{\text{fit}} \approx 4.58 \times 10^{-3}$ (-3%)
- *Exclude Rosetta 2005*: $\beta_{\text{fit}} \approx 4.64 \times 10^{-3}$ (-2%)
- *Exclude Cassini, Galileo 1992, or MESSENGER*: $\beta_{\text{fit}} \approx 4.74 \times 10^{-3}$ (<0.1%)

The stability coefficient is 0.092, indicating moderate robustness (values < 0.5 are treated as acceptable in Step 008). NEAR is the dominant lever on the pooled scale; the three near-zero-prediction rows are negligible levers, as expected from their formal uncertainties.

Effect size: Cohen's d compares each detection to the null-result population using the pooled standard deviation of the two groups:

$$d = \frac{\Delta v_{\text{det}} - \mu_{\text{null}}}{\sigma_{\text{pooled}}}, \quad \sigma_{\text{pooled}} = \sqrt{\frac{(n_{\text{det}} - 1)s_{\text{det}}^2 + (n_{\text{null}} - 1)s_{\text{null}}^2}{n_{\text{det}} + n_{\text{null}} - 2}} \quad (28)$$

The null population comprises the published flybys with sub-threshold anomalies ($n_{\text{null}} = 3$: Rosetta 2007, Rosetta 2009, Juno; $\mu_{\text{null}} = 0.007$ mm/s, $s_{\text{null}} = 0.012$ mm/s). The detection population ($n_{\text{det}} = 6$, $\mu_{\text{det}} = 2.10$ mm/s, $s_{\text{det}} = 6.30$ mm/s) yields $\sigma_{\text{pooled}} \approx 5.33$ mm/s. The resulting Cohen's d values are:

- NEAR: $d = +2.53$ (large effect)
- Galileo 1990: $d = +0.73$ (medium effect)
- Galileo 1992: $d = -0.87$ (large effect, negative branch)
- Cassini: $d = -0.38$ (small effect, negative branch)
- Rosetta 2005: $d = +0.34$ (small effect)
- MESSENGER 2005: $d = +0.003$ (negligible effect)

NEAR remains strongly distinguishable from the null population, and Galileo 1992 now appears as a large negative-branch effect — both consistent with the sign structure the TEP kernel assigns. Rosetta 2005 shows a small effect and MESSENGER a negligible one, reflecting their proximity to the null-population mean. The spread in d values is consistent with the spread in gated fitted β_{fit} , confirming geometry-dependent modulation rather than a perfectly universal effective coupling at fixed envelope.

4.9 Resolution of Beta Heterogeneity

The multiplicative spread in gated fitted β_{fit} values is partially summarized through a four-stage decomposition (Step 009). This unified analysis consolidates structural physics modulation, observational pipeline effects, environmental modulation, and statistical limitations into a coherent framework. The apparent scatter is not treated as pure noise: it is the object of the envelope construction. See Section 4.3 for the detailed variance decomposition analysis.

4.10 PPN Compliance and Global State

Model comparison: Stable four-tier model comparison (Step 026) compares the Null, Anderson empirical, TEP restricted, and TEP flexible models on the full $n = 9$ catalog. All non-null tiers compress the catalog by $\Delta\text{BIC} \gtrsim 1.66 \times 10^4$ against the Null (TEP restricted ≈ 16689.5 , TEP flexible ≈ 16653.6 , Anderson ≈ 16694.3), confirming that trajectory asymmetry carries genuine signal. The Anderson empirical model — fitted with two free parameters on this same catalogue — attains the nominally smallest BIC, at parity-level separation from TEP restricted ($\Delta\text{BIC} \approx 4.8$); that ordering and its systematic-convention dependence are discussed in Section 4.5.

Formal correlation analysis: Pearson correlation tests quantify relationships between the per-flyby effective geometry factor G_{eff} (Section 4.2.1) and the physical parameters, as computed in Step 015 on the six-member ensemble:

Table 6: Geometry-Factor Correlation Results (n = 6 gated ensemble members, Step 015; correlations are underpowered and should be interpreted cautiously)

Parameter	Pearson r	p-value	Interpretation
Perigee altitude	-0.13	0.82	No altitude trend detected
Velocity	+0.90	0.011	Moderate positive correlation; the only approach to significance
Trajectory asymmetry	-0.28	0.62	Weak (G_{eff} already incorporates asymmetry by construction)

With $n = 6$ the correlation analysis is exploratory: the velocity trend ($r = 0.90$, $p = 0.011$) is the only coefficient approaching significance and is consistent with velocity-modulated response in the Temporal Topology screening framework, but a six-point sample cannot establish it. A multiple regression of $\ln G_{\text{eff}}$ on altitude, velocity, and asymmetry reaches $R^2 = 0.97$ (Step 015), dominated by the two divergent low-prediction rows and reported for completeness only.

Prediction intervals: Uncertainty propagation yields prediction intervals for additional flybys:

- Representative $\beta = 4.74 \times 10^{-3} \pm 5.83 \times 10^{-4}$ (Step 008 weighted mean and formal uncertainty; random-effects uncertainty 1.28×10^{-3})
- 68% bootstrap interval (Step 008): $[1.86 \times 10^{-3}, 2.52 \times 10^{-3}]$
- 95% bootstrap interval (Step 008): $[1.81 \times 10^{-3}, 8.97 \times 10^{-3}]$

The prediction intervals bracket the amplitude-informative gated fitted β_{fit} values and illustrate residual width driven largely by Galileo 1990's high per-flyby coupling.

Sensitivity analysis: All model parameters show stable results across plausible variation ranges:

Table 7: Parameter Sensitivity

Parameter	Range Tested	Stability
Phase-boundary factor $\Delta R/R$	0.25 – 0.45	Stable (all results within fitted range)
Relaxation length λ_T	3000 – 5000 km	Stable (weak dependence)
J2 coefficient	1.0 – 1.1	Stable (J2 dominates)

Model adequacy tests: Shapiro–Wilk on the gated prediction residuals yields $p \approx 0.95$ in Step 008 ($n = 6$; consistent with normality, small- n caution). Heterogeneity diagnostics (Cochran Q, I^2) dominate the interpretation relative to classical normality tests.

The preceding sections have established that the TEP model reproduces the observed anomalies and satisfies PPN constraints. The following section tests a deeper prediction: that the *residual* discrepancy between observation and prediction should correlate with the geometry of velocity in the ambient-field rest frame, approximated by the CMB dipole frame.

4.11 Cosmographic Temporal Shear Modulation Analysis

A key prediction of the TEP framework is that temporal shear should depend on the total velocity of the Earth-Moon system relative to the ambient-field rest frame, not merely the spacecraft velocity relative to Earth. If the cosmic microwave background (CMB) dipole frame approximates this frame, the ~ 370 km/s bulk motion of the Solar System toward (RA, Dec) = (167.94°, −6.93°) would provide a cosmographic modulation of the velocity-activated response branch. Additionally, Earth's elliptical orbit produces a heliocentric distance-dependent modulation via solar scalar topology. Under the canonical $B(\phi)$ envelope the propagating-sector amplitude at flyby velocities is negligible (§3.5), so this section tests these predictions as template-level directional hypotheses rather than as signals of a detected disformal channel, using full three-dimensional spacecraft state vectors extracted from JPL Horizons archival ephemeris.

4.11.1 3D State Vector Extraction

Raw JPL Horizons ephemeris files were parsed for each flyby mission, extracting geocentric apparent right ascension, declination, range, and range-rate at 1-minute intervals. Cartesian position and velocity vectors were reconstructed in the J2000 equatorial frame and rotated to the ecliptic frame using the obliquity of the ecliptic $\varepsilon = 23.439281^\circ$. Perigee state vectors were identified by minimum geocentric range. All twelve catalogued flybys carry reconstructed 3D state vectors (Step 038 refuses declination-only fallbacks); the correlation sample of this section is the $n = 9$ subset that also carries a catalogued anomaly value or published bound — Stardust, OSIRIS-REx, and BepiColombo have no published anomaly and so do not enter the ratio tests. Earth heliocentric position and velocity were computed via a low-precision analytical ephemeris with proper elliptical orbit mechanics, yielding non-zero radial velocity components up to ± 0.5 km/s consistent with Earth's orbital eccentricity $e = 0.0167$.

4.11.2 Cosmographic Modulation Factors

For each flyby, three classes of modulation proxies were computed:

1. Heliocentric distance modulation: The solar scalar field density scales as r^{-2} , yielding a modulation proxy $M_{\odot} = 1/r_{\text{AU}}^2$.
2. Solar scalar wind factor: Earth's orbital speed relative to the Sun modulates the scalar wind experienced by the spacecraft, approximated as $v_{\text{orb}}/29.78 \text{ km/s}$.
3. CMB dipole projection: The total velocity of the spacecraft in the CMB rest frame is $v_{\text{total}} = v_{\text{CMB}} + v_{\text{Earth}} + v_{\text{sc}}$. The component along the CMB dipole direction n_{CMB} defines the modulation factor $M_{\text{CMB}} = (v_{\text{total}} \cdot n_{\text{CMB}}) / 369.82 \text{ km/s}$.

The TEP disformal coupling scales as v^2 in the ambient-field rest frame. The CMB-rest-frame disformal enhancement factor is $f_{\text{enh}} = |v_{\text{total}}|^2 / |v_{\text{sc}}|^2$, ranging from ~ 350 to ~ 1300 across the sample. (Under the canonical $B(\phi)$ envelope this channel carries negligible absolute amplitude at flyby velocities — §3.5 — so the modulation explored here is a template-level directional check, not a predicted propagating-sector signal.) Because the 370 km/s CMB bulk velocity is nearly constant, the dominant variation in the effective coupling comes from the *direction* of the spacecraft velocity relative to the CMB dipole, quantified by $\cos \theta_{\text{SC-CMB}} = (v_{\text{sc}} \cdot n_{\text{CMB}}) / |v_{\text{sc}}|$.

4.11.3 Results

Table 8: Cosmographic Modulation Parameters and Residual Ratios

Mission	r_{AU}	v_{rad} (km/s)	$\cos \theta_{\text{SC-CMB}}$	$v_{\text{SC,CMB}}$ (km/s)	f_{enh}	Both Aligned	Obs (mm/s)	Pred (mm/s)	Ratio
NEAR	0.984	+0.178	+0.244	+3.10	973	YES	13.46	0.304	44.33
Galileo 1990	0.985	-0.217	+0.261	+3.57	866	YES	3.92	0.027	143.65
Cassini	1.012	-0.342	-0.957	-18.07	324	no	-2.00	-0.000	188850.40
Galileo 1992	0.985	-0.227	+0.864	+12.12	861	YES	-4.60	-0.007	664.54
Rosetta 2005	0.992	+0.445	-0.573	-6.02	1243	no	1.80	0.035	51.49
Rosetta 2007	0.990	-0.404	+0.611	+7.58	1056	YES	0.02	0.000	1512.56
Rosetta 2009	0.990	-0.385	+0.796	+10.59	931	YES	0.00	0.000	0.00
MESSENGER 2005	1.015	-0.219	+0.004	+0.04	1148	no	0.02	0.000	339558.33
Juno	0.999	-0.506	-0.454	-6.74	645	no	0.00	0.003	0.00
Stardust	0.984	+0.115	+0.519	+5.35	1514	no	—	-0.000	—
OSIRIS-REx	1.004	-0.486	+0.877	+7.47	2015	no	—	0.000	—
BepiColombo	1.002	+0.500	+0.999	+7.59	2308	no	—	-0.000	—

4.11.4 Correlation Analysis

Pearson correlation tests were performed between the observed-to-predicted ratio and each cosmographic modulation factor ($n = 9$). With the corrected catalogue, several ratio correlations are formally significant — but the residual ratio is dominated by two divergent rows (Cassini and MESSENGER, whose reference predictions are ≈ 0 and whose ratios reach $\sim 4 \times 10^4$ and $\sim 7 \times 10^4$), so these correlations are outlier-driven and reported for transparency only. The strongest individual correlations are:

Table 9a: Individual Correlation between Residual Ratio and Cosmographic Modulation Factors (formal values; divergent-ratio caveat applies)

Modulation Factor	Pearson r	p-value	Interpretation
Ratio vs heliocentric distance	+0.894	0.001	Formally significant; outlier-dominated
Ratio vs Earth orbital speed	−0.894	0.001	Formally significant; outlier-dominated
Ratio vs heliocentric modulation	−0.889	0.001	Formally significant; outlier-dominated
Ratio vs Earth CMB projection	−0.875	0.002	Formally significant; outlier-dominated
Ratio vs SC-orbital alignment	−0.854	0.003	Formally significant; outlier-dominated

Modulation Factor	Pearson r	p-value	Interpretation
Ratio vs CMB modulation factor	−0.759	0.018	Formally significant; outlier-dominated
Ratio vs SC-CMB $\cos \theta$	−0.360	0.341	Not significant at n = 9
Ratio vs radial velocity	−0.136	0.727	Not significant at n = 9

4.11.5 Directional Consistency: The Both-Aligned Test

Within the fitted response template, the velocity-activated branch depends on the *total* CMB-frame velocity, which is the vector sum of the Earth's orbital velocity and the spacecraft velocity, both projected onto the CMB dipole direction. When *both* the spacecraft velocity and Earth's orbital velocity are aligned with the CMB dipole apex ($\cos \theta_{\text{SC-CMB}} > 0$ and $\mathbf{v}_{\text{Earth}} \cdot \mathbf{n}_{\text{CMB}} > 0$), the two velocity components add constructively in the ambient-field rest frame, boosting the template's velocity factor. When one or both are anti-aligned, the components partially cancel, suppressing the template factor. This is a directional hypothesis about the phenomenological envelope only: the CMB-rest-frame enhancement multiplies a canonical-branch amplitude that is negligible at flyby velocities (§3.5), so even the largest f_{enh} factors remain far below any propagating-sector signal.

This template hypothesis was tested by defining a binary "both-aligned" flag for each flyby, equal to 1 when both projections are positive and 0 otherwise. The correlation between this flag and the residual ratio is:

Both-aligned flag: Pearson $r = -0.57$, $p = 0.11$ ($n = 9$)
Mann-Whitney U (aligned > unaligned): $U = 7.5$, $p = 0.79$ (exact test)

On the corrected catalogue the both-aligned flag shows no positive association with the residual ratio — the formal sign is negative and the Mann–Whitney test is null ($p = 0.79$). NEAR and Galileo 1990 retain the largest non-divergent observed-to-predicted ratios in the sample, but the directional flag does not separate them. The test remains exploratory pending additional flybys with published anomalies and full 3D trajectory reconstructions.

4.11.6 Multivariate Geometric Regression

A multivariate ordinary least squares regression was fitted to test whether a linear combination of geometric alignment factors can explain the residual ratio:

$$\text{ratio} = b_0 + b_1 \cos \theta_{\text{SC-CMB}} + b_2 (\mathbf{v}_{\text{Earth}} \cdot \mathbf{n}_{\text{CMB}} / 30) + b_3 (\text{SC-orbital alignment}) + \varepsilon$$

The fitted coefficients are $b_0 = +3.28 \times 10^4$, $b_1 = +1.42 \times 10^4$, $b_2 = -3.76 \times 10^4$, $b_3 = -1.58 \times 10^4$. The model achieves $R^2 = 0.96$ (adjusted $R^2 = 0.94$) and reduces the residual standard deviation by 80.7%. These figures are reported for completeness only: the response variable contains divergent ratios from the two near-zero-prediction rows (Cassini $\approx 3.8 \times 10^4$, MESSENGER $\approx 6.8 \times 10^4$), so a four-parameter fit can score a high R^2 by tracking the two outlier points while leaving the physical content empty. The regression does not constitute evidence for a CMB-frame coupling at this sample size.

Table 9b lists observed-to-predicted ratios for the nine flybys with usable 3D vectors. The multivariate fit is reported for transparency; it should not be over-interpreted at this sample size.

Table 9b: Multivariate Geometric Regression Predictions

Mission	Observed Ratio	Predicted Ratio	Residual
NEAR	44.33	-12877.80	+12922.13
Galileo 1990	143.65	-15048.85	+15192.50
Cassini	188850.40	156100.72	+32749.68
Galileo 1992	664.54	-22015.75	+22680.30
Rosetta 2005	51.49	40751.58	-40700.09
Rosetta 2007	1512.56	28876.50	-27363.94
Rosetta 2009	0.00	5362.63	-5362.63
MESSENGER 2005	339558.33	344532.90	-4974.58
Juno	0.00	5143.36	-5143.36

4.11.7 Optimal Weighted Combination

The relative weighting of the spacecraft and Earth CMB projections was determined by scanning the coefficient w in the linear combination $E = \cos \theta_{\text{SC-CMB}} + w (\mathbf{v}_{\text{Earth}} \cdot \mathbf{n}_{\text{CMB}} / 30)$ and selecting the value that maximizes $|r(E, \text{ratio})|$. On the corrected catalogue the optimum lies at the grid boundary ($w = -2.0$ of the scanned range $[-2, +2]$), yielding:

Optimal combination: $E = \cos \theta_{\text{SC-CMB}} - 2.0 (\mathbf{v}_{\text{Earth}} \cdot \mathbf{n}_{\text{CMB}} / 30)$ (grid boundary)
Pearson $r = +0.96$, $p = 5.5 \times 10^{-5}$ ($n = 9$); permutation $p = 0.003$

The apparently strong optimum is not evidence of a calibrated coupling: the maximizing weight sits at the boundary of the scanned grid, and the correlation is driven by the divergent residual ratios flagged in Table 9a. The artifact's own sensitivity exclusion of the near-zero-prediction Rosetta 2007 row leaves the heliocentric-distance correlation essentially unchanged ($r = 0.89$, $p = 0.003$, $n = 8$), confirming that the signal is carried by the Cassini/MESSENGER divergent-ratio points rather than by a generic property of the set. The scan is reported as an exploratory diagnostic only.

4.11.8 Interpretation

The corrected catalogue changes the exploratory picture in two ways. First, several formal correlations are now nominally significant, but they are carried by the divergent residual ratios of Cassini and MESSENGER (predictions ≈ 0) rather than by a directional property of the detection population. Second, the directional tests proper are null or negative:

1. Both-aligned flag ($r = -0.57$, $p = 0.11$; Mann–Whitney $U = 7.5$, $p = 0.79$): the flag shows no positive association with the residual ratio; NEAR and Galileo 1990 retain the largest non-divergent ratios but are not separated by alignment.
2. Multivariate regression ($R^2 = 0.96$, adjusted $R^2 = 0.94$): the nominal fit is outlier-driven — four parameters track the two divergent-ratio rows — and is reported for transparency only.
3. Optimal weighted combination ($r = +0.96$, $p = 5.5 \times 10^{-5}$, permutation $p = 0.003$): the optimum lies at the scanned grid boundary and is likewise carried by the divergent rows; it does not define a calibrated CMB-frame coefficient.

Caveats: With $n = 9$ flybys, cosmographic modulation remains exploratory. Neither the both-aligned test nor the directional correlations support a CMB-rest-frame claim on the corrected catalogue, and the formally significant ratio correlations are artefacts of near-zero denominators rather than directional evidence. Additional flybys with published anomalies and full 3D trajectory reconstructions are required before elevating this sector above the core altitude–asymmetry law.

5. Discussion

5.1 Physical Interpretation: The Clock-Sector Artefact Mechanism

The TEP framework supplies a clock-sector reduction artefact for links that carry the spacecraft's own time reference, distinct from the empirical geometry template fitted to the historical flyby catalogue. In standard General Relativity, orbit determination reconstructs spacecraft velocities under an assumed GR proper-time standard. In TEP, the dynamical proper time field ϕ couples to the causal matter metric $\tilde{g}_{\mu\nu} = A^2(\phi)g_{\mu\nu} + B(\phi)\nabla_\mu\phi\nabla_\nu\phi$; under Paper 0's amplitude–shear split (v0.16), the Earth-vicinity Temporal Shear is strongly screened while the clock-amplitude channel S_A remains unsuppressed. A fast, highly asymmetric hyperbolic flyby therefore accumulates a rapid proper-time offset at perigee that — where the tracking link references the spacecraft's own time standard — the pre- and post-encounter fits misread as a ΔV discontinuity — a clock-sector artefact misread as ΔV without an unmodeled force. Which measurement classes carry the term is decided by the transponder physics analysed in Section 5.2, and that classification is a load-bearing part of the interpretation: in the fully coherent two-way class of the published catalogue the endpoint cancellation removes the term identically, so the mechanism's claim on the catalogue is a prediction for clock-carrying links rather than an attribution of the recorded anomalies.

Screening projection notice. Screening in TEP is represented at theory level by the environmental operator $S_\Sigma(E)$. Quantities such as p_T , $R_T(M)$, $S_\oplus(r)$, compactness Φ/c^2 , local stellar density, thermal epoch, coherence length, proximity, and boundary geometry are domain-specific projections of E , not independent screening mechanisms and not interchangeable universal thresholds.

Four key refinements to the model:

1. Clock-sector artefact mechanism: The apparent velocity anomaly arises from accumulated proper-time offsets through the unsuppressed amplitude channel S_A , a consequence of the universal conformal coupling established in the Jakarta axioms. The perigee-centred offset integral mimics a small shift in reconstructed GM ; its non-radial, geometry-dependent component produces the observable velocity discontinuity *in the measurement classes that carry the spacecraft's own time reference* (one-way, non-coherent, or regenerative-ranging links and onboard-clock telemetry); in the fully coherent two-way class the term is absent at the observable, per the channel analysis of Section 5.2.

2. TEP relaxation length: The terrestrial length scale is fixed by the UCD/GNSS calibration pathway: the GNSS coherence-length measurement underwrites the UCD-inferred ρ_T , which fixes the geometric transition scale $R_T(M_\oplus)$ used here. This shared-input dependency is disclosed explicitly to avoid circular claims of independent validation.
3. Trajectory asymmetry: The factor $\cos \delta_{\text{in}} - \cos \delta_{\text{out}}$ determines how asymmetrically the spacecraft samples Earth's oblate (J_2) field. This factor—taken from Anderson et al. (2008)—is the dominant source of inter-flyby variation.
4. Disformal coupling: The full TEP metric includes a disformal term $B(\phi)\partial_\mu\phi\partial_\nu\phi$ that produces velocity-dependent effects; within the fitted response template, the velocity-activated factor can reverse the reference prediction sign for high-velocity anti-aligned trajectories. (The canonical-action status of this sector is assessed in §3.5: no kilometre-per-second disformal transition exists under the admissible envelope, and the hierarchical fit returns b_{disf} consistent with zero, so the velocity template is a phenomenological descriptor rather than a detected propagating channel.) The Cassini sign evaluation is physical rather than conventional: Table 1a shows that an independent JPL Horizons state-vector reconstruction yields the same trajectory asymmetry $\cos \delta_{\text{in}} - \cos \delta_{\text{out}} = -0.022$ as the published Anderson et al. (2008) value, so the weakly negative geometry factor is a property of the actual trajectory, not a literature convention artifact — and on the re-transcribed catalogue the anomaly itself is negative (-2 ± 1 mm/s), so the kernel's negative prediction agrees in sign, with the residual tension confined to magnitude. The sign test therefore applies to the shear-channel reference prediction — a channel that is screened in the Earth vicinity and is not the claimed anomaly mechanism. The clock-sector channel, the claimed mechanism, returns the reconstruction in the link class that carries it: Step 043 gives $\Delta V_{\text{app}} = +0.01$ mm/s against the catalogued -2.00 ± 1.00 mm/s (Table 3h), classified by Section 5.2's channel accounting: the catalogued value was measured in the coherent two-way class, where the clock term is absent, so the comparison is a consistency check of the mechanism's sign and scale rather than an attribution of the recorded anomaly.

The physical interpretation is that a spacecraft traversing Earth's oblate gravitational field accumulates a non-radial proper-time offset through the field's amplitude excursion; the offset reaches the tracking record in the clock-carrying link classes enumerated in Section 5.2. For symmetric trajectories the non-radial contribution cancels, naturally explaining the pattern of detections and null results where the term is present. Fit stability of the gated fits is verified in Section 4.6.

5.2 Comparison with Other Proposed Explanations

Several alternative explanations for the flyby anomaly have been proposed in the literature. A systematic comparison is essential for assessing the relative merit of the TEP framework:

Standard physics systematic effects:

- *Atmospheric drag*: Independent first-principles simulation (Step 022) computes atmospheric density from a piecewise-exponential layer atmosphere and integrates the drag acceleration over each osculating hyperbolic arc at a generic ballistic coefficient ($C_d A/m = 2.2 \times 10^{-2}$ m²/kg). For NEAR (539 km altitude), the integrated contribution is 0.11 mm/s — 0.8% of the observed 13.46 mm/s anomaly and opposite in sign, since drag removes energy while the NEAR anomaly is an energy gain. Across the catalogue the contributions span 10^{-30} –6.4 mm/s, the upper end belonging to the single lowest-altitude pass: Galileo 1992 at 303 km returns ~ 6.4 mm/s — of order its published -4.60 mm/s anomaly and matching its energy-removing sign. Drag is therefore excluded ($\lesssim 1\%$ of the anomaly, or wrong-signed) for every flyby above ~ 500 km, but it is not order-of-magnitude excluded for the 303 km member. That exclusion stays honest on two grounds: the catalogued anomalies are post-fit residuals of orbit solutions that already model drag at low altitude, so the operative bound on a residual drag term is the drag-model error rather than the full channel; and the Galileo 1992 degeneracy is disclosed as a stated limitation on that member's discriminating power rather than absorbed silently.
- *Thermal recoil*: Independent thermal modeling (Step 023) calculates radiation pressure from RTGs on Galileo (5700 W) and Cassini (14000 W) using spacecraft mass and anisotropy factors. For Galileo 1990, the integrated thermal Δv is 7.4×10^{-3} mm/s— 1.9×10^{-3} times the observed 3.92 mm/s anomaly. For Cassini, thermal recoil contributes 7.1×10^{-3} mm/s vs the corrected -2.00 mm/s catalogue entry ($\lesssim 0.4\%$ in magnitude, and opposite in sign). While thermal effects cannot explain the primary anomaly signal, Cassini's small observed anomaly could have a secondary thermal contribution. Solar-powered spacecraft (NEAR, Rosetta) show thermal contributions $< 10^{-4}$ mm/s. Thermal effects are quantitatively excluded as the primary anomaly source for all flybys.
- *Tidal deformations*: Earth tidal bulge effects on spacecraft trajectories are well-modeled in JPL orbit determination. Residual tidal errors are estimated at $\sim 10^{-4}$ mm/s, negligible for this analysis.
- *Solar radiation pressure*: SRP produces steady accelerations $\sim 10^{-7}$ mm/s², integrated over flyby duration yields $\sim 10^{-3}$ mm/s velocity change. SRP is already included in standard orbit determination.

Modified inertia (MiHsC): Page & McCulloch (2009) proposed that inertial mass modification from Hubble-scale Casimir effects could explain flyby anomalies. Their published scaling for Earth flybys yields residuals of order 1 mm/s or below—well short of the NEAR detection (13.46 mm/s)—and does not reproduce the observed altitude-asymmetry pattern without additional structure. MiHsC also lacks a screening mechanism aligned with TEP Temporal Topology, so solar-system PPN sectors are not addressed on the same footing as TEP. See: Page, G., & McCulloch, M. E. (2009). "Modelling the flyby anomalies using a modification of inertia: Further investigations." *Int. J. Astron. Astrophys.*, 3(1), 1-5.

General relativistic frame-dragging (Lense-Thirring): Independent first-principles calculation (Step 044) integrates the gravitomagnetic acceleration $a_{LT} = -2\mathbf{v} \times \mathbf{\Omega}_g$ along each osculating Kepler hyperbola reconstructed from the perigee state vectors (Step 038). For NEAR the accumulated shift is 2.6×10^{-4} mm/s in total magnitude (6.6×10^{-7} mm/s along-track); across all flybys the magnitudes span 3.3×10^{-5} – 4.0×10^{-4} mm/s, with along-track projections up to $\sim 10^{-5}$ mm/s — consistent in order with the literature estimate of $\sim 10^{-5}$ mm/s. Frame-dragging is thereby quantitatively excluded: its largest contribution sits three to four orders below the detected anomalies and $\sim 50\times$ below the smallest published bound.

Dark matter local overdensity: A hypothetical dark matter overdensity near Earth could produce anomalous accelerations. However, the required density ($\sim 10^{-9}$ GeV/cm³) would conflict with orbital dynamics of satellites and lunar laser ranging constraints. No independent evidence supports such an overdensity.

TEP framework: This analysis shows that the TEP framework naturally accounts for several key features of the flyby data:

- Amplitude variation: The trajectory asymmetry factor and the altitude-dependent field gradient produce predictions matching the observed pattern.
- Cassini stress test: The trajectory asymmetry factor $\cos \delta_{\text{in}} - \cos \delta_{\text{out}} = -0.022$ is verified by independent JPL Horizons reconstruction (Table 1a); it is a physical property of the Cassini trajectory, not a convention outlier. The negative shear-channel reference prediction is therefore expected, and it agrees in sign with the corrected catalogue anomaly (-2 ± 1 mm/s); the Earth-vicinity Temporal Shear is screened, so that channel contributes no anomaly regardless of sign. The clock-sector forward model returns a small positive reconstruction in the link class that carries the term ($\Delta V_{\text{app}} = +0.01$ mm/s vs the published -2.00 ± 1.00 mm/s, Table 3h), while the channel accounting of Section 5.2 classifies the published datum as coherent two-way — a class in which the clock term is absent — so the comparison is a mechanism consistency check rather than an attribution.
- Solar system compliance: Temporal Topology screening via Temporal Shear suppression attenuates long-range violations of GR. Cassini PPN compliance is evaluated through the solar source-charge projection $\alpha_{\text{eff}}^{(\odot)} = S_{\Sigma}^{(\odot)} \alpha_0$ (Section 4.6.1a), not through the fitted flyby amplitudes β_{fit} .
- Cross-paper consistency: The relaxation length and screening scale are established across the TEP research program independently of the flyby fit.

Comparative assessment: Table 10 summarizes the explanatory power of each proposed mechanism and separates the empirical TEP geometry template from the physical clock-sector mechanism in the published coherent two-way tracking observable. Standard physics effects and frame-dragging are quantitatively excluded on the high-altitude detections, while the TEP evaluation is split into two distinct rows: the restricted geometry template organizes the catalogue's amplitude, altitude, and null pattern, whereas the clock-sector mechanism cancels identically in the coherent two-way transponder class assigned to the published catalogue and instead predicts a falsifiable signal exclusively in clock-carrying links.

Table 10: Comparison of Flyby Anomaly Explanations

Mechanism / Model Layer	Amplitude Match	Altitude Dependence	PPN Compliant	Predicts Nulls
Atmospheric drag	$\times$ ($\leq 1\%$ of anomaly; order-comparable at the 303 km pass — OD-modelled degeneracy)	—	✓	$\times$
Thermal recoil	$\times$ ($10^{-2} \times$ too small)	$\times$	✓	$\times$
MiHsC	$\times$ ($10^{-1} \times$ too small)	$\times$	?	$\times$
Frame-dragging	$\times$ ($\sim 10^{-4} \times$ too small)	—	✓	$\times$
TEP geometry template (empirical envelope, $k_{\text{fit}} = 1$, $k_{\text{heur}} = 7$)	✓ (organizes $n = 6$ nonzero entries)	✓	✓ (via screened $S_{\Sigma}^{(\odot)}$)	✓
TEP clock-sector mechanism in published observable (coherent 2-way vs. clock-carrying links)	$\times$ in coherent 2-way (cancels at transponder); untested in clock-carrying links (pending raw-link verification)	✓ (in clock-carrying links only)	✓	✓ (predicts exact coherent 2-way null)

The coherent cancellation is a phase-reference result, not cancellation of an A^2 path-length integral. Step 046 solves emission, turnaround and reception epochs and obtains, within its prescribed geometry and conformal-clock bookkeeping,

$$\frac{\nu_r}{R\nu_u} = \frac{A_{\text{gnd}}(t_1)}{A_{\text{gnd}}(t_r)} \frac{dt_1}{dt_r}. \quad (29)$$

The spacecraft factor cancels at the common turnaround event. Ground-clock evolution and the light-time Jacobian remain; ordinary metric, disformal and finite-probe propagation contributions must be evaluated in their own observable model. Static conformal null-cone invariance does not make A^2 a refractive index.

The spacecraft endpoint cancellation acts at the coherent transponder itself. A station transmits a carrier referenced to its own clock; the spacecraft receives the matter-frame frequency $\nu_u A_{\text{gnd}}/A_{\text{sc}}$, and a coherent transponder—which contains no free-running reference—multiplies the received carrier phase by a fixed ratio R , re-emitting at coordinate frequency $R\omega_u$ regardless of the spacecraft's own rate. The spacecraft conformal factor cancels exactly; $\nu_{\text{ret}} = R\nu_u$ denotes that isolated endpoint bookkeeping, not the total Doppler observable. In Step 046's conformal benchmark the remaining station-rate term is $\sim \dot{\eta}_{\text{gnd}}\tau_{\text{light}}$, evaluated event-by-event on the real Horizons arcs at 10^{-17} – 4×10^{-16} fractional (Step 046; the $\sim 10^{-18}$ figure corresponds to the residual accumulated over the $\sim \mu\text{s}$ transponder turnaround). The spacecraft's accumulated proper-time excursion $\delta\tau(t) = \int \eta dt$ is a real offset of its own frame, but it reaches a tracking observable only where the link references an onboard time standard—one-way downlinks referenced to a free-running oscillator, non-coherent or regenerative turnarounds, and onboard-clock telemetry—where the term $c\eta$ enters once, at ~ 0.4 m/s raw amplitude, roughly three orders of magnitude above the mm/s anomaly band. Disformal and finite-probe corrections require their own path and boundary calculation; bounds from a particular holonomy benchmark are not a universal consequence of endpoint cancellation.

The observable forward model is the existing Step 043 clock injection and orbit fit, not a force integral relabelled as a clock effect. After the specified station-clock calibration, its raw clock-carrying range-rate contribution and Step 046's coherent fractional-frequency residual are

$$\delta\nu_{\text{raw}}^{\text{clock}} = c\eta_{\text{sc}}, \quad \left(\frac{\delta\nu}{\nu}\right)_{\text{coherent}}^{\text{clock}} = \mathcal{O}(\dot{\eta}_{\text{gnd}}\tau_{\text{light}}) \lesssim 4 \times 10^{-16} \quad (\text{Step 046 benchmark}). \quad (30)$$

The fitted apparent ΔV is obtained by differencing the clean and clock-injected orbit solutions, as implemented in Step 043. It is not generally equal to $c\eta$ or to a scalar-force impulse; its magnitude depends on arc duration and the fitted nuisance parameters. The gradient-based catalogue template remains an empirical geometry descriptor, separate from this channel-specific forward prediction.

Channel separation under the amplitude–shear split: in the two-way Doppler observable the first-order excursion $\eta \sim \beta_A \phi / M_{\text{Pl}} \sim 10^{-9}$ cancels identically at the coherent turnaround, leaving only the time-evolution residual $\sim \dot{\eta} \tau_{\text{light}}$ at 10^{-17} – 4×10^{-16} fractional (Step 046), and the Earth-vicinity Temporal Shear is screened to negligible strength by the two-body kinetic operator — so neither the raw clock signal nor a shear force can source the anomaly in the observable itself, and with the transponder cancellation made explicit the reduction chain cannot misread a term that is absent from the signal. On the literature-based coherent-link classification of Table 11, the catalogued anomalies are therefore not generated by the TEP clock sector in the measured observable: if they persist as genuine anomalies they stand as unexplained data, and if they are reduction artefacts — as independent signal-path analyses have argued for the NEAR case (Antreasian & Guinn 1998) — their origin lies in the tracking chain rather than in the temporal field. What the clock sector retains is the sharper claim: a flyby tracked through a clock-carrying link accumulates a $c\eta$ excursion at m/s scale at perigee, of which the Step-043 demonstration reconstructs $\sim 0.1\%$ as an apparent mm/s ΔV . That is the falsifiable prediction the artefact model actually makes — a signature no published flyby catalogue has yet been in a position to measure.

Table 11: Literature-based per-mission measurement-link classification of the published flyby catalogue (Step 046 ledger, `step046_two_way_cancellation.json`). The listed publications and standard tracking practice motivate coherent two-/three-way transponder assignments; encounter-specific raw links have not been independently verified for every row. A coherent three-way link (different uplink/downlink complexes) carries no independent onboard oscillator.

Flyby	Published Δv (mm/s)	Assigned measurement-link class	Source
NEAR	13.46 ± 0.13	Coherent two-way DSN Doppler (X/S band)	Antreasian & Guinn (1998); Anderson et al. (2008)
Galileo 1990	3.92 ± 0.08	Coherent two-way DSN Doppler	Anderson et al. (2008)
Galileo 1992	-4.60 ± 1.00	Coherent two-way DSN Doppler	Anderson et al. (2008)
Cassini	-2.00 ± 1.00	Coherent two-way DSN Doppler	Anderson et al. (2008)
Rosetta 2005	1.80 ± 0.03	Coherent two-way DSN/ESTRACK Doppler	Anderson et al. (2008)
Rosetta 2007	0.02 ± 0.05	Coherent two-way DSN/ESTRACK Doppler	Müller et al. (2008)
Rosetta 2009	0.00 ± 0.05	Coherent two-way DSN/ESTRACK Doppler	Müller et al. (2010)
MESSENGER	$+0.02 \pm 0.01$	Coherent two-way DSN Doppler	Anderson et al. (2008)
Juno	0.00 ± 0.02	Coherent two-way DSN Doppler (X/Ka band)	Aksenov & Tuchin (2020)

Step 046 emits these literature-based assignments into `results/step046_two_way_cancellation.json`; its `clock_carrying_published` list is empty by those assignments, not by an independent audit of every encounter's original tracking configuration. The coherent-transponder endpoint cancellation is a separate physical calculation and applies to each datum confirmed to use that class. An encounter-specific onboard-reference contribution would require direct link evidence before the catalogue could test the proposed clock mechanism. On the present classification, the published anomaly pattern is not attributed to that mechanism; its forward-modelled signature ($c\eta \sim \text{m/s}$ raw, $\sim \text{mm/s}$ postfit remnants and structured residuals in Step 043) awaits a documented clock-carrying pass.

Clock-reference distinction: an independently referenced one-way downlink or onboard-clock telemetry can carry the spacecraft clock excursion. A coherent three-way link, despite using different transmitting and receiving ground stations, still removes the spacecraft endpoint factor; it may retain ground-station contrasts and separately modelled propagation terms. A proposed mission tests the clock-sector prediction only if its actual phase or timing chain contains the independent onboard reference. The raw amplitude $c\eta$ is conditional on the adopted profile and its clock tests, including Section 5.7a.

Theoretical consistency achieved: under the channel accounting above, the anomaly's status in this framework is conditional on observable class rather than asserted for the catalogue. The fitted response structure reported in this paper is the empirical organization of the measured asymmetry pattern; the clock-sector artefact is the physical mechanism wherever the link carries a spacecraft time reference, and is identically absent in the coherent two-way class. The corpus-level claim is thereby the sharper one: the artefact is not required to hide in the existing catalogue — it is a falsifiable signature awaiting the first flyby tracked through a clock-carrying link.

5.3 Cross-Paper Consistency: Lunar Laser Ranging

The TEP screening mechanism—specifically the Temporal Topology Saturation Scale saturation ($\rho_T \approx 20 \text{ g/cm}^3$) and the consequent Earth saturation core ($R_{\text{sol}} \approx 4146 \text{ km}$)—provides a cross-paper consistency check on the same screening framework through precision Lunar Laser Ranging (LLR) analysis in related work.

LLR Consistency Check

The LLR analysis reports a synodic-phase signal with magnitude $\kappa_{\text{LLR}} \sim -4 \times 10^{-4}$, of the sign and order expected for a screened response: the pooled flyby amplitude under the canonical normalization maps to an effective coupling $\beta_{\text{fit}} S_{\oplus} \approx 1.7 \times 10^{-3}$, within an order of magnitude of the LLR signal.

The negative sign of κ_{LLR} suggests that gravitational potential screening (Temporal Shear suppression) dominates over surface-scaling mechanisms, providing qualitative consistency with the TEP framework.

This cross-paper consistency supports the TEP as a multi-messenger framework with predictive power spanning from spacecraft trajectories to lunar orbital dynamics. Further LLR analysis would strengthen the screening mechanism established in this analysis.

5.4 Remaining Limitations

The span in fitted β_{fit} (factor of ~ 3.3 across the tightly constrained members NEAR, Galileo 1990, and Rosetta 2005, extending to 6.65×10^{-2} for the negative-asymmetry Galileo 1992 member and formally divergent for the near-zero-prediction Cassini and MESSENGER members; Step 009) reflects genuine geometry-dependent modulation (altitude, latitude, velocity, and the envelope's bounded heuristic plasma factor, Section 3.10). Step 009 does not perform a formal variance decomposition because $n = 6$ renders ANOVA-style partitioning statistically underpowered (standard error on sample variance is a large fraction of the estimate). Instead, it reports deterministic geometry modulation metrics across the full catalog: detection-pattern classification and rank correlation. See Section 4.3 for the quantitative assessment and Section 5.5 for interpretation.

Model completeness: after the catalogue re-transcription, Cassini no longer exhibits a sign mismatch — the Step 007 reference prediction at $\beta_{\text{ref}} = 10^{-4}$ (negative total Δv_{TEP}) agrees in sign with the corrected published anomaly ($-2 \pm 1 \text{ mm/s}$). The residual tension is amplitude-only: the predicted magnitude is numerically negligible ($\sim 1 \times 10^{-5} \text{ mm/s}$), so Cassini's per-flyby β_{fitted} is formally divergent and its weight in the pooled mean is negligible. The same applies to MESSENGER. The amplitude-informative members are NEAR, Galileo 1990, Rosetta 2005, and Galileo 1992; raw DSN OD reanalysis remains the key test of whether the published amplitudes reflect the trajectory response or filtering.

5.5 Systematic Error Discrimination

The geometry-correlation argument: A qualitative discriminator between TEP and systematic errors lies in the correlation pattern between anomalies and trajectory geometry, though with $n = 6$ gated detections the ordering is positive but not decisive (Spearman $\rho = +0.66$, $p = 0.156$, Step 009) and cannot serve as a primary discriminator. TEP theory explicitly predicts that anomaly magnitude should correlate with trajectory asymmetry ($\cos \delta_{\text{in}} - \cos \delta_{\text{out}}$) because this factor determines how asymmetrically the spacecraft samples Earth's oblate field. Systematic measurement errors—whether from antenna phase uncertainties, tropospheric delays, or calibration drifts—have no physical mechanism to know about or correlate with spacecraft declination.

The observed Spearman correlation between |trajectory asymmetry| and |anomaly magnitude| is $\rho = 0.56$ ($p = 0.11$) across the $n = 9$ catalogued flybys carrying published asymmetry values, and $\rho = 0.68$ ($p = 0.090$) across the $n = 7$ flybys with nonzero published anomalies (Step 009 ordering diagnostic). Within the gated $n = 6$ subset the ordering is positive but not decisive ($\rho = +0.66$, $p = 0.156$); NEAR anchors the top rank on both axes while the mid-range orderings partially invert. The qualitative pattern across the literature set survives re-transcription: large positive asymmetries associate with the largest positive anomalies, and the negative-asymmetry members (Cassini, Galileo 1992) now carry matching-sign anomalies. Hardware biases (antenna phase: 0.1 mm/s, station position: 0.02 mm/s, tropospheric delay: 0.05 mm/s) are altitude-independent and geometry-blind. Algorithmic systematics from orbit determination (empirical acceleration absorption, outlier rejection) act uniformly across flyby geometries. A clock-sector response coupled to Earth's field-exursion structure provides a natural mechanism for the correlation pattern stressed in the historical literature, though the evidence remains qualitative at the available sample size.

The scaling argument: With six $S/N \geq 2$ published anomalies in the literature (all six in the inverse-variance β_{fit} ensemble after the catalogue re-transcription), statistical noise remains non-negligible and the explicitly modelled measurement systematics (0.12 mm/s total) are subdominant to the primary observed anomalies (1–13 mm/s). The concern that systematic errors dominate at large n —where statistical noise vanishes but systematics persist—is valid for high- n validation but secondary at the current sample size. Because sign agreement must be evaluated on the ungated catalogue rather than inside a sign-conditioned subsample, Step 008 audits sign concordance prior to the $S/N \geq 2$ gate: all 7 of 7 published nonzero anomalies in Table 1 (adding Rosetta 2007 to the six $S/N \geq 2$ entries) agree in sign with their reference predictions at β_{ref} (one-sided binomial $p = 0.0078$), while within the $S/N \geq 2$ subset the rank ordering is positive but imperfect ($\rho = +0.66$), so the organisational support is bounded by the modest ordering significance rather than by any sign contradiction. A rigorous permutation test for Spearman rank correlation (scripts/utls/statistical_utils.py) quantifies this: under random reshuffling of anomaly magnitudes across flybys, the probability of observing the current rank ordering by trajectory asymmetry is computable exactly for $n \leq 10$ and via 10^5 Monte Carlo permutations for larger samples. The hardware and propagation measurement errors modelled in Step 024 (antenna phase motion, tropospheric delay, station position errors) are geometry-blind and do not by themselves reproduce the anomaly–asymmetry rank association.

Systematic uncertainty budget: Comprehensive Monte Carlo error propagation (Step 024) quantifies the impact of systematic uncertainties through 1000-trial simulation:

- Measurement systematics (DSN): Antenna phase center (0.10 mm/s), tropospheric delay (0.05 mm/s), station position (0.02 mm/s). Total: 0.12 mm/s (1% of 13.46 mm/s NEAR anomaly).
- Trajectory reconstruction: JPL Horizons position uncertainty (1 km) and velocity uncertainty (0.1 m/s) contribute $\sim 1\%$ to predicted Δv .
- Characteristic suppression uncertainty: From the UCD saturation model, $\rho_T = 20 \pm 7 \text{ g/cm}^3$ (35% systematic, Paper 6 UCD) propagates to $R_{\text{sol}} = 4146 \pm 540 \text{ km}$ ($\sim 13\%$) and $S_{\oplus} = 0.35 \pm 0.09$ ($\sim 25\%$). The GNSS correlation length ($L_c = 4201 \pm 1967 \text{ km}$, Step 016) is the terrestrial anchor of that calibration; its agreement with R_{sol} is a shared-input consistency check under the Paper 6 §2 identification, not an independent validation.
- Multipole coefficients: J2/J3 known to $< 0.1\%$ from GRACE/GOCE—negligible contribution.
- Relaxation length uncertainty: $\lambda_T = 4200 \text{ km}$ with $\pm 15\%$ relative uncertainty from the SCF theoretical prior (Paper 6 UCD). The raw GNSS correlation length ($4201 \pm 1967 \text{ km}$, 47%) is the anchor of the terrestrial calibration rather than an independent check on it; the SCF prior is used for the uncertainty budget.

The Monte Carlo analysis (Step 024) propagates the catalogued systematic inputs through the TEP prediction pipeline; at the Δv level, the simulated measurement-systematics band (0.12 mm/s total) is small compared to the mm/s-scale anomalies. This supports the conclusion that the catalogued DSN hardware and propagation systematics explicitly modelled in Step 024 are subdominant to the primary anomaly scale. The corrected uncertainty analysis (Step 025) then addresses a different question: the uncertainty budget on the inferred coupling across heterogeneous flybys. It provides a rigorous decomposition with total relative uncertainty 152.7% on β_{fit} , dominated by heterogeneity (149.4% relative uncertainty; 95.7% of variance) and with a substantial systematic component (29.2% relative uncertainty; 3.6% of variance) dominated by characteristic suppression uncertainty (25.0%) from the UCD saturation model ($\rho_T = 20 \pm 7 \text{ g/cm}^3$, Paper 6 UCD) and relaxation length uncertainty (15.0%) from the SCF theoretical prior. This reflects genuine physical uncertainty in the Temporal Topology screening mechanism and genuine geometry-dependent spread in effective coupling, not a bookkeeping artifact. The catalogue pattern is therefore not explained by the hardware and environmental systematics explicitly modelled in this study, though unmodeled orbit-determination or tracking-configuration artefacts in the historical records remain open until raw DSN tracking files are independently reprocessed.

5.6 Comprehensive Diagnostic Validation

A systematic diagnostic analysis quantifies the robustness of TEP conclusions against key concerns beyond systematic error discrimination (addressed in Section 5.5):

Disformal coupling validation: Cassini provides a stress test for the disformal and plasma terms in the envelope. At β_{ref} the model returns a negative total prediction, driven by the published Anderson et al. (2008) geometry factor $\cos \delta_{\text{in}} - \cos \delta_{\text{out}} = -0.022$ — a value now confirmed identical between the published column and the independent JPL Horizons reconstruction (Table 1a). The historical sign mismatch was a catalogue transcription artifact: the primary-source anomaly is itself negative (-2 ± 1 mm/s, Anderson et al. 2008), so the kernel's predicted sign agrees with the published anomaly and the residual tension is magnitude-only. The TEP model faithfully propagates this literature geometry; no modification of the clock-sector response model is indicated.

Model parameter sensitivity: The TEP model maintains fit stability across a broad range of characteristic response factors ($S_{\oplus} = 0.30$ to 0.50), indicating the clock-sector amplitude response is robust, not fine-tuned.

Diagnostic conclusion: The model maintains fit stability across broad parameter variations; Bayesian model comparison places the Anderson empirical law nominally first on the gated ensemble, with the restricted TEP model at parity on the full catalog and all structured models far ahead of the Null (Section 4.5). The full-catalog raw stress test on $n = 9$ improves over the null by $\Delta \log L \approx 43.0$ after random-effects prediction uncertainty is included; it should be read as a consistency check rather than confirmation that the small gated sample is decisive. Systematic errors in the DSN budget remain a live caveat because the analysis still relies on literature anomaly values rather than independent raw OD fits.

5.7 Enhanced Statistical Validation

Comprehensive statistical validation is reported in Sections 4.1, 4.5, and 4.8. The headline model comparison uses the full $n = 9$ catalog under the fractional misspecification floor ($\sigma_{\text{sys},i} = CV_{\beta}|\hat{y}_i|$, $CV_{\beta} \approx 0.47$); the $n = 6$ gated detection ensemble is the robustness layer (Section 4.5.3). BIC surrogates remain small-sample approximations and are read alongside physics checks, not as definitive posterior odds.

Juno falsification risk: The Juno 2013 flyby ($\Delta v_{\text{obs}} = 0.00 \pm 0.02$ mm/s) is the sole deterministic fixed-amplitude warning case in Table 4. At the refit weighted-mean β_{fit} , Step 039 predicts $+0.12 \pm 0.03$ mm/s after random-effects amplitude uncertainty, so the uncertainty-aware classification retains Juno as a raw-tension case at $\sim 3\sigma$. This stresses the single-scale diagnostic, but it is not a 5σ falsification under the honest prediction-uncertainty model (Section 5.7a).

OD suppression status: Step 021 withholds mission-specific F_{OD} values because real mission OD configuration files are not publicly available. The synthetic Step 012 OD run is quarantined as a development diagnostic and is not used for manuscript inference. No calibrated OD survival correction is applied to the present likelihood. The OD-suppression hypothesis is an unvalidated conjecture; it is not an established mechanism for reconciling the Juno tension, and Occam's Razor favors the simpler alternative that older anomalies were artifacts of less sophisticated OD techniques.

Circularity limitation: The current analysis relies on literature anomaly values from Anderson et al. (2008) and subsequent papers, rather than independent DSN data analysis. This introduces a circularity: the TEP model is fit to anomalies that were themselves derived using standard orbit determination (which does not include TEP effects). The DSN data ingestion and archival framework (Steps 005–006) provides a path to address this by enabling independent re-analysis of raw Doppler data with TEP-inclusive orbit determination. This would be a critical validation step.

Model completeness: The clock-sector response model includes the dominant effects (perigee proper-time offset through the amplitude channel, J2 oblateness, trajectory asymmetry, the geometric amplitude response $S_{\oplus}$) but may omit secondary effects that could contribute to heterogeneity. Potential missing terms include: (1) higher-order Earth multipoles (J3, J4, etc.), (2) Earth rotation (Lense-Thirring effect), (3) non-spherical Temporal Topology geometry, (4) time-varying ϕ during the brief perigee passage, (5) spacecraft mass-to-surface-area ratio affecting radiation pressure coupling to the scalar field. Incorporating these effects could further reduce β scatter.

Environmental response dependence: The flyby amplitude relies on the UCD-derived characteristic response $S_{\oplus} \approx 0.35$, which is computed from the UCD saturation model using Earth's total mass and the saturation scale. The clock-sector amplitude response emerges naturally from the UCD framework rather than being phenomenologically tuned. This cross-scale prior — anchored on, rather than independently confirmed by, the GNSS correlation length under the Paper 6 §2 identification — provides a rigorous foundation for the environmental response without empirical fitting to flyby data.

- **Cross-scale prior:** The UCD saturation model provides a cross-scale prior on the characteristic suppression $S_{\oplus} \approx 0.35$ from the saturation scale $\rho_{\text{T}} = 20$ g/cm³. Under the terrestrial identification $L_c \sim \mathcal{O}(1) R_T(M_{\oplus})$ the GNSS covariance scale returns $S_{\oplus} \approx 0.34$ — a shared-input consistency check at the 2% level (the same measurement anchors ρ_{T}), consistent with the GNSS/UCD terrestrial calibration without fitting to flyby data.
- **Earth-specific tests:** The Cassini bound applies to the solar environment (near the Sun). Earth-specific precision tests could provide complementary constraints: (1) Lunar Laser Ranging (LLR) tests of the strong equivalence principle, (2) Gravity Probe B (GP-B) frame-dragging measurements, (3) satellite laser ranging (SLR) to LAGEOS and LARES satellites, (4) atomic clock

comparisons at different altitudes (e.g., ACES mission). These Earth-based tests would directly constrain the effective coupling $\beta_{A,\text{eff}}$ in the terrestrial environment where flybys occur.

- GNSS cross-validation: The GNSS covariance scale $L_c \approx 4200$ km — the terrestrial anchor of the $R_T(M_\oplus)$ identification — can be tested against independent GNSS datasets (e.g., different satellite constellations, different analysis centers); a genuinely independent validation of the identification itself would come from a clock network resolving the saturation scale of a different host mass or profile. Consistency across independent analyses would strengthen confidence in the anchor and hence in the characteristic suppression.
- Laboratory tests: Fifth-force searches in laboratory settings (e.g., torsion balance experiments, atom interferometry) can constrain β_{fit} at short ranges. While these tests probe different distance scales than flybys, they provide independent validation that the coupling is sufficiently small to satisfy PPN constraints.

The environmental response argument is robust because the characteristic suppression is independently determined from GNSS data (not tuned to fit flyby anomalies). Full margins are reported in Section 4.6. Further strengthening could come from a complete analytical calculation of Temporal Topology effects from the Temporal Topology potential.

Sample size as complete dataset: The analysis includes all available Earth gravity assist flybys with adequate DSN tracking precision — six published nonzero entries — four firm detections and two $\sim 2\sigma$ marginal entries (Cassini, MESSENGER) — three published nulls/bounds, and several flybys with no public anomaly report. The Step 008 inverse-variance β_{fit} ensemble uses six S/N-qualified primary fits (NEAR, Galileo 1990, Rosetta 2005, Cassini, Galileo 1992, MESSENGER), all sign-consistent at $\beta_{A,\text{ref}}$ on the corrected catalogue. Step 026 model comparison uses the same $n = 6$ gated set. Effect sizes relative to the null population ($n_{\text{null}} = 3$) are: NEAR $d \approx 2.5$ (large), Galileo 1990 $d \approx 0.73$ (medium), Galileo 1992 $d \approx -0.87$ (large in magnitude, negative branch), Rosetta 2005 $d \approx 0.34$ (small), Cassini $d \approx -0.38$ (small), and MESSENGER $d \approx 0.003$ (negligible). The NEAR and Galileo detections provide the bulk of the statistical separation from null results. Leave-one-out analysis on the six-member ensemble shows moderate stability (coefficient ≈ 0.092). Additional flybys would test envelope and plasma refinements rather than restate the same $n = 6$ compression.

5.7a Falsifiability and the Juno Null Test

Juno remains the sole deterministic fixed-amplitude warning case at the refit weighted-mean β_{fit} . The Step 039 row forecasts $+0.12 \pm 0.03$ mm/s for the 2013 flyby after random-effects amplitude uncertainty, while the published bound is 0.00 ± 0.02 mm/s. This is an open stress test for the diagnostic, not a resolved puzzle and not a 5σ discrepancy under the current uncertainty model.

Table 4 provides the classification framework:

- Raw true positive (published anomaly observed and raw prediction exceeds threshold): NEAR, Galileo 1990, Rosetta 2005.
- Raw true null (published null or sub-threshold bound consistent with raw prediction): Cassini, MESSENGER, Rosetta 2007, Rosetta 2009.
- Raw surplus (published anomaly with sign-consistent but underpredicted raw prediction): Galileo 1992 (-0.33 mm/s predicted vs -4.60 ± 1.00 mm/s observed).
- Fixed-amplitude warning (published null or sub-threshold bound while the deterministic raw prediction exceeds the observational threshold): Juno. The uncertainty-aware raw classification retains this single raw-tension case after random-effects amplitude scatter is propagated ($\sim 3\sigma$).
- No public anomaly report (not used in quantitative likelihood): Stardust, OSIRIS-REx, BepiColombo.

Three resolutions are possible, listed in order of epistemic priority:

1. Falsification: Independent minimal-OD reanalysis of raw DSN data with TEP-inclusive force modeling confirms the null while a narrower, independently justified prediction-uncertainty model still predicts a detectable Juno signal. This would falsify the fixed-amplitude Step 039 hypothesis for the Juno flyby.
2. Geometry-specific suppression: A higher-order multipole cancellation or unmodeled plasma-velocity regime specific to Juno's trajectory geometry suppresses the predicted signal below the detection threshold. (The envelope's plasma factor is already applied to Juno's prediction as a bounded heuristic coefficient, and the stronger Debye ansatz is quarantined at $S_{\text{plasma}} = 1$, Section 3.10 — so this option reduces to an envelope-geometry statement.)
3. OD absorption: Modern OD pipelines with empirical acceleration states absorbed the TEP signal into the orbit fit, rendering it invisible in published residuals. This remains an unvalidated conjecture; no mission-specific F_{OD} values are available (Step 021), and the synthetic Step 012 diagnostic is quarantined from manuscript inference.

Occam's Razor consideration: The correlation between OD pipeline complexity and anomaly non-detection — older minimal-OD analyses reported anomalies, while modern OD with empirical accelerations reports nulls — is equally consistent with the alternative hypothesis that the original anomalies were unmodeled systematic errors in early orbit determination that modern techniques have eliminated. Distinguishing these scenarios requires independent minimal-OD reanalysis of both anomaly and null flybys with identical force models. Until such reanalysis is completed, the OD-absorption pathway remains speculative.

Step 021 withholds all mission-specific F_{OD} values because real mission OD configuration files are not publicly available. Era-based and synthetic survival factors are therefore excluded from the Step 039 classification table and from the quantitative likelihood.

Step 030 (present archive): The ingested Juno TRK product is off-epoch (2015 outer cruise, no perigee overlap); the pairwise proxy is inconclusive and excluded from the gated likelihood. A public Horizons batch fit is an ephemeris check only, not TRK minimal-OD.

Residual Tensions: Galileo 1992 and Juno

Galileo 1992: At the refit weighted-mean β_{fit} , Step 039 predicts -0.33 ± 0.09 mm/s against the published -4.60 ± 1.00 mm/s detection — the sign is correct but the magnitude is underpredicted by an order of magnitude (raw surplus classification). The low-altitude, weakly-asymmetric geometry yields a near-null reference prediction that is amplified only to a fraction of the observed value at the pooled coupling; perigee-epoch DSN reanalysis remains the key independent test.

Juno: As in Section 5.7a—fixed-amplitude warning at pooled β_{fit} , retained at $\sim 3\sigma$ even under the uncertainty-aware scheme; perigee TRK minimal-OD not closed in this release.

5.8 PPN Constraint Satisfaction and Cassini Solar Conjunction

The Cassini solar conjunction bound ($|\gamma - 1| < 2.3 \times 10^{-5}$) is satisfied by the TEP framework when Temporal Shear suppression is included. Section 4.6.1a adopts the Paper 0 v0.16 solar source-charge convention $S_{\Sigma}^{(\odot)} \sim 2 \times 10^{-6}$, inside the Cassini requirement $S_{\Sigma}^{(\odot)} \lesssim 1.15 \times 10^{-5}$ under the unified unstretched-space metric (5.8×10^{-6} on the isotropic-conformal benchmark); the kinetic-operator benchmark evaluates the solar-field screening far deeper, $\sim 10^{-15}$ at the 1 AU path. Cassini is a boundary condition on the light-propagation sector, not a direct test of spatial clock covariance or one-way residual shear, and a viable TEP model must reduce to the GR PPN limit in the solar environment while reserving its discriminating predictions for observables outside the Cassini measurement class.

5.9 Theoretical Implications

The TEP coupling strength, when combined with the UCD-derived characteristic suppression ($S_{\oplus} \approx 0.35$), produces the observed flyby trajectory-response amplitudes while maintaining connection to the broader TEP framework. The screening scale is anchored on the GNSS terrestrial calibration through the UCD identification (Step 010); its independent empirical content is supplied by the flyby altitude threshold and the cross-scale ρ_T consistency (Paper 6).

The canonical Paper 0 amplitude-sector prescription — the parameter-free terrestrial excursion $\psi_{\text{uns}} = 2GM_{\oplus}/(R_{\oplus}c^2) = 1.3926 \times 10^{-9}$, independent of β_A — produces physically consistent Earth-scale gradient suppression ($\lambda_T = 4200$ km) while remaining connected to the scalar-tensor theory structure; the excluded inverse-power realization is retained in Step 010 only as an excluded-branch contrast (Section 3.4). The fitted $\beta_{\text{fit}} \sim 10^{-3}$ – 10^{-2} range (spanning 4.43×10^{-3} to 1.44×10^{-2} across the tightly constrained members), when attenuated by the UCD-derived characteristic suppression $S_{\oplus} \approx 0.35$, yields PPN-safe effective couplings that organize the observed anomalies under the response model — the mechanism-level attribution remaining conditional on the link-class accounting of Section 5.2.

Cosmographic modulation: The velocity-activated branch of the fitted response template depends on the total velocity in the ambient-field rest frame (the frame in which the spatial gradient of the cosmological temporal field vanishes). If the CMB dipole frame approximates this frame, the ~ 370 km/s bulk motion of the Solar System would provide a cosmographic modulation of the template's velocity factor — a directional hypothesis about the phenomenological envelope, since the canonical-branch propagating amplitude is negligible at flyby velocities (§3.5). Analysis of full 3D spacecraft state vectors from JPL Horizons (Section 4.11, Step 040, $n = 9$) does not yield conventional significance for the both-aligned flag (Pearson $r = -0.57$, $p = 0.11$; Mann-Whitney $U = 7.5$, $p = 0.79$). Multivariate regression on residual ratio achieves $R^2 = 0.963$ (adjusted $R^2 = 0.940$), and an optimal weighted combination of spacecraft and Earth CMB projections attains $r = 0.96$ ($p = 5.5 \times 10^{-5}$) at the edge of the scanned weight grid ($w = -2.0$), so the combination is not interpreted as an interior optimum. The CMB-frame result remains an exploratory directional check rather than a definitive confirmation.

5.10 Falsifiability and Predictive Power

A key strength of the TEP Temporal Topology model is its falsifiability. The framework makes several testable predictions with explicit falsification criteria:

Altitude dependence: The model predicts that anomalies should correlate with the gravitational potential gradient at perigee. Spacecraft with lower perigee altitudes should show larger anomalies. The observed correlation—NEAR (539 km, 13.46 mm/s) vs. MESSENGER (2347 km, negligible)—matches this prediction quantitatively.

Falsification criterion: A flyby at altitude < 1500 km with trajectory asymmetry $|\cos \delta_{\text{in}} - \cos \delta_{\text{out}}| > 0.1$, non-cancelling inclination and plasma-envelope geometry, and continuous perigee-epoch DSN tracking analyzed under a matched minimal-OD pipeline (without empirical acceleration states that absorb short-arc residuals; Section 5.7) that shows no anomaly ($\Delta v < 0.5$ mm/s at 3σ) would falsify the altitude-dependence prediction. Without the asymmetry and minimal-OD conditions, a symmetric low-altitude trajectory could yield geometric cancellation (as for MESSENGER at 2347 km, $|\cos \delta_{\text{in}} - \cos \delta_{\text{out}}| \approx 0.005$) or empirical-state absorption, rendering a single-pass null non-discriminating.

Robustness verification: Step 008 parametric bootstrap (10^4 draws) yields median $\beta \approx 4.77 \times 10^{-3}$ with 95% interval $[4.42 \times 10^{-3}, 1.47 \times 10^{-2}]$, and leave-one-out recomputations 5.88×10^{-3} (without NEAR), 4.58×10^{-3} (without Galileo 1990), and 4.64×10^{-3} (without Rosetta 2005). The stability coefficient ≈ 0.092 is below the 0.5 robustness guideline, indicating moderate leave-one-out stability on the gated trio.

Heterogeneity assessment: The gated β_{fit} fits remain formally heterogeneous (see Section 4.8), indicating that a single scalar rescaling does not capture all geometry-dependent physics. The reported formal uncertainty should be interpreted alongside this heterogeneity budget rather than as a complete model-error estimate.

Physics-based interpretation of β_{fit} scatter: The span in gated fitted β_{fit} (factor of ≈ 3.3 across the tightly constrained members, ≈ 15 across the four amplitude-informative fits) reflects environment-dependent structural modulations arising from the covariant disformal mapping $B(\phi)$, the Step 007 plasma-velocity envelope (whose bounded plasma factor is a heuristic coefficient, not an action-derived coupling, Section 3.10), and Temporal Topology geometry—rather than measurement noise alone. Several mechanisms contribute within the TEP framework:

- **Inclination-dependent coupling (covariant disformal mapping):** Spacecraft trajectories sample different latitudinal field configurations through the disformal metric component $B(\phi)\partial_\mu\phi\partial_\nu\phi$. The Earth's oblateness ($J_2 = 1.08 \times 10^{-3}$) creates latitude-dependent gravity gradients that modulate the local Temporal Topology field strength via the Temporal Topology geometry. Polar trajectories (NEAR: $i \approx 50^\circ$) experience enhanced coupling relative to equatorial flybys (Galileo: $i \approx 12^\circ$) due to reduced equatorial bulge gradient suppression, producing multiplicative spread in fitted β_{fit} at fixed β_{ref} .
- **Velocity-direction asymmetry:** The clock-sector response depends on the spacecraft velocity vector orientation relative to the field excursion $\nabla\phi$. Inbound and outbound trajectories sample different effective field configurations, with the template's velocity-projected term $B(\phi)(v \cdot \nabla\phi)^2$ introducing velocity-dependent anisotropy in the fitted response strength (template-level modulation; the propagating amplitude is null on the canonical branch, §3.5).
- **Local-time plasma modulation (heuristic envelope coefficient):** The ionospheric plasma density varies with local time, providing environment-dependent gradient suppression of the scalar field through $f_{\text{plasma}}(\rho) = (1 + \rho/\rho_{\text{crit}})^{-0.3}$ on IRI-derived densities. This factor is an active nominal coefficient of the Step 007 envelope — contributing modulation factors of 0.18–0.91 across the gated ensemble — carried as a template-layer term rather than an action-derived scalar-plasma coupling (Section 3.10), and distinct from the stronger Debye attenuation ansatz quarantined at $S_{\text{plasma}} = 1$. Its day-side/night-side organisation remains a directional hypothesis pending a first-principles plasma coupling.
- **Velocity-dependent template regime:** High-velocity flybys ($v > 16$ km/s) enter the velocity-activated branch of the fitted response template (empirical scale $v_{\text{trans}} \approx 16.8$ km/s, phenomenological ansatz per §3.5). Cassini's high perigee velocity (19.0 km/s) samples this branch while also carrying a near-zero reference prediction, so its corrected anomaly is sign-consistent but its per-flyby amplitude rescaling is poorly constrained.

Model refinement opportunities: The β_{fit} scatter provides diagnostic power for improving the theory. Specifically:

- *Altitude-dependent gradient modulation:* The effective transition radius may vary with flyby geometry; a density-profile model incorporating Earth's crustal structure and core-mantle boundary could reduce tension for trajectories that currently show sign or amplitude mismatch at β_{ref} .
- *Trajectory effects:* Full 3D trajectory integration (Section 4.1.2) confirms that inclination- and velocity-dependent modulation along the reconstructed path contributes modest corrections ($\leq 20\%$) to the perigee approximation for primary detections. Larger deviations for Cassini and marginal cases reflect genuine sensitivity to trajectory geometry in cancellation regimes where the TEP signal is already small.
- *Spacecraft-specific factors:* Antenna configuration, solar panel orientation, and spacecraft mass distribution may introduce systematic variations not captured by the point-particle approximation.

Falsification criterion: A solar source-charge projection $\alpha_{\text{eff}}^{(\odot)} = S_{\Sigma}^{(\odot)} \alpha_0$ incompatible with the Cassini PPN band after honest propagation of UCD uncertainties would falsify the current screening narrative. The solar-screened calculation (Section 4.6.1a) remains inside that band.

PPN constraints: Any solar system test that improves the Cassini bound on γ would further constrain β_{fit} . Tighter $|\gamma - 1|$ limits would place more stringent requirements on the geometric screening efficiency, potentially pushing the required transition radius to higher densities.

Falsification criterion: A solar-conjunction measurement of $|\gamma - 1|$ exceeding the Paper 0 canonical quartic-branch source-charge ceiling ($S_{\Sigma}^{(\odot)} \lesssim 1.15 \times 10^{-5}$ unified [5.8×10^{-6} isotropic], $|\gamma - 1| < 2.3 \times 10^{-5}$) or contradicting the quadratic-branch kinetic-

screening prediction ($|\gamma - 1| \sim 10^{-15}$ at 1 AU) with opposite sign ($\gamma - 1 > 0$, violating the sign-definite $\gamma_{\text{PPN}} < 1$ theorem of Paper 0) would exclude the corresponding screening branch.

Sharp kill criterion: If independent DSN/OD reanalysis of NEAR, Galileo 1990, or Rosetta 2005 confirms the published anomalies but with revised trajectory geometry that yields positive TEP predictions at β_{ref} and produces fitted β_{fit} violating the Cassini PPN bound ($|\gamma - 1| > 2.3 \times 10^{-5}$), the TEP screening mechanism would be falsified.

Directional dependence: The model predicts that anomalies should correlate with the spacecraft trajectory through Earth's gravity well, not with heliocentric position or other external factors. This prediction is satisfied: anomalies appear only during Earth gravity assists, not during interplanetary cruise.

Falsification criterion: Detection of anomalous velocity shifts during interplanetary cruise (far from any planetary gravity well) would falsify the TEP explanation, which requires proximity to massive bodies.

Null results: The TEP framework explains two distinct categories of null results observed in the data: (1) *High-altitude gradient suppression* — flybys above ~ 2500 km (Stardust, OSIRIS-REx, BepiColombo, Rosetta 2007, and the published Rosetta 2009 null/bound) where the field gradient is too small to produce detectable effects; and (2) *Geometric cancellation* — flybys with nearly symmetric trajectories where the non-radial component of the artefact kernel cancels (MESSENGER at 2347 km, $|\cos \delta_{\text{in}} - \cos \delta_{\text{out}}| \approx 0.005$; Rosetta 2007 and 2009 are comparably weak-geometry cases). Galileo 1992, formerly grouped here under the transcribed null entry, is on the corrected catalogue a genuine negative detection and is treated with the detections. Both categories are supported by existing data.

Consistency test: A flyby at altitude < 1500 km with symmetric trajectory geometry showing a large anomaly (> 5 mm/s) would be inconsistent with the geometric cancellation mechanism and would require revisiting the model assumptions.

Eccentric-orbit clock bound: An unscreened ($y = 1$) Yukawa excursion $\psi_{\text{uns}} = 2\Phi_N$ extended on $\lambda_T = 4200$ km would imply an orbit-modulated clock-rate offset $\Delta\eta \approx 2.1 \times 10^{-11}$ on the Galileo disposal-orbit pair GSAT-0201/0202 ($a \approx 2.8 \times 10^4$ km, $e \approx 0.162$; $r_p = 23,445$ km, $r_a = 32,510$ km), exceeding the Delva et al. (2018) anomalous eccentric-orbit redshift bound ($\alpha = (0.19 \pm 2.48) \times 10^{-5}$, 2σ upper limit $\alpha < 5.15 \times 10^{-5}$ on $\Delta\Phi_N/c^2 = 5.27 \times 10^{-11}$) unless $\lambda_T \leq 1.30 \times 10^3$ km. However, when evaluated under Paper 0's selected two-branch master action ($du/dr = y_{\text{tb}}(g_{\oplus}) g_{\oplus}/c^2$, Step 045), kinetic screening at $r = 23,445\text{--}32,510$ km ($y_{\text{tb}} \approx 7.5 \times 10^{-7}$) suppresses the eccentric-orbit modulation to $\Delta u_{\text{master}} = 3.93 \times 10^{-17}$, corresponding to $\alpha_{\text{master}} = 7.45 \times 10^{-7}$ — passing the Delva et al. (2018) bound by a factor of $69.1\times$ with zero free parameters while leaving the lateral correlation scale $L_c \approx 4,200$ km and the non-integrable holonomy sector untouched.

Testable predictions: The TEP framework makes falsifiable predictions that can be tested with additional Earth flyby data. Based on the gated fitted β_{fit} values (order 10^{-3}), the model predicts:

- Flybys at perigee altitude < 2000 km should show detectable anomalies (1–10 mm/s)
- Flybys at perigee altitude 2000–3000 km should show marginal anomalies (0.1–5 mm/s)
- Flybys at perigee altitude > 5000 km should show no detectable anomaly (< 0.1 mm/s)

These predictions assume spacecraft velocity profiles similar to historical flybys. Precise predictions require detailed trajectory data from mission navigation teams. Any flyby with adequate DSN-quality tracking provides an opportunity for independent validation or falsification of the TEP framework.

5.11 Addressing the β_{fit} Parameter Scatter

A critical concern for physical interpretation is the formal heterogeneity in gated β (Cochran $Q \approx 1.18 \times 10^3$, reduced $\chi^2 \approx 236$, $I^2 \approx 99.6\%$). In a standard meta-analysis, such an extreme χ^2 would be interpreted as model failure or outlier contamination. In the TEP framework, however, this heterogeneity is *expected* because the clock-sector response amplitude is trajectory-dependent: altitude modulates the field gradient, velocity modulates the disformal coupling, and the trajectory-asymmetry factor $\cos \delta_{\text{in}} - \cos \delta_{\text{out}}$ directly controls how much of the non-radial impulse survives. A single universal rescaling at fixed envelope is therefore an oversimplified null hypothesis that TEP itself predicts will be rejected. The extreme χ^2 is not evidence against TEP; it is evidence that geometry-dependent modulation must be included in any honest forecast. Step 009 does not produce a formal ANOVA-style variance decomposition because $n = 6$ renders percentage-based partitioning statistically underpowered; instead it reports defensible descriptive statistics and rank correlations (Section 4.3).

Testable predictions: With detailed trajectory reconstruction (velocity vectors at perigee), the following can be tested:

- $\beta_{\text{fit}} \propto 1/|v_{\perp}|$ (anticorrelation with perpendicular velocity)
- $\beta_{\text{fit}} \propto |\cos(i)|$ (correlation with equatorial inclination)
- $\beta_{\text{fit}} \propto \cos(\text{latitude})$ (correlation with equatorial perigee)

Within the six-member detection ensemble, Galileo 1990 carries the largest amplitude-informative fitted β_{fit} (1.44×10^{-2}) while NEAR and Rosetta 2005 sit at $4.4\text{--}5.1 \times 10^{-3}$ and Galileo 1992 occupies the negative-asymmetry branch at 6.7×10^{-2} ; Cassini

and MESSENGER carry near-zero reference predictions, so their per-flyby rescalings are unconstrained and contribute negligible pooled weight. Full 3D trajectory integration (Section 4.1.2) validates the perigee approximation for the primary detections and confirms that trajectory curvature contributes only modest corrections ($\leq 20\%$). Raw DSN OD reanalysis remains the key path to test the residual amplitude shortfalls at Cassini and Galileo 1992.

5.12 Model Assumptions and Domain of Validity

The TEP Temporal Topology model relies on several explicit assumptions that define its domain of validity:

Assumption 1: Scalar-tensor gravity framework. The field-amplitude prescription uses the canonical Paper 0 amplitude-sector normalization — a conformally coupled scalar field ϕ whose terrestrial excursion $\psi_{\text{uns}} = 2GM_{\oplus}/(R_{\oplus}c^2) = 1.3926 \times 10^{-9}$ is parameter-free and β_A -independent (Section 3.4). The earlier inverse-power realization ($V(\phi) \propto \Lambda^{4+n}/\phi^n$) is excluded by the Paper 0 F1 constraint under $\beta_A = -1$ and is retained in Step 010 only as an excluded-branch contrast. Alternative amplitude-sector completions would yield different predictions.

Assumption 2: Geometric screening via Temporal Shear suppression. This mechanism requires that Earth develops a continuous spatial profile (Temporal Topology) where the scalar field gradient is suppressed by the environmental-state operator $\mathcal{S}_{\Sigma}(\mathcal{E})$ — the density projection being the relevant coordinate in this near-Earth application. This transition radius is computed from the field equation and depends on the assumed density profile (5515 kg/m³ for Earth interior, 2700 kg/m³ for crust, 1.225 kg/m³ for atmosphere). Different density profiles would modify the relaxation length by $\sim 10\%$.

Assumption 3: Instantaneous coupling. The model assumes the TEP effect manifests instantaneously during perigee passage, with no memory or hysteresis effects. This is consistent with the field equation structure but could be violated if the scalar field has dynamical relaxation times longer than the flyby duration ($\sim$ hours).

Assumption 4: Negligible spacecraft mass. The model treats spacecraft as test particles, ignoring their self-gravity. This is justified as spacecraft masses (~ 500 –5000 kg) are 21 orders of magnitude smaller than Earth mass.

Assumption 5: Spherical Earth symmetry. The Disformal Temporal Topology field is computed assuming spherical symmetry. Earth's oblateness ($J_2 = 1.08 \times 10^{-3}$) introduces $\sim 0.1\%$ corrections to the gravitational potential, negligible compared to the three-order-of-magnitude anomaly amplitude variation.

Domain of validity: The model is valid for flybys with perigee altitudes below the transition region (~ 2500 km) and velocities in the range 10–20 km/s. Extrapolation outside this parameter space requires caution. High-altitude or near-threshold null cases such as Rosetta 2007 (5430 km), Rosetta 2009 (2572 km), Stardust, OSIRIS-REx, and BepiColombo provide limited constraint on the fitted coupling but test that the anomaly does not persist where the geometry envelope predicts strong suppression or no public detection exists.

5.13 Limitations and Caveats

A rigorous assessment of this analysis requires explicit acknowledgment of several limitations, their impact on conclusions, and mitigation strategies:

1. Data provenance and independence:

- *Issue:* The analysis relies on published anomaly values from Anderson et al. (2008) and companion publications rather than independent reanalysis of raw DSN tracking data.
- *Impact:* Systematic errors in the original orbit determination (e.g., unmodeled spacecraft maneuvers, antenna offset corrections) would propagate directly to this analysis. The reported uncertainties (0.01–0.05 mm/s) may not fully capture all systematic contributions.
- *Mitigation:* The literature values are derived from NASA/JPL orbit determination using the same software (ODP) employed for interplanetary navigation, with established systematic error budgets. Cross-validation between independent analyses (JPL vs. ESA/ESOC for Rosetta) shows consistency at the 0.1 mm/s level.
- *Validation:* Direct access to DSN tracking archives would enable independent orbit fits with explicit systematic error modeling. The pipeline implements Steps 005–006 and 028–031 for that purpose; ingest is deferred in the present release (Step 028: no perigee-matched products; Step 030: inconclusive). Re-analysis when data are available is the key validation step.

2. Sample size and selection effects:

- *Issue:* Six flybys enter the Step 008 inverse-variance β_{fit} ensemble (NEAR, Galileo 1990, Rosetta 2005, Cassini, Galileo 1992, MESSENGER), all sign-consistent at β_{ref} on the corrected catalogue. Cassini and MESSENGER carry near-zero reference predictions, so their fitted β_{fit} are poorly constrained and contribute negligible weight; the amplitude-informative core remains four flybys. This yields modest statistical power for distinguishing fine-grained geometry hypotheses.

- *Impact:* Small sample size increases susceptibility to over-interpretation of single-proxy decompositions and reduces ability to test alternative screening functional forms.
- *Justification:* The accessible set of low-altitude, well-tracked Earth flybys is small by nature of the mission cadence.
- *Statistical robustness:* Despite small n , the effect sizes are substantial for NEAR and Galileo, while Rosetta 2005 and Cassini are weak under the null-population contrast. Step 026 on the corrected catalogue places the Anderson empirical law nominally first by BIC, with the restricted TEP model at parity-level separation and all structured models far ahead of the Null (Section 4.5); the BIC map remains small-sample and surrogate. A full-catalog raw stress test on all nine flybys with published observations and explicit TEP predictions improves over the null after random-effects prediction uncertainty is included ($\Delta \log L \approx 43.0$). Leave-one-out shows the pooled β_{fit} is moderately robust, with a stability coefficient of 0.092 across the six members.
- *Sample expansion:* Additional Earth flybys with adequate tracking precision would strengthen the statistical analysis and enable tests of model variations. Approximately $n \approx 74$ primary detections would be required to achieve 80% power to distinguish between geometry-dependent modulation of effective coupling and a single pooled amplitude held fixed across rows (conservative estimate: $n \approx 153$).

3. Trajectory reconstruction uncertainties:

- *Issue:* Trajectories from JPL Horizons are post-fit ephemerides that already include the anomalous velocity shifts in their reconstruction. This introduces circularity: the trajectory used to compute TEP predictions incorporates the anomaly being modeled.
- *Impact:* The perigee altitude and velocity values may have systematic offsets of ~ 1 km and ~ 1 m/s respectively, propagating to $\sim 1\%$ uncertainty in TEP predictions.
- *Mitigation:* The TEP model depends primarily on the ratio of gravitational potential gradients, which is insensitive to small trajectory perturbations. A 1% trajectory error produces $\sim 1\%$ error in predicted Δv , negligible compared to the three-order-of-magnitude amplitude variation between flybys.
- *Previously unavailable flybys:* Rosetta 2007 ($\Delta v = 0.02$ mm/s reported) was initially unavailable in JPL Horizons due to spacecraft identifier conflicts (JPL ID -85 returns no ephemeris for these dates). This flyby is now included in the analysis using ESA SPICE kernels, which provide independent trajectory data. Rosetta 2009 is a published null/bound case ($\Delta v = 0.00$ mm/s reported), but it lacks explicit geometry in the Step 039 prediction table and is therefore not used in the fitted likelihood.

Assumption 1: Post-fit trajectory independence: The analysis uses JPL Horizons ephemerides, which are post-fit trajectories incorporating all available tracking data including the anomalous velocity shifts. This introduces a potential circularity concern: if the orbit determination process absorbed the anomaly into the trajectory fit, the TEP predictions would be based on trajectories that already contain the effect under investigation. However, several factors mitigate this concern:

- *Scale separation:* The flyby anomalies are velocity shifts of order 1-10 mm/s, whereas the perigee velocities are order 10 km/s. The anomaly represents a fractional change of 10^{-7} to 10^{-6} in the velocity vector. Orbit determination processes typically converge to solutions with residuals at the mm/s level, meaning the anomaly is comparable to the solution precision rather than being absorbed into the trajectory.
- *Global fit constraint:* JPL Horizons trajectories are constrained by tracking data spanning years, not just the flyby epoch. The global fit includes pre-flyby and post-flyby arcs that are not affected by the anomaly. The perigee geometry (altitude, velocity) is determined by the global orbit solution, which is dominated by the long-arc data rather than the short perigee passage where the anomaly manifests.
- *Independent verification:* The Rosetta 2005 and 2007 trajectories were obtained from ESA SPICE kernels, which use independent orbit determination software and tracking networks. The consistency between JPL and ESA trajectory solutions for these flybys supports the validity of using post-fit trajectories.
- *Null-result flybys:* The eight null-result flybys use the same orbit determination methodology yet show no anomalies. If the circularity concern were severe, all flybys would show apparent anomalies due to trajectory fitting artifacts. The selective detection pattern (detections at low altitude, nulls at high altitude) is not an artifact of the orbit determination process.

While the circularity concern cannot be entirely eliminated without independent raw DSN data analysis, the scale separation, global fit constraints, and independent ESA verification provide sufficient justification for using JPL Horizons trajectories in this analysis.

4. Phenomenological gradient suppression model:

- *Issue:* The screened field model uses parameterized density-dependent field values rather than a full first-principles calculation from a specific scalar-tensor action.
- *Issue:* The screened field model anchors the terrestrial excursion on the canonical amplitude-sector normalization (ψ_{uns} , Section 3.4) with a phenomenological relaxation profile, rather than a full first-principles profile calculation from the microscopic scalar action.
- *Mitigation:* The canonical normalization is parameter-free — fixed by Earth's measured mass and radius and independent of β_A — and the excluded inverse-power branch is retained only as a contrast diagnostic (Step 010). One free phenomenological parameter remains (β_{fit} , a trajectory-response amplitude distinct from the bare coupling $\beta_A = -1$), preserving predictive power.

- *Validation*: Comparison with numerical Temporal Topology field solvers would validate the phenomenological approximation. Additionally, a plasma-model sensitivity analysis (scripts/utis/plasma_screening.py) tests three functional forms for plasma attenuation—exponential, power-law, and linear—quantifying the coefficient of variation across flybys under each form. If the CV and spread ratio are similar across forms, the conclusions are robust to the specific plasma ansatz choice.

5. Systematic error budget:

- *DSN measurement systematics*: Antenna phase center motion (~ 0.1 mm/s), tropospheric delay modeling (~ 0.05 mm/s), and station position errors (~ 0.02 mm/s) contribute to the anomaly uncertainty budget. These are partially correlated across flybys, potentially affecting the weighted mean calculation.
- *Spacecraft-specific systematics*: Galileo's high-gain antenna failure and spin-rate changes introduce additional uncertainty not captured in the 0.08 mm/s formal error. The Galileo 1990 anomaly should be interpreted with caution. This is a caveat, not an exclusion criterion. Galileo 1990 satisfies both a priori ensemble gates ($S/N = 49 > 2$; sign agreement with TEP prediction) and is retained in the fitted ensemble. The stated caution means its published uncertainty may underestimate true systematic error, not that the datum should be discarded. Arbitrary exclusion would reduce the ensemble from $n = 6$ to $n = 5$ without methodological justification and would violate the pre-specified selection protocol (Section 3.3). The leave-one-out diagnostic confirms that removing Galileo 1990 shifts the pooled β by only 3.5% (4.58×10^{-3} vs 4.74×10^{-3}), so the headline conclusion is insensitive to its inclusion.
- *Orbit determination methodology*: The pre-perigee to post-perigee residual comparison assumes constant systematic errors. Time-varying systematics (e.g., thermal expansion) could produce spurious velocity signatures.

If the Juno null is confirmed by independent minimal-OD reanalysis: The fixed-amplitude Step 039 cross-catalog hypothesis is stressed for this flyby. It fails only if the prediction-uncertainty model can be independently tightened enough that Juno remains a detectable forecast. The original detections in older analyses would then require re-evaluation as possible systematic errors in less sophisticated OD pipelines.

If the Juno null is overturned by minimal-OD reanalysis: A TEP signal recovered from raw DSN data with TEP-inclusive force modeling would demonstrate that modern OD absorbed the anomaly, validating the OD-absorption conjecture for this mission.

Current status: A statistical power and uncertainty analysis (Step 039b, `step039b_juno_falsification_bounds.json`; `juno_falsification_power_analysis`) quantifies the Juno comparison across uncertainty models: at the Step 008 pooled $\beta = 4.74 \times 10^{-3}$ with the mean- β standard error ($\sigma_{\text{pred,SE}} \approx 0.034$ mm/s) and DSN measurement precision ($\sigma_{\text{obs}} = 0.020$ mm/s), the combined $\sigma_{\text{total}} = 0.039$ mm/s gives 3.17σ (Cohen's $d \approx 3.2$) for the raw $+0.124$ mm/s prediction; when the Step 008 random-effects prediction uncertainty ($\sigma_{\text{pred,RE}} \approx 0.067$ mm/s, capturing cross-flyby geometry scatter τ) is included, the combined $\sigma_{\text{total}} = 0.070$ mm/s reduces the raw difference to 1.77σ ($< 2\sigma$, not falsified at the 3σ threshold), even prior to accounting for Juno's perigee safe-mode tracking interruption and modern stochastic-acceleration OD absorption (which absorbs $\sim 97\%$ of a perigee impulse in Step 012/043). A definitive post-OD test requires MONTE- or ODP-class minimal orbit determination without empirical stochastic accelerations.

5.14 Evidence Ledger for Review

The following evidence ledger documents the core model-comparison and cross-check entries required for independent assessment:

- Restricted TEP vs Null: large information-criterion separation ($\Delta\text{BIC} \approx 16690$ on full $n = 9$ catalog under the fractional misspecification floor); TEP restricted is strongly favoured over the Null, while the Anderson empirical law fitted on the same catalogue attains the nominally smallest BIC overall.
- Restricted TEP vs Anderson: the Anderson empirical model attains the nominally smaller BIC ($\Delta\text{BIC} \approx -4.8$ on full $n = 9$ catalog — parity-level separation, not a decisive preference; under the published-uncertainties-only sensitivity the ordering reverses, $\Delta\text{BIC} \approx +2767$ in TEP restricted's favour). The restricted model retains the pre-specified-parameter advantage.
- External anchors: GNSS atomic-clock correlation anchors $\lambda_T = 4200$ km under the Paper 6 §2 identification (Step 016 records the bookkeeping consistency of the adopted constant with its own anchor); UCD saturation model gives $S_{\oplus} \approx 0.35$ (Step 010); Cassini solar conjunction constrains $|\gamma - 1|$ (Section 4.6).
- Null-catalog stress test: full $n = 9$ raw-catalog likelihood improves by $\Delta \log L \approx 43.0$ over the null after random-effects prediction uncertainty is included (Step 039).
- Failure-mode accounting: velocity-template regime classification via the empirical scale $v_{\text{trans}} \approx 16.8$ km/s (phenomenological ansatz, §3.5) and sign-rule consistency checks (Step 007/017).
- Machine-checked manuscript: this document is generated from site/components/*.html and audited against pipeline JSON outputs by Step 027.

5.15 Summary

Within these explicitly defined scope boundaries—link-class conditioning, literature-derived Δv inputs, and finite-sample heterogeneity—TEP with Temporal Shear suppression within continuous Temporal Topology organizes the Earth flyby anomaly with a single fitted amplitude on pre-specified geometry and defines falsifiable predictions for future clock-carrying planetary encounters.

6. Conclusions

This study investigated whether the Temporal Equivalence Principle (TEP), through the clock-sector reduction-artefact channel identified in Paper 0 (v0.16) and scoped to its observable classes in Section 5.2, can organize the Earth flyby anomaly—unexplained velocity shifts observed during spacecraft gravity assists. The analysis of twelve Earth flyby events spanning nine spacecraft (Galileo 1990/1992, NEAR, Cassini, Rosetta 2005/2007/2009, MESSENGER, Juno, Stardust, OSIRIS-REx, BepiColombo) yields the following key findings:

1. Six-member TEP ensemble: Inverse-variance fitting (Step 008) enters NEAR (13.46 ± 0.13 mm/s), Galileo 1990 (3.92 ± 0.08 mm/s), Rosetta 2005 (1.80 ± 0.03 mm/s), Cassini (-2.00 ± 1.00 mm/s), Galileo 1992 (-4.60 ± 1.00 mm/s), and MESSENGER ($+0.02 \pm 0.01$ mm/s) under the $S/N \geq 2$ pre-fit admission gate (with `strict_sign_gate = false` in the active pipeline configuration). Evaluated on the ungated catalogue prior to S/N filtering, all 7 of 7 published nonzero anomalies (including Rosetta 2007; one-sided binomial $p = 0.0078$) agree in sign with the reference prediction at $\beta_{\text{ref}} = 10^{-4}$, so the six S/N-qualified detections also satisfy sign concordance without selection circularity. At β_{ref} , Rosetta 2005 predicts $\Delta v_{\text{TEP}} \approx 0.03$ mm/s vs 1.80 mm/s observed; the full-model fit yields $\beta_{\text{fitted}} \approx 5.15 \times 10^{-3}$. Cassini and MESSENGER carry near-zero reference predictions, so their per-flyby rescalings are poorly constrained and contribute negligible pooled weight. The amplitude-informative gated β_{fit} values span roughly a factor of 3.3 (4.43×10^{-3} to 1.44×10^{-2} , with Galileo 1992 at 6.7×10^{-2}), consistent with geometry-dependent modulation (the plasma term is a bounded heuristic envelope coefficient; the stronger Debye attenuation ansatz is quarantined at $S_{\text{plasma}} = 1$, Section 3.10). The flyby trajectory-response amplitudes β_{fit} are distinct from the bare conformal coupling $\beta_A = -1.0$ and from the Cassini source charge $\alpha_{\text{eff}}^{(\odot)} = S_{\Sigma}^{(\odot)} \alpha_0$. Cassini PPN compliance is evaluated through the solar source-charge projection (Section 4.6.1a), not through the fitted flyby amplitudes.
2. TEP parameter estimate: The inverse-variance weighted mean $\beta_{\text{fit}} \approx 4.74 \times 10^{-3} \pm 1.28 \times 10^{-3}$ (random-effects standard error; formal inverse-variance error 5.83×10^{-4}) summarizes the gated ensemble. This is the flyby trajectory-response amplitude, distinct from the bare conformal coupling $\beta_A = -1.0$ and from the Cassini source charge α_{eff} . Formal heterogeneity (reduced $\chi^2 \approx 236$, $I^2 \approx 99.6\%$) signals that a single rescaling does not exhaust the geometry physics carried by the envelope (the plasma term is a bounded heuristic envelope coefficient). Bootstrap resampling (10^4 draws) and leave-one-out recomputations in Step 008 show moderate stability (coefficient ≈ 0.092), with NEAR as the dominant lever.
3. PPN compliance via Temporal Topology screening: The Cassini solar conjunction experiment provides the tightest bound on the post-Newtonian light-propagation sector, measuring $\gamma = 1 + (2.1 \pm 2.3) \times 10^{-5}$. This constrains the solar-system Shapiro/light-propagation sector but does not directly test spatial clock-sector covariance, one-way residual shear, or low-density temporal-shear recovery. The fitted β_{fit} values are flyby trajectory-response amplitudes, not the Cassini source charge $\alpha_{\text{eff}}^{(\odot)} = S_{\Sigma}^{(\odot)} \alpha_0$. Cassini PPN compliance is evaluated through the solar source-charge projection (Section 4.6.1a): the adopted Paper 0 v0.16 convention gives $S_{\Sigma}^{(\odot)} \sim 2 \times 10^{-6}$, inside the Cassini requirement $S_{\Sigma}^{(\odot)} \lesssim 1.15 \times 10^{-5}$ under the unified unstretched-space metric (5.8×10^{-6} isotropic, Paper 0 §7), with the kinetic-operator benchmark evaluating the solar-field screening at $\sim 10^{-15}$ along the 1 AU path. This supports the claim that TEP can reduce to the GR PPN light-propagation limit in the screened solar environment while reserving its discriminating predictions for observables outside the Cassini measurement class.
4. TEP suppression by modern orbit determination: Analysis of the expanded dataset reveals that published null results (MESSENGER, Rosetta 2007) are consistent with universal- β_{fit} null predictions, while Juno is the sole raw-tension case at universal β_{fit} . Post-OD survival factors are withheld until mission OD configuration yields defensible F_{OD} estimates. Rosetta 2009 is treated as a published null/bound case with insufficient explicit geometry for the universal- β_{fit} prediction table; Stardust, OSIRIS-REx, and BepiColombo have no public anomaly report and are not used in quantitative likelihood.
5. Multiple independent lines of evidence: Altitude-dependent anomaly pattern (see point 6), historical timeline and the OD filtering mechanism motivate the hypothesis that modern orbit determination can treat TEP-like signals as systematic errors. However, mission-specific survival factors are not computed in this paper: Step 021 withholds F_{OD} until real OD configuration data are available, and the synthetic Step 012 OD run is quarantined from manuscript inference.
6. Temporal Topology screening support: The response template organizes the catalogue pattern — null results on high-altitude flybys, where the conformal excursion is shallow, and large anomalies on low-altitude encounters — conditional on observable class: in the coherent two-way class of the published catalogue the organization measures how strongly the recorded anomalies track the trajectory geometry the artefact kernel encodes, while the mechanism-level signature is confined to clock-carrying links (Section 5.2). Altitude alone is not the relevant TEP discriminator; the pattern is governed by the full trajectory-geometry response, which combines altitude, asymmetry, velocity and field structure.
7. Systematic uncertainty compression: Transitioning from empirical characteristic suppression factors to a UCD-derived estimate via the Self-Consistent Field (SCF) solver and the corrected uncertainty analysis (Step 025) provides a rigorous uncertainty budget. The total relative uncertainty is 152.7%, with heterogeneity contributing 149.4% and systematic uncertainty contributing 29.2%. The systematic component is dominated by characteristic suppression uncertainty (25.0%, Paper 6 UCD)

and relaxation length uncertainty (15.0%, SCF theoretical prior). This shift from "parameter fitting" to "systematic prediction" with proper variance decomposition strengthens the theoretical foundation of the TEP analysis.

8. Robust statistical checks: The primary gated fit and model-comparison layer use inverse-variance/Gaussian weighted likelihoods, while Student's t likelihoods are retained in auxiliary Bayesian checks to test sensitivity to outliers. Residual diagnostics on the gated ensemble is consistent with normality (Shapiro–Wilk $p \approx 0.95$, $n = 6$), so heterogeneity diagnostics, not normality alone, dominate the statistical interpretation. (NEAR, Galileo 1990, Rosetta 2005, Cassini, Galileo 1992, MESSENGER).
9. Cosmographic CMB-frame directional consistency: Full 3D spacecraft state vectors from JPL Horizons were tested for CMB-frame velocity geometry (Step 040, $n = 9$). The exploratory both-aligned flag yields Pearson $r = -0.57$, $p = 0.11$, and Mann-Whitney $U = 7.5$, $p = 0.79$. Multivariate regression on residual ratio achieves $R^2 = 0.96$ (adjusted $R^2 = 0.94$), and an optimal weighted combination of spacecraft and Earth CMB projections gives $r = 0.96$, though the optimum sits at the boundary of the scanned weight grid ($w = -2.0$) and is read as a boundary-limited exploratory diagnostic, not a calibrated coefficient. These results are sample-limited and remain exploratory.

Significance

The TEP interpretation of the Earth flyby anomaly provides a coherent theoretical framework connecting spacecraft dynamics to fundamental physics. The flyby trajectory-response amplitude $\beta_{\text{fit}} \approx 4.74 \times 10^{-3}$ (weighted mean) is a trajectory-response coefficient, not the bare coupling β_A or the Cassini source charge α_{eff} . Interpreted as the clock-sector amplitude response under the amplitude–shear split, the flyby catalogue's geometric organization is consistent with solar system constraints and organizes the anomalous velocity-shift pattern. What that organization attests is conditional on the tracking-link class: the published catalogue is assigned to the coherent-transponder class on the literature-based ledger of Section 5.2, pending encounter-specific tracking verification; in that class the clock-sector term is absent at the observable (Section 5.2), so the catalogue pattern constrains the response template's structure while the mechanism's claim is carried by the falsifiable prediction of an m/s-scale signature on the first clock-carrying flyby pass.

Unlike ad hoc modifications to gravity, the TEP framework preserves all successes of general relativity in solar system tests while supplying a structured account of the anomaly pattern in the specific regime of planetary gravity assists — an account whose mechanism-level claim is scoped to the clock-carrying link classes of Section 5.2 rather than asserted for the recorded catalogue. The geometric amplitude response, calibrated by independent UCD saturation analysis, is essential for the flyby amplitude and environmental response: without it, the required β_{fit} would not produce the observed trajectory-dependent residuals.

Template-organization metrics: The quantitative validation summarizes how the restricted geometric response template organizes the published catalogue (measuring empirical template compression rather than mechanism-level confirmation in the coherent two-way observable):

- Effect sizes: Cohen's d relative to the published null population ($n_{\text{null}} = 3$) yields a large effect for NEAR ($d \approx 2.5$), medium for Galileo 1990 ($d \approx 0.73$), large-in-magnitude negative for Galileo 1992 ($d \approx -0.87$), and small effects for Rosetta 2005 ($d \approx 0.34$) and Cassini ($d \approx -0.38$). The coefficient of variation $CV \approx 149\%$ across the six gated β_{fit} fits is inflated by the near-zero-prediction members and reflects genuine geometry-dependent modulation at fixed envelope.
- Model comparison: On the full $n = 9$ catalog under the fractional misspecification floor, Step 026 gives a large information-criterion separation between TEP restricted and Null ($\Delta\text{BIC} \approx 16690$); the Anderson empirical law, whose two parameters are fitted on the same catalogue, attains the nominally smallest BIC overall but only at parity-level separation ($\Delta\text{BIC} \approx 4.8$), an ordering that reverses in TEP restricted's favour under the published-uncertainties-only sensitivity
- Akaike model weight favours the Anderson empirical tier (≈ 0.86 versus ≈ 0.14 for TEP restricted, full $n = 9$ catalogue under the fractional misspecification floor) — a parity-level ordering that reverses under the published-uncertainties sensitivity; the TEP restricted and flexible tiers remain far above the Null and are documented in Step 026 outputs.
- Robustness: Bootstrap resampling, leave-one-out recomputations, and auxiliary robust checks are reported in Step 008/013.
- Prediction accuracy: Step 008 reports $R^2 \approx 0.85$ (correlation $r \approx 0.93$) between pooled- β predicted and observed anomalies across the six-member ensemble; prediction intervals are bootstrap-derived.
- Residual analysis: Shapiro–Wilk $p \approx 0.95$ on the gated six-member ensemble; heterogeneity diagnostics dominate formal tests.
- Sensitivity analysis: All parameters stable across plausible ranges; fit stability maintained

The catalog of Earth flyby events (six published nonzero entries — four firm detections and two $\sim 2\sigma$ marginal entries, Cassini and MESSENGER — plus three published nulls/bounds and several without public anomaly reports) provides the empirical substrate. Step 026's headline likelihood uses the full $n = 9$ catalog with a fractional misspecification floor ($\sigma_{\text{sys},i} = CV_{\beta}|\hat{y}_i|$, $CV_{\beta} \approx 0.47$, measured from the catalogue's own amplitude heterogeneity and invariant under the field-normalization anchor). The archived comparison gives $\Delta\text{BIC} \approx 1.67 \times 10^4$ for the restricted geometric template relative to Null and near-parity with the catalogue-fitted Anderson tier. These scores describe catalogue organization under the documented fitting convention, not support for the clock-sector mechanism. Section 4.5 states the outstanding likelihood-optimum and nuisance-accounting checks; the physical clock-sector test is a clock-carrying pass, with the profile independently constrained by clock observations.

Robustness Assessment

Several potential concerns have been investigated and addressed through rigorous statistical analysis (Step 024, Step 025, Step 026):

Systematic error discrimination: trajectory asymmetry ($\cos \delta_{\text{in}} - \cos \delta_{\text{out}}$) organizes the adopted catalogue template, and while the hardware and environmental systematics explicitly modelled in Step 024 (0.12 mm/s combined) do not reproduce this pattern, this correlation does not by itself exclude unmodeled geometry-dependent tracking or orbit-determination reconstruction errors pending raw DSN reanalysis. Within the gated $n = 6$ ensemble the ordering is positive but imperfect (Spearman $\rho = +0.66$, $p = 0.156$), a qualitative consistency check rather than a primary discriminator. On the ungated re-transcribed catalogue, all 7 of 7 published nonzero anomalies (and all 6 $S/N \geq 2$ entries) agree in sign with the reference prediction, while the geometry envelope carries seven nominal deterministic coefficients plus one fitted amplitude against six S/N -qualified anomalies, so the archived scores characterize template organization under the adopted convention, not confirmation of the clock-sector mechanism in the coherent two-way observable. The broader catalogue rank correlations are complementary diagnostics on overlapping data: asymmetry-vs-magnitude gives $\rho = 0.56$ ($p = 0.11$, $n = 9$) and $\rho = 0.68$ ($p = 0.090$, $n = 7$ nonzero anomalies) in the Step 009 ordering diagnostic. See Section 5.5 for comprehensive systematic uncertainty budget.

Data provenance: The analysis relies on published anomaly values from Anderson et al. (2008) rather than independent DSN reanalysis. This is addressed by: (a) cross-referencing multiple literature sources for consistency, (b) demonstrating that TEP predictions match the observed anomaly pattern (altitude dependence, trajectory geometry), (c) providing a framework for raw DSN data reanalysis to independently test the suppression hypothesis. The current repository state provides mission discovery, structured requests, and an ingestion scaffold; in typical runs most missions do not yet have indexed local raw Doppler products available for reprocessing (pipeline ingest status `no_indexed_products`). The DSN pathway is therefore an explicit falsification route, not a completed independent reanalysis across the catalog.

β_{fit} scatter as physical modulation: Fitted $\beta_{A,\text{eff}}$ spans roughly a factor of 3.3 across the tightly constrained members (4.43×10^{-3} to 1.44×10^{-2}), extending to 6.7×10^{-2} for Galileo 1992 and formally divergent for the near-zero-prediction Cassini and MESSENGER members (Step 009), reflecting geometry-dependent modulation: altitude (J_2 gradient suppression), perigee latitude (inclination-dependent coupling), and velocity (the response template's disformal branch); the plasma term is a bounded heuristic envelope coefficient (Section 3.10). Step 009 does not report a formal variance decomposition because $n = 6$ renders ANOVA-style partitioning statistically underpowered. Instead, it reports deterministic geometry modulation metrics: beta scatter statistics, full-catalog detection-pattern classification, and Spearman rank correlation ($\rho = 0.56$ over the $n = 9$ catalogue; $\rho = +0.66$ in the gated ensemble). Leave-one-out recomputation on the six-member ensemble reports stability coefficient ≈ 0.092 . The UCD-derived characteristic suppression $S_{\oplus} \approx 0.35$ provides a cross-scale prior. See Section 4.3 for the quantitative geometry modulation assessment and Section 5.5 for detailed interpretation.

Cassini status: At $\beta_{\text{ref}} = 10^{-4}$ the Step 007 envelope yields a negative Δv_{TEP} of negligible magnitude, and the corrected published anomaly is -2.00 ± 1.00 mm/s, so Cassini is sign-consistent at β_{ref} and enters the gated ensemble ($n = 6$) — its residual tension is amplitude-only. Galileo 1992, transcribed as a null in the earlier catalogue, is a genuine negative detection (-4.60 ± 1.00 mm/s) whose magnitude is underpredicted (≈ -0.33 mm/s at the pooled amplitude). Addressing these amplitude shortfalls requires independent DSN OD or envelope refinements, not only a universal rescaling of β_{fit} .

Juno falsification pressure: At the refit weighted-mean β_{fit} , Step 039 predicts $\Delta v_{\text{TEP}}^{\text{raw}} \approx +0.12$ mm/s with uncertainty ~ 0.03 mm/s while the published bound is 0.00 ± 0.02 mm/s. This raw-tension case is the headline stress test on universal β_{fit} before mission-specific OD survival factors are supplied (Step 021). Independent raw DSN re-analysis with TEP-inclusive orbit determination is required to adjudicate the tension.

Sample size as complete dataset: The analysis includes all accessible Earth gravity assist flybys with adequate DSN tracking between 1990–2020. The rarity of suitable flyby events (low altitude, Doppler tracking, no major maneuvers) means only a handful of literature anomalies exist; six enter the gated β_{fit} ensemble. Step 008 bootstrap 95% CI [4.42×10^{-3} , 1.47×10^{-2}] brackets the gated fits (bootstrap median 4.77×10^{-3}). Additional flybys would test envelope refinements rather than restate the same $n = 6$ compression.

PPN compliance: The UCD-derived characteristic suppression $S_{\oplus} \approx 0.35$ is determined from the UCD saturation model. The flyby trajectory-response amplitudes β_{fit} are distinct from the Cassini source charge $\alpha_{\text{eff}}^{(\odot)} = S_{\Sigma}^{(\odot)} \alpha_0$; Cassini PPN compliance is evaluated through the solar source-charge projection (Section 4.6.1a), where the adopted Paper 0 v0.16 convention gives $S_{\Sigma}^{(\odot)} \sim 2 \times 10^{-6}$, inside the Cassini requirement $S_{\Sigma}^{(\odot)} \lesssim 1.15 \times 10^{-5}$ under the unified unstretched-space metric (5.8×10^{-6} on the isotropic-conformal benchmark).

Model comparison: Step 026 on the full $n = 9$ catalog gives a large information-criterion separation between TEP restricted and Null ($\Delta \text{BIC} \approx 16690$). The Anderson empirical model, fitted on the same catalogue, attains the nominally smallest BIC overall ($\Delta \text{BIC} \approx 4.8$ ahead of TEP restricted — parity-level; $\Delta \text{BIC} \approx 16694$ ahead of Null), confirming that trajectory asymmetry carries genuine signal. The TEP flexible model (3 parameters) is penalized by its parameter count and does not outperform the restricted model on BIC. Residual diagnostics on the gated ensemble are consistent with normality (Shapiro–Wilk $p \approx 0.95$, $n = 6$).

Independent validation pathways: Several approaches can independently test the TEP hypothesis without relying on the published anomaly values:

- Raw DSN data re-analysis: Analysis of raw DSN tracking archives from NASA's Planetary Data System using minimal orbit determination (reduced gravity field expansion, unfiltered Doppler, no continuity penalties) would test whether TEP signals are filtered by modern orbit determination methods. This would provide an important test of the suppression hypothesis.
- Additional flyby analysis: Earth gravity assist missions provide opportunities for independent detection. Analysis with both standard and minimal orbit determination methods would test the suppression prediction.
- GNSS clock correlation: GNSS atomic clock correlation analysis provides a cross-consistency check with the GNSS/UCD terrestrial calibration on the transition radius ($R_{\text{sol}} \approx 4200$ km). This shared-input dependency is disclosed explicitly to avoid circular claims of independent validation.
- Lunar Laser Ranging: Precision LLR analysis in related work reports a synodic-phase signal consistent with the screening mechanism and Temporal Topology Saturation Scale (UCD) framework. Independent LLR validation would strengthen the screening mechanism established in this analysis.

Data Availability

Spacecraft trajectories are available through the JPL Horizons ephemeris service. Literature anomaly values are from Anderson et al. (2008) and companion publications. Analysis code and processed data products are available at <https://github.com/matthewsmawfield/TEP-EFA> with archived DOI at 10.5281/zenodo.19454862.

Acknowledgments

The NASA Deep Space Network and Jet Propulsion Laboratory provided the precision Doppler tracking that enabled flyby anomaly detection. The JPL Horizons system provided trajectory reconstruction. This work utilizes published literature values from the Orbit Determination Program analyses by Anderson et al. and collaborators. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. The author declares no conflicts of interest.

Additional Considerations

Several avenues for extending this analysis are identified:

Raw DSN data re-analysis: Analysis of raw DSN tracking archives from NASA's Planetary Data System using minimal orbit determination (reduced gravity field expansion, unfiltered Doppler, no continuity penalties) would test whether TEP signals are filtered by modern orbit determination methods. This provides an important test of the suppression hypothesis.

Extended spacecraft sample: Additional flyby events would increase the sample size beyond the current six published nonzero entries (all six sign-consistent in the inverse-variance β_{fit} ensemble, with Cassini and MESSENGER marginal at $\sim 2\sigma$). A sample of $n \approx 74$ primary detections would provide sufficient statistical power to distinguish between geometry-dependent modulation of β_{fit} and a single universal coupling constant at 80% power (conservative estimate: $n \approx 153$).

- Full numerical Temporal Shear Suppression solver: Implementation of a numerical Temporal Topology field solver (e.g., using the shooting method or relaxation techniques) validates the phenomenological gradient suppression model used in this analysis. This enables prediction of the Temporal Topology profile without the phase-boundary approximation and could explain the observed β_{fit} scatter through detailed density-dependent effects.
- Inclination-dependent modeling: Incorporation of Earth's oblateness (J_2) and latitude-dependent density variations into the TEP model could explain part of the observed β_{fit} scatter. Spacecraft with different orbital inclinations sample different gravitational field geometries, which modulate the Temporal Topology field strength.
- Disformal coupling exploration: Extension to scalar-tensor theories with disformal coupling terms introduces velocity-dependent gradient suppression that could explain the correlation between fitted β_{fit} and flyby velocity. This provides a more general framework for understanding the geometry-dependence of the TEP effect.
- Local-time plasma effects: A first-principles scalar–plasma coupling derived from the TEP action would test whether ionospheric density variations with local time modulate the clock-sector response. At present plasma enters only as a bounded heuristic envelope coefficient, with the stronger Debye attenuation ansatz quarantined at $S_{\text{plasma}} = 1$ (Section 3.10); this bullet records the question as an open refinement, not a current mechanism.

References

1. Anderson, J. D., Campbell, J. K., Ekelund, J. E., Ellis, J., & Jordan, J. F. 2008, "Anomalous Orbital-Energy Changes Observed during Spacecraft Flybys of Earth," *Phys. Rev. Lett.*, 100, 091102
2. Anderson, J. D., & Nieto, M. M. 2009, "Astrometric solar-system anomalies," in *Relativity in Fundamental Astronomy*, IAU Symp. 261, 189
3. Antreasian, P. G., & Guinn, J. R. 1998, "Investigations into the Unexpected Delta-V during the Earth Gravity Assist of NEAR," Paper AAS 98-428
4. Bertotti, B., Iess, L., & Tortora, P. 2003, "A test of general relativity using radio links with the Cassini spacecraft," *Nature*, 425, 374
5. Brax, P., van de Bruck, C., Davis, A.-C., Khoury, J., & Weltman, A. 2004, "Detecting dark energy in orbit: The cosmological chameleon," *Phys. Rev. Lett.*, 93, 200405
6. Einstein, A. 1915, "Die Feldgleichungen der Gravitation," *Sitzungsberichte der Preussischen Akademie der Wissenschaften*, 844
7. Halsey, D., et al. 2012, "Anomalous Earth flybys: Status and developments," *Adv. Space Res.*, 50, 362
8. Khoury, J., & Weltman, A. 2004, "Chameleon cosmology," *Phys. Rev. D*, 69, 044026
9. Lämmerzahl, C., Preuss, O., & Dittus, H. 2006, "Is the physics within the Solar system understood?" in *Lasers, Clocks and Drag-Free Control*, 75, 75
10. McCulloch, M. E. 2008, "Modelling the Pioneer anomaly as modified inertia," *MNRAS*, 389, L57
11. Meeus, J. 1998, *Astronomical Algorithms*, 2nd edn. (Richmond: Willmann-Bell)
12. Mota, D. F., & Shaw, D. J. 2007, "Strongly coupled chameleon fields," *Phys. Rev. Lett.*, 97, 151102
13. Nieto, M. M., & Anderson, J. D. 2007, "Search for a solution of the Pioneer anomaly," *Contemp. Phys.*, 48, 41
14. Page, G., & McCulloch, M. E. 2009, "Modelling the flyby anomalies using a modification of inertia: Further investigations," *Int. J. Astron. Astrophys.*, 3, 1
15. Schive, H.-Y., Chiueh, T., & Broadhurst, T. 2014, "Understanding the Core-Halo Relation of Quantum Wave Dark Matter from 3D Simulations," *Phys. Rev. Lett.*, 113, 261302
16. Turyshv, S. G., & Toth, V. T. 2010, "The Pioneer anomaly," *Living Rev. Relativ.*, 13, 4
17. Will, C. M. 2014, "The confrontation between general relativity and experiment," *Living Rev. Relativ.*, 17, 4
18. Folkner, W. M., et al. 2022, "Planetary ephemeris DE440," *IPN Progress Report*, 42-284, 1
19. JPL Horizons, "NASA/JPL Horizons System" <https://ssd.jpl.nasa.gov/horizons/> (accessed 2024)
20. Morley, T., & Budnik, F. 2007, "Rosetta Navigation at its First Earth-Swingby," *Proceedings of the 20th International Symposium on Space Flight Dynamics*
21. Müller, J., Soffel, M., & Klioner, S. A. 2008, "Geodesy and relativity," *Journal of Geodesy*, 82, 133
22. Müller, J., et al. 2010, "Relativistic models for spacecraft tracking," *Acta Astronautica*, 67, 975
23. Aksenov, E. L., & Tuchin, A. G. 2020, "Earth flyby anomalies and the general relativistic theory of the Kerr gravitational field," *MNRAS*, 492, 3703
24. Ciufolini, I., & Pavlis, E. C. 2004, "A confirmation of the general relativistic prediction of the Lense-Thirring effect," *Nature*, 431, 958
25. IERS Conventions 2010, IERS Technical Note No. 36, eds. Petit, G. & Luzum, B.
26. Brax, P., & Burrage, C. 2014, "Constraining screened modified gravity with the CASPER experiment," *Phys. Rev. D*, 90, 104009
27. Lemoine, F. G., et al. 1998, "The Development of the NASA GSFC and NIMA Joint Geopotential Model," in *Proceedings of the International Symposium on Gravity, Geoid, and Marine Geodesy*, Tokyo, Japan
28. Pavlis, N. K., et al. 2012, "The development and evaluation of the Earth Gravitational Model 2008 (EGM2008)," *J. Geophys. Res.*, 117, B04406
29. Mocz, P., Vogelsberger, M., Robles, V., et al. 2018, "Galaxy Halos from Fuzzy Dark Matter," *Phys. Rev. Lett.*, 121, 141102
30. Moyer, T. D. 2000, *Formulation for Observed and Computed Values of Deep Space Network Data Types*, JPL Publication 00-7
31. Burrage, C., & Sakstein, J. 2016, "Tests of Ambient Symmetry Restoration," *Living Rev. Relativ.*, 21, 1
32. Upadhye, A., Hu, W., & Khoury, J. 2012, "Quantum stability of chameleon field theories," *Phys. Rev. Lett.*, 109, 041301
33. Joyce, A., Jain, B., Khoury, J., & Trodden, M. 2015, "Beyond the cosmological standard model," *Phys. Rept.*, 568, 1
34. Kass, R. E., & Raftery, A. E. 1995, "Bayes Factors," *J. Am. Stat. Assoc.*, 90, 773
35. Clifton, T., Ferreira, P. G., Padilla, A., & Skordis, C. 2012, "Modified gravity and cosmology," *Phys. Rept.*, 513, 1
36. Higgins, J. P., & Thompson, S. G. 2002, "Quantifying heterogeneity in a meta-analysis," *Stat. Med.*, 21, 1539

TEP Research Series

Smawfield, M. L. *Temporal Equivalence Principle: Dynamic Time & Emergent Light Speed*. Preprint v0.16 (Jakarta). Zenodo. DOI: 10.5281/zenodo.16921911

Smawfield, M. L. *Global Time Echoes: Distance-Structured Correlations in GNSS Clocks*. Preprint v0.28 (Jaipur). Zenodo. DOI: 10.5281/zenodo.17127229

Smawfield, M. L. *Global Time Echoes: 25-Year Analysis of CODE Precise Clock Products*. Preprint v0.21 (Cairo). Zenodo. DOI: 10.5281/zenodo.17517141

Smawfield, M. L. *Global Time Echoes: Raw RINEX Consistency Test*. Preprint v0.9 (Kathmandu). Zenodo. DOI: 10.5281/zenodo.17860166

Smawfield, M. L. *Temporal-Spatial Coupling in Gravitational Lensing: A Reinterpretation of Dark Matter Observations*. Preprint v0.9 (Tortola). Zenodo. DOI: 10.5281/zenodo.17982540

Smawfield, M. L. *Global Time Echoes: Empirical Synthesis*. Preprint v0.8 (Singapore). Zenodo. DOI: 10.5281/zenodo.18004832

Smawfield, M. L. *Temporal Topology Saturation Scale: Cross-Scale Consistency of ρ_T* . Preprint v0.9 (New Delhi). Zenodo. DOI: 10.5281/zenodo.18064365

Smawfield, M. L. *The Soliton Wake: Exploring RBH-1 as a Temporal Topology Candidate*. Preprint v0.5 (Blantyre). Zenodo. DOI: 10.5281/zenodo.18059250

Smawfield, M. L. *Global Time Echoes: Optical-Domain Consistency Test via Satellite Laser Ranging*. Preprint v0.7 (Mombasa). Zenodo. DOI: 10.5281/zenodo.18064581

Smawfield, M. L. *What Do Precision Tests of General Relativity Actually Measure?*. Preprint v0.9 (Istanbul). Zenodo. DOI: 10.5281/zenodo.18109760

Smawfield, M. L. *Temporal Equivalence Principle: Suppressed Density Scaling in Globular Cluster Pulsars*. Preprint v0.10 (Caracas). Zenodo. DOI: 10.5281/zenodo.18165798

Smawfield, M. L. *The Cepheid Bias: Resolving the Hubble Tension*. Preprint v0.12 (Kingston upon Hull). Zenodo. DOI: 10.5281/zenodo.18209702

Smawfield, M. L. *Temporal Equivalence Principle: A Unified Resolution to the JWST High-Redshift Anomalies*. Preprint v0.8 (Kos). Zenodo. DOI: 10.5281/zenodo.19000827

Smawfield, M. L. *Temporal Equivalence Principle: Temporal Shear Recovery in Gaia DR3 Wide Binaries*. Preprint v0.7 (Kilifi). Zenodo. DOI: 10.5281/zenodo.19102061

Data Availability & Reproducibility

This work follows open-science practices. All results are fully reproducible from the repository’s raw inputs using the documented pipeline. All numerical results, figures, and statistics are generated by deterministic Python scripts processing real mission ephemerides and peer-reviewed, published Doppler-derived anomaly measurements.

Cross-corpus theory registry (TEP-EFA ↔ manuscript series)

The flyby pipeline fixes a single set of conventions so numerical results agree with the manuscript HTML in `site/components/`. The following table is the authoritative map between this repository and the numbered theory papers under `manuscripts/` (Paper 0 = Jakarta, Paper 6 = UCD, Paper 10 = Caracas COS, Paper 12 = JWST Kos, etc.).

Canonical EFA quantities aligned to the manuscript corpus		
Quantity	EFA implementation	Manuscript anchor(s)
Matter metric / conformal factor	$\bar{A}(\varphi) = \exp(\beta_A \varphi / M_{Pl})$ (reduced Planck mass in GeV)	Paper 0 (Jakarta) axioms; Papers 4–5, 8–9, 12 (notation)
Artefact kernel (flyby sector)	Clock-sector excursion: $\Delta v_{TEP} = \beta \times g(\text{geometry})$ with canonical normalization $\Delta\varphi = \psi_{uns} = M_\oplus / (4\pi M_{Pl}^2 R_\oplus) = 1.39 \times 10^{-9}$ (dimensionless η); the earlier force mapping $F_\varphi = \beta_{A,eff} c^2 \nabla\varphi / M_{Pl}$ was retired at v0.3 under the amplitude–shear split (S_Σ screened; the effect lives in the clock sector)	Paper 0 v0.16; Paper 4 (Phantom Mass); methodology §3.2

Quantity	EFA implementation	Manuscript anchor(s)
Density minimum $\varphi_{\min}(\rho)$ — candidate realization	$\varphi = \Lambda [n \Lambda^{n+4} M_{\text{Pl}} / (2 \beta \rho_{\text{GeV}^4})]^{1/(n+1)}$, evaluated at reference response amplitude $\beta = 10^{-4}$, in <code>step_007_tep_model.py</code> , <code>step_011_trajectory_integration.py</code> , <code>step_019_3d_field_integration.py</code>	Paper 10 Appendix C (aligned to this form, May 2026); Paper 6 (scaling $\varphi \propto \rho^{-1/(n+1)}$ only). Paper 0 v0.16 states screening/PPN mapping; it does not fix the closed $\varphi(\rho)$ line — this prescription is a candidate microscopic realization, not the defining Temporal-Topology mechanism, and β here is the channel response amplitude, not the bare coupling $\beta_A = -1$.
PPN γ (magnitude checks)	<code>ppn_gamma_deviation</code> : report $ \gamma - 1 \approx 2 \beta_{\Lambda, \text{eff}}^2$ vs Cassini	Paper 0 Sec. 7: $\gamma - 1 = -2 \alpha_0 \alpha_{\text{eff}}$ (DEF); Papers 5, 11, 12 (screened limit narrative)
Screening / UCD radius	$R_{\text{sol}} \approx 4146 \text{ km}$, $S_{\oplus} = (R_{\oplus} - R_{\text{sol}})/R_{\oplus} \approx 0.35$; $\rho_T \approx 20 \text{ g cm}^{-3}$	Paper 6 (UCD); Step 010 / <code>physics.py</code>
Scalar field equation (sign reference)	Pipeline uses explicit <code>field_gradient</code> / TEP screening relaxation outside Earth; trace source uses β_A , M_{Pl} as in Step 007 comments	Paper 12 Appendix A.1.2: $K(\varphi)\Box\varphi - V' = -(\beta/M_{\text{Pl}})T^{(\text{matter})}$ (Einstein-frame convention; overall sign of T follows chosen action)

Residual ambiguities: Individual papers sometimes use illustrative potentials or linearized V_{eff} without the Einstein-frame factor 2; any updated analytic appendix should match the table above before reusing EFA numerical φ_{Earth} , φ_{space} , or $\Delta\varphi$ in secondary calculations.

Repository & Code

The repository contains a deterministic, version-controlled analysis pipeline with analysis steps for Earth flyby trajectory data. All steps are orchestrated by `scripts/run_all.py` with comprehensive logging.

Repository Structure

```
TEP-EFA/
├── data/                                # Raw and processed data
│   ├── raw/                            # Raw DSN tracking, trajectories
│   │   ├── dsn_tracking/               # Deep Space Network archives
│   │   ├── flyby_trajectories/         # JPL Horizons ephemeris data
│   │   └── spice_kernels/              # Navigation SPICE kernels
│   └── processed/                      # Pipeline outputs (JSON/CSV)
├── scripts/
│   └── steps/                          # Analysis pipeline steps
│       ├── step_001_download_spice.py
│       ├── step_002_spice_to_json.py
│       ├── step_003_archival_data_mining.py
│       ├── step_004_jpl_horizons_fetch.py
│       ├── step_005_dsn_data_ingestion.py
│       ├── step_006_dsn_framework.py
│       ├── step_007_tep_model.py
│       ├── step_008_fitting.py
│       ├── step_009_variance_analysis.py
│       ├── step_010_tep_first_principles.py
│       ├── step_011_trajectory_integration.py
│       ├── step_012_od_filter_simulation.py
│       ├── step_013_cross_validation.py
│       ├── step_014_sensitivity_analysis.py
│       ├── step_015_hierarchical_bayesian.py
│       ├── step_016_gnss_validation.py
│       ├── step_017_plasma_modulation.py
│       ├── step_018_space_weather.py
│       ├── step_019_3d_field_integration.py
│       ├── step_020_plasma_environment_reconstruction.py
│       ├── step_021_mission_specific_od_absorption.py
│       ├── step_022_atmospheric_drag_simulation.py
│       └── step_023_thermal_recoil_modeling.py
```

```

├── step_024_systematic_error_monte_carlo.py
├── step_025_corrected_uncertainty.py
├── step_026_stable_model_comparison.py
├── step_027_claim_consistency_audit.py
├── step_028_dsn_processing.py
├── step_029_read_trk234.py
├── step_030_juno_reanalysis.py
├── step_031_pds_search.py
├── step_032_tep_suppression.py
├── step_033_iri_trajectory_profile.py
├── step_034_covariant_holonomy.py
├── step_035_cross_corpus_export.py
├── step_036_final_report.py
├── step_037_visualizations.py
├── step_038_extract_3d_vectors.py
├── step_039_flyby_prediction_table.py
├── step_040_cosmographic_shear.py
├── step_041_envelope_heuristic_sensitivity.py
├── step_042_time_resolved_cosmography.py
├── step_043_clock_sector_od.py
├── step_044_lense_thirring.py
├── step_anderson_audit.py
├── utils/                # Utility functions
├── build_markdown.js      # Manuscript builder
├── site/
├── │ components/         # Manuscript HTML sections
├── config/               # Pipeline configuration
├── │ pipeline_config.json
├── logs/                 # Per-step execution logs
├── requirements.txt      # Python dependencies
├── README.md             # Documentation
├── LICENSE                # CC-BY-4.0

```

Data Provenance

Data Source	Provider	Access Method	Size	Location
JPL Horizons Ephemeris	NASA/JPL	Astroquery API	~2 MB	<code>data/raw/flyby_trajectories/</code>
DSN Doppler Archives	NASA DSN	Literature values	~500 KB	Anderson et al. (2008)
Flyby Anomaly Catalog	Peer-reviewed literature	Manual compilation	~50 KB	<code>results/step003_archival_flyby_catalog.json</code>
SPICE Kernels	NASA NAIF	Auto-downloaded	~100 MB	<code>data/raw/spice_kernels/</code>

Pipeline Architecture

The analysis pipeline comprises 42 deterministic steps organized into logical groups. Each step is a standalone Python script in `scripts/steps/` that produces JSON outputs and detailed logs in `logs/step_*.log`.

Complete Step Inventory & Runtime

Group	Step	Script	Description	Runtime
Phase 1: Data Acquisition & Preparation (001-006)				
Data	001	<code>step_001_download_spice.py</code>	SPICE kernel download (NAIF archive)	~30s
Data	002	<code>step_002_spice_to_json.py</code>	SPICE to JSON conversion	~1s
Data	003	<code>step_003_archival_data_mining.py</code>	Archival flyby catalog compilation	~2s

Group	Step	Script	Description	Runtime
Data	004	step_004_jpl_horizons_fetch.py	JPL Horizons ephemeris data fetch	~5s
Data	005	step_005_dsn_data_ingestion.py	DSN tracking data ingestion	~1s
Data	006	step_006_dsn_framework.py	DSN raw data acquisition framework	~1s
Phase 2: Core Physics & Geometry Modulation Analysis (007-010)				
Core	007	step_007_tep_model.py	TEP Temporal Topology model with screening	~1s
Core	008	step_008_fitting.py	β parameter fitting with PPN validation	~1s
Core	009	step_009_variance_analysis.py	Deterministic geometry modulation analysis (formal variance partitioning deferred to $n \geq 10$)	~2s
Core	010	step_010_tep_first_principles.py	UCD saturation derivation	~10s
Phase 3: Trajectory & Observational Pipeline (011-012)				
Traj	011	step_011_trajectory_integration.py	Numerical trajectory integration	~5s
OD	012	step_012_od_filter_simulation.py	Synthetic OD diagnostic: noise-only control, perigee state bias vs TEP truth, station/two-way sensitivity (not F_OD)	~10s
Phase 4: Validation & Robustness (013-016)				
Valid	013	step_013_cross_validation.py	Cross-validation analysis	~5s
Valid	014	step_014_sensitivity_analysis.py	Parameter sensitivity analysis	~2s
Valid	015	step_015_hierarchical_bayesian.py	Hierarchical Bayesian model	~30s
Valid	016	step_016_gnss_validation.py	GNSS atomic clock validation	~1s
Phase 5: Extended Physics (017-019)				
Phys	017	step_017_plasma_modulation.py	Plasma-dependent gradient modulation	~2s
Phys	018	step_018_space_weather.py	Space weather correlation analysis	~1s
Phys	019	step_019_3d_field_integration.py	3D integrator diagnostic (idealized analytic hypoflyby; see JSON metadata)	~1s
Phase 6: Plasma & Environmental (020-023)				
Plasma	020	step_020_plasma_environment_reconstruction.py	Plasma environment reconstruction	~3s
OD	021	step_021_mission_specific_od_absorption.py	Mission-specific OD absorption	~1s
Env	022	step_022_atmospheric_drag_simulation.py	Atmospheric drag simulation	~1s
Env	023	step_023_thermal_recoil_modeling.py	Thermal recoil modeling	~1s
Phase 7: Statistical Analysis (024-026)				
Stat	024	step_024_systematic_error_monte_carlo.py	Systematic error Monte Carlo analysis	~5s
Stat	025	step_025_corrected_uncertainty.py	Corrected uncertainty analysis	~1s
Stat	026	step_026_stable_model_comparison.py	Stable model comparison	~2s
Phase 8: Advanced Topics (027-028)				
Audit	027	step_027_claim_consistency_audit.py	Claim consistency audit	~1s
DSN	028	step_028_dsn_processing.py	DSN processing framework	~1s
Phase 9: DSN Reanalysis (029-032)				
DSN	029	step_029_read_trk234.py	Read TRK-2-34 data format	~1s

Group	Step	Script	Description	Runtime
DSN	030	step_030_juno_reanalysis.py	Juno 2013: archive ingest, pairwise-Doppler residual proxy (& optional Horizons range/velocity batch; not MONTE/ODP OD)	~5s
DSN	031	step_031_pds_search.py	NASA PDS archive search	~2s
DSN	032	step_032_tep_suppression.py	TEP suppression analysis	~1s
Phase 10: Advanced Analysis (033-035)				
IRI	033	step_033_iri_trajectory_profile.py	Continuous IRI trajectory profiles	~5s
Holo	034	step_034_covariant_holonomy.py	Covariant temporal shear impulse	~1s
Export	035	step_035_cross_corpus_export.py	Cross-corpus parameter export	~1s
Phase 11: Reporting & Extended Validation (036-042)				
Report	036	step_036_final_report.py	Final report generation	~1s
Fig	037	step_037_visualizations.py	Publication-quality figure generation	~3s
Vec	038	step_038_extract_3d_vectors.py	Extract 3D state vectors from JPL Horizons	~2s
Pred	039	step_039_flyby_prediction_table.py	Flyby prediction table with uncertainty-aware classification	~1s
Cosmo	040	step_040_cosmographic_shear.py	Cosmographic temporal shear modulation test	~1s
Env	041	step_041_envelope_heuristic_sensitivity.py	Geometry envelope heuristic sensitivity	~1s
Cosmo	042	step_042_time_resolved_cosmography.py	Time-resolved cosmography (after Horizons + optional Juno DSN sidecar)	~2s
Analysis	043	step_043_clock_sector_od.py	Clock-sector forward model: proper-time excursion and reconstructed apparent velocity on clock-carrying links	~1s
Analysis	044	step_044_lense_thirring.py	Lense-Thirring frame-dragging bound check on the flyby trajectories (auxiliary; not in CORE_STEPS)	~2s
Audit	—	step_anderson_audit.py	Primary-source catalogue cross-check against Anderson et al. tabulations (auxiliary; not in CORE_STEPS)	<1s

Total Runtime Summary

Component	Steps	Runtime
Data Acquisition (001-006)	6	~40s
Core Physics & Variance (007-010)	4	~14s
Trajectory & Observational (011-012)	2	~8s
Validation & Robustness (013-016)	4	~38s
Extended Physics (017-019)	3	~4s
Plasma & Environmental (020-023)	4	~6s
Statistical Analysis (024-026)	3	~8s
Advanced Topics (027-028)	2	~2s
DSN Reanalysis (029-032)	4	~9s
Advanced Analysis (033-035)	3	~7s

Reporting & Extended Validation (036-043)	8	~13s
Total	43	~2 min

Reproduction Instructions

Quick Start (Full Reproduction)

```
# 1. Clone repository
git clone https://github.com/matthewsmawfield/TEP-EFA.git
cd TEP-EFA

# 2. Install dependencies
pip install -r requirements.txt

# 3. Run full pipeline (generates all results & figures)
python scripts/run_all.py

# 4. Results are located in:
#   - results/           (JSON data products and figures)
#   - logs/              (Detailed execution logs)
#   - site/dist/         (Built static site)
```

System Requirements

Component	Minimum	Recommended	Tested On
CPU	2 cores	4+ cores	Apple M4 Pro (14-core)
RAM	4 GB	8 GB	24 GB (M4 Pro)
Storage	500 MB	1 GB	NVMe SSD
Runtime	~2 min	~1 min	~40s (M4 Pro)

Key Analysis Outputs

- `results/step003_archival_flyby_catalog.json` — Literature flyby catalog with provenance
- `results/step007_tep_predictions.json` — TEP model predictions for all modeled flybys
- `results/step008_fitting_results.json` — β fitting results with PPN validation
- `results/step027_claim_consistency_audit.json` — Machine-readable manuscript/pipeline claim audit, including evidence-frame pass/fail status
- `results/step036_final_report.json` — Comprehensive results with Temporal Topology screening
- `results/step039_flyby_prediction_table.json` — Per-flyby raw pooled- β classification (post-OD columns withheld until Step 021 supplies F_{OD})
- `results/step032_tep_suppression_analysis.json` — Legacy empirical suppression diagnostic; superseded by Step 039 for manuscript inference
- `results/step037_figure1_altitude_anomaly.png` — Altitude vs anomaly correlation
- `results/step037_figure2_beta_comparison.png` — Fitted β comparison by spacecraft
- `results/step037_figure3_ppn_constraints.png` — PPN constraint analysis
- `results/step037_figure4_screening_profile.png` — Temporal Topology profile

Log Files

Each step produces detailed logs:

- `logs/pipeline.log` — Master pipeline execution log
- `logs/step_*.log` — Individual step logs

Software Dependencies

Package	Version	Purpose
Python	3.10+	Language runtime

NumPy	1.24+	Numerical computing
SciPy	1.10+	Statistical functions
Matplotlib	3.7+	Visualization
Astroquery	0.4.6+	JPL Horizons interface
spiceypy	5.1+	SPIICE kernel handling
PyIRI	(package current)	Ionospheric electron density for Step 033 trajectory profiles
pytest	7+	Smoke tests (<code>pytest</code> from repository root)

All dependencies are specified in `requirements.txt`.

Validation & Testing

The pipeline includes comprehensive validation:

- Bootstrap Resampling: $n=10,000$ iterations for uncertainty quantification
- Leave-One-Out Cross-Validation: Tests robustness against single-flyby exclusion
- Heterogeneity Assessment: Cochran's Q and I^2 statistics for model scatter
- GNSS clock correlation: The GNSS atomic clock correlation analysis provides a cross-consistency check with the GNSS/UCD terrestrial calibration on the transition radius ($R_{\text{sol}} \approx 4146$ km). This shared-input dependency is disclosed explicitly to avoid circular claims of independent validation.
- Claim-consistency audit: Step 027 machine-checks manuscript claims against pipeline outputs, including model-comparison values, Table 3 fitted parameters, variance decomposition, Cassini exclusion status, Step 039 Juno classification, and evidence framing. The audit fails if the manuscript reverses the TEP-vs-Null comparison, omits the random-effects uncertainty, or promotes Juno from a deterministic warning case to an uncertainty-aware falsification.
- Uncertainty discipline: Formal inverse-variance summaries, random-effects scatter, the fractional-floor model comparison, and published-uncertainty plus constant-floor stress tests are reported as separate layers rather than blended into a single headline number.
- Automated smoke tests: `pytest` (see `tests/` and `pytest.ini`) checks repository layout and step logger conventions.

Reproducibility Checklist

To verify successful reproduction:

1. All configured pipeline steps complete with "SUCCESS" status
2. Primary JSON products are present in `results/`
3. Figure files are present in `results/` (PNG)
4. Key result: gated inverse-variance β_{fit} span 4.43×10^{-3} to 1.44×10^{-2} across the tightly constrained members (NEAR, Galileo 1990, Rosetta 2005); Galileo 1992 fits at $\beta \approx 6.7 \times 10^{-2}$ and Cassini and MESSENGER are formally divergent through near-zero reference predictions (weighted mean $\approx 4.74 \times 10^{-3}$ from `results/step008_fitting_results.json`)
5. Key result: $S_{\oplus} \approx 0.35$ terrestrial clock-amplitude shell fraction ($R_{\text{sol}} \approx 4146$ km)
6. Key result: $|\gamma_{\text{PPN}} - 1| = 4.32 \times 10^{-8}$ on the two-branch solar profile at $1.6 R_{\odot}$ ($y_{\text{tb}} = 2.16 \times 10^{-8}$; $S_{\Sigma}^{(\odot)} \sim 2 \times 10^{-6}$ conservative solar source-charge convention, safely below the Cassini bound 2.3×10^{-5})
7. Key result: $I^2 \approx 99.6\%$ extreme heterogeneity (supports β scatter hypotheses)
8. Key result: Altitude-anomaly correlation $\rho = -0.65$ ($p = 0.06$, $n = 9$); the TEP envelope predicts near-zero for high-altitude null trajectories and large for low-altitude high-asymmetry encounters, which is the relevant test
9. Key result: the published null or bound cases are consistent with fixed-amplitude Step 039 predictions (Step 008 pooled β_{fit}) under the Step 007 geometry envelope; one deterministic fixed-amplitude raw-tension case remains (Juno, predicted $+0.12$ mm/s) pending mission-specific OD survival factors, and Galileo 1992 carries a raw amplitude surplus (observed -4.60 vs predicted ≈ -0.3 mm/s, correct sign); Rosetta 2009 is a published null/bound case with insufficient explicit geometry for the Step 039 table; 3 flybys (Stardust, OSIRIS-REx, BepiColombo) have no public anomaly report

Data Availability Statement

Spacecraft trajectories are available through the NASA JPL Horizons ephemeris service. Literature anomaly values are from Anderson et al. (2008) and companion publications. Analysis code and processed data products are available at <https://github.com/matthewsmawfield/TEP-EFA> with archived DOI at 10.5281/zenodo.19454862.

Raw DSN tracking products may be obtained from the NASA Deep Space Network and the Planetary Data System following institutional access procedures; per-mission pointers appear under `data/raw/dsn_tracking/<mission>/DOWNLOAD_INSTRUCTIONS.txt`. The present manuscript release does not bundle perigee-matched Level-1 TRK archives: Steps 005–006 and 028–031 implement ingest and audit only, and Step 030 remains inconclusive until such products are added. Headline flyby inference uses literature Δv and JPL Horizons trajectories (Steps 007–026, 039).