



Temporal Equivalence Principle: Temporal Shear Recovery in Gaia DR3 Wide Binaries

Matthew Lukin Smawfield

Version: v0.7 (Kilifi)

First published: 19 March 2026 · Last updated: 17 October 2026

DOI: 10.5281/zenodo.19102061

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

Abstract

The Gaia DR3 catalog of over one million wide binaries opens a precise window onto gravity in the weak-field regime ($a \lesssim 10^{-10} \text{ m s}^{-2}$), yet whether the observed velocity excess reflects modified gravity or unresolved systematics remains contested.

In the Temporal Equivalence Principle (TEP), a conformal scalar field modulates matter proper time as

$$d\tau/dt \approx A(\phi), \quad A(\phi) = \exp(\beta_A \phi / M_{\text{Pl}}). \quad (1)$$

The Cepheid-calibrated response scale is denoted κ_{Cep} , while the wide-binary transition is parameterized independently by the velocity-profile saturation amplitude α_{sat} , not by a bare scalar coupling. This paper tests whether the Gaia wide-binary anomaly is better described as smooth Temporal Shear recovery in weak-field environments.

From 341,315 high-purity systems, the analysis identifies a screening transition at $R_s = 2,646 \pm 182 \text{ AU}$ (statistical; $\pm 609 \text{ AU}$ total), strongly preferred over both a flat Newtonian profile ($\Delta\chi^2 = 14,845$) and a constant boost ($\Delta\chi^2 = 3,583$). The transition is strongly preferred under diagonal, AR(1) and Gaussian-process covariance treatments; even the most conservative covariance-aware comparison gives $\Delta\chi^2 \geq 1,073$ against a flat Newtonian profile and $\Delta\chi^2 \geq 284$ against a constant boost. At large separation the profile saturates at $\alpha_{\text{sat}} = 0.366 \pm 0.012$, roughly 35–40% above the Keplerian baseline. Broader smooth-transition fits preserve the same few-thousand-AU onset.

The signal also shows the environmental ordering required by TEP. With a non-circular metallicity guardrail that uses a conservative external β_{MLR} prior unless independent spectroscopic metallicities are cached, the lower-density high- $|Z|$ population transitions at smaller radius than the higher-density midplane. The headline comparison is the joint profile-likelihood fit, which shares one saturation amplitude across environments—the physically prescribed constraint, since screening modulates the recovery scale while the ambient-floor-limited plateau is a shared response at fixed environment—and returns $R_s = 3,447 \pm 366$ versus $4,912 \pm 792 \text{ AU}$ with $\Delta\chi^2 = 60.3$ ($p = 8.1 \times 10^{-15}$), the ordering preserved in 100% of resamples. The model-free permutation test confirms the ordering at $p \leq 10^{-4}$ (0 exceedances in 10,000 realizations;

fixed-amplitude protocol: $R_s = 4,662 \pm 196$ versus $7,131 \pm 1,341$ AU), as does the solar-track control ($R_s = 4,145 \pm 276$ versus $6,856 \pm 920$ AU; $p \leq 10^{-3}$), whose zero-metallicity-correction selection additionally breaks the $|Z|$ -metallicity correlation by construction. The transition scale furthermore rises with system mass ($R_s = 2,778, 2,966, 4,848$ AU across mass terciles; measured exponent $n = 0.64$ with 86% of resamples positive), consistent with the derived radius law $R_s \propto M^{1/2}$. A five-bin $|Z|$ stratification confirms the density-dependent scaling: the environmental ordering is quantified by an empirical density-scaling exponent, $R_s \propto \rho_{\text{amb}}^p$ with $p \approx 0.4 \pm 0.1$ under free fits of the same stratified radii—consistent with the square-root benchmark $p = 1/2$ and the $\rho^{1/3}$ ambient channel alike (TEP does not commit to a specific microscopic mechanism; the inverse-power chameleon form that motivated the parameterization is excluded for the TEP coupling $\beta_A = -1$, Smawfield 2025a, F1). The ordering itself is selection-independent: restricted to the window below 9 kAU where the purity cut removes fewer than 2% of pairs, the per-separation median- $\tilde{v}$ ordering is monotone in $|Z|$ in every bin, reaches $+0.043 \pm 0.006$ (6.9σ , 100% bootstrap-positive) at $s > 3$ kAU, and is unchanged with no purity cut applied at all. Under bound-probability weighting the profiles rise monotonically to the widest measured separations, the absolute high-versus low- $|Z|$ contrast grows to $+0.094 \pm 0.018$ (5.3σ) at $s = 15\text{--}30$ kAU, and a mixture deconvolution leaves the bound-component median elevated at $s = 20\text{--}50$ kAU under interloper fractions up to three times the prior estimate. Scrambling tests and phase-mixed Newtonian orbital forward models fail to reproduce the observed screening preference. The wide-binary anomaly is therefore not a generic low-acceleration excess but a structured, environmentally modulated screening transition—one whose morphology, onset scale, and environmental ordering are quantitatively consistent with the conformal scalar field of TEP and are not reproduced by the Newtonian orbital-projection or MOND/EFE parameterizations tested here.

Keywords: Temporal Equivalence Principle – wide binaries – Gaia DR3 – weak-field gravity – Temporal Shear recovery – environmental transition morphology – Temporal Topology – Temporal Shear – modified gravity – MOND

1. Introduction

Gaia Data Release 3 (DR3) provided astrometry for more than a million wide binary stars, opening a direct laboratory for gravity in the extreme weak-field regime ($a \lesssim 10^{-10}$ m/s²). Yet the physical interpretation of the resulting signal remains unsettled.

Chae (2023, 2024a) and Hernandez (2023) report a significant anomalous velocity boost at wide separations ($s > 3,000$ AU), interpreting it as evidence for Modified Newtonian Dynamics (MOND) and a challenge to dark matter at sub-galactic scales. Hernandez et al. (2024) further showed through statistical analysis that the anomaly persists in kinematically cleaner subsamples. Chae (2024b) confirmed the anomaly using normalized velocity profiles with increasingly stringent quality cuts, arguing that triple contamination alone cannot account for the observed profile shape. By contrast, Banik et al. (2024) argue that their Gaia DR3 analysis is consistent with Newtonian gravity once hierarchical triple contamination is modeled explicitly, and Pittordis et al. (2025) reach a similar conclusion using realistic triple population synthesis, attributing the apparent velocity excess to unresolved inner binaries that inflate photometric masses and bias the velocity ratio. The disagreement therefore hinges on whether the observed signal survives stringent quality cuts and whether its separation-dependent morphology can be quantitatively reproduced by triple contamination alone. Independent constraints from the vertical dynamics of the Galactic disk add further context, suggesting that the weak-field regime may harbor signals beyond the reach of standard dark-matter models.

This paper argues that the Temporal Equivalence Principle (TEP) offers a more physically specific resolution. By predicting smooth Temporal Shear recovery in weak-field environments, TEP naturally explains a partial, environmentally modulated velocity boost—one that neither pure MOND nor a purely artifactual GR interpretation accommodates as readily.

The extension from the Temporal Topology saturation scale framework (Smawfield 2025g) to wide binaries follows directly. In the conformal sector of TEP (written in isotropic shorthand $\tilde{g}_{\mu\nu} = A^2(\phi) g_{\mu\nu}$ with $A(\phi) = \exp(\beta_A \phi / M_{\text{Pl}})$, and resolved in the operative preferred-foliation realization of Smawfield 2025a, §2.2 into the local lapse $A_{\text{loc}}(u_{\text{grav}})$ and spatial ruler factor $A_{\text{sp}}(u_{\text{drift}})$), matter proper time in the non-relativistic weak-field limit at a fixed cosmic epoch is modulated as $d\tau/dt \approx A(\phi)$. In dense environments the source-charge sector continuously suppresses the locally observable Temporal Shear $\nabla \ln A$; in diffuse environments that suppression weakens and a residual temporal gradient emerges. The wide-binary test therefore probes the radius at which this pre-screened conformal field first becomes kinematically visible against the Galactic background.

The velocity-profile saturation amplitude α_{sat} and the transition scale R_s function as the kinematic realizations of the abstract screening operator $\mathcal{S}_{\Sigma}(\mathcal{E})$ for virialized orbital phase space. Rather than a flat scalar coupling, these parameters trace the continuous un-screening of the Temporal Shear in the ultra-diffuse, weak-field environments of the galactic halo.

The claim-discipline framework for the TEP corpus, including the scope limitations of canonical precision tests, is established in TEP-EXP (Paper 9).

2. Theoretical Framework: Temporal Shear Recovery

Unlike MOND, which is organized around a universal acceleration threshold $a_0 \approx 1.2 \times 10^{-10}$ m/s² (Milgrom 1983), the Temporal Equivalence Principle (TEP) is organized around Temporal Shear recovery in weak-field environments of a conformal scalar field. In the TEP framework (Smawfield 2025a), the isotropic conformal-disformal expression $\tilde{g}_{\mu\nu} = A^2(\phi) g_{\mu\nu} + B(\phi) \nabla_{\mu} \phi \nabla_{\nu} \phi$

provides the historical benchmark shorthand, while the operative realization adopts the preferred-foliation decomposition $\tilde{g}_{\mu\nu} = -A_{\text{loc}}^2(u_{\text{grav}})U_\mu U_\nu + A_{\text{sp}}^2(u_{\text{drift}})\Pi_{\mu\nu} + B(\phi)\Pi_\mu^\alpha \Pi_\nu^\beta \nabla_\alpha \phi \nabla_\beta \phi$ along the cosmic temporal flow $U_\mu \equiv -\nabla_\mu \bar{\phi}/\sqrt{2\bar{X}}$ (Paper 0, §2.2), where $A(\phi) = \exp(\beta_A \phi/M_{\text{Pl}})$ governs local proper-time rates and $B(\phi)$ encodes small disformal deformations. Wide binaries probe non-relativistic orbital separations and velocities at a fixed cosmic epoch ($A_{\text{sp}}(t_0) = 1$), where the disformal sector is negligible and both metric realizations yield the identical slow-matter lapse modulation $d\tau/dt \approx A(\phi)$, with $\alpha(\phi) \equiv d \ln A/d\phi$ the conformal coupling strength. The underlying screening transition is associated with a Temporal Topology saturation scale $\rho_T \approx 20 \text{ g/cm}^3$ (Smawfield 2025g), which is not a binary on/off threshold but a non-linear saturation scale of the conformal-factor sector. In the language of scalar-tensor theory, the gravitational metric $g_{\mu\nu}$ corresponds to the Einstein frame and the matter metric $\tilde{g}_{\mu\nu}$ to the Jordan frame; the kinematic observables measured here—orbital velocities and projected separations—are defined in the Jordan frame, where matter propagates and clocks tick. The conformal factor $A(\phi)$ therefore enters the velocity ratio $\tilde{v}$ directly through its effect on matter-frame proper time, not merely as a post-Newtonian correction to the Einstein-frame potential. This frame distinction is what ties the kinematic enhancement to the clock-rate signature in Section 2.3.

In the weak-field limit, writing $\Theta = \ln A$, the matter-frame potential is $\Phi_{\text{eff}} \simeq \Phi_N + c^2 \Theta$. Matter-frame motion therefore responds to

$$\mathbf{a}_{\text{eff}} = -\nabla \Phi_N - c^2 \boldsymbol{\Sigma}^{\text{obs}}, \quad \boldsymbol{\Sigma}^{\text{obs}} = \mathcal{S}_\Sigma(\mathcal{E}) \nabla \ln A \quad (2)$$

where $\mathcal{S}_\Sigma(\mathcal{E})$ is the environmental screening factor that suppresses the locally observable Temporal Shear. For an approximately radial binary response,

$$v_c^2(s) \simeq \frac{GM}{s} + c^2 s \hat{\mathbf{s}} \cdot \boldsymbol{\Sigma}^{\text{obs}} \quad (3)$$

The observed velocity enhancement is therefore the kinematic projection of recovered Temporal Shear, while α_{sat} parameterizes the wide-binary channel response rather than the microscopic coupling β_A .

2.1 Consistency with Solar-System PPN Constraints

The Cassini Shapiro-delay bound $|\gamma_{\text{PPN}} - 1| < 2.3 \times 10^{-5}$ (Bertotti, Iess & Tortora 2003) constrains the source-charge sector $\mathcal{S}_\Sigma$ in the solar-vicinity environment, not the wide-binary environment probed here. Two features of the Cassini measurement make the bound weaker than a naive reading suggests for TEP. First, the measurement is a round-trip radio ranging experiment performed during solar conjunction: the signal path samples the Sun's deep potential well, which in TEP corresponds to the screened regime where $\mathcal{S}_\Sigma$ is suppressed. 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), $\gamma_{\text{PPN}} - 1 = -2\beta_A^2 \mathcal{S}_\Sigma^{(\odot)} / (1 + 2\beta_A^2 \mathcal{S}_\Sigma^{(\odot)})$, so the bound $\mathcal{S}_\Sigma^{(\odot)} \lesssim 1.15 \times 10^{-5}$ (5.8×10^{-6} on the isotropic-conformal benchmark; with the solved two-branch solar profile giving $y_{\text{tb}}(1.6 R_\odot) = 2.16 \times 10^{-8}$ and $|\gamma_{\text{PPN}} - 1| = 4.32 \times 10^{-8} \ll 2.3 \times 10^{-5}$) applies to the solar-vicinity environment along the signal path, not to the dilute interstellar medium where the Temporal Shear recovers and wide-binary dynamics originate. Second, the round-trip geometry is reciprocity-even: the conformal clock variation enters as a gradient integral $\int \nabla(\ln A) \cdot d\mathbf{l}$, and the outbound and inbound contributions have opposite signs because the path is traversed in opposite directions, so the clock-sector contribution cancels exactly in the round-trip observable (Smawfield 2025a, Smawfield 2025i). The wide-binary signal, by contrast, is a kinematic observable in the dilute interstellar medium where $\mathcal{S}_\Sigma$ has recovered; it probes a different environment and a different measurement geometry from the Cassini bound. The ten-order-of-magnitude gap between the Cassini and wide-binary Yukawa mass requirements is therefore environmental, not sectoral: $\mathcal{S}_\Sigma$ is suppressed in the high-density solar vicinity and recovers in the low-density interstellar medium. The density-dependent minimum of the TEP field equation $\nabla^2 \phi = V_{,\phi} + \rho_* A_{,\phi}$ provides the environmental transition; this is a general property of the bidirectional conformal field equation, not specific to any particular microscopic completion.

The reconciliation has a precise spatial form. The corpus's master-action screening closure derives a pairwise response of the form $\mathcal{R}(s) = \mathcal{S}_\Sigma(X_{\text{env}})^2 y(s)$ (Smawfield 2025a, §7): the mutual shear profile $y(s)$ —the exact flux-conserving solution $y[1 + y^2(r_*/r)^4] = 1$ of the stated kinetic completion $P = X - V + X|X|/\Lambda_X^4$, with $\phi' \propto r^{-2/3}$ deep inside the shell—propagated with the ambient inverse-stiffness factor $\mathcal{S}_\Sigma(X_{\text{env}}) = [1 + 2X_{\text{env}}/\Lambda_X^4]^{-1}$ at the emission and response vertices. The ambient X_{env} is the field of the hierarchy level above the pair, excluding the pair's own co-generated mutual field. Two regimes of the same operator follow. For a hierarchical pair—a test member embedded in a dominant member's pre-existing nonlinear shell—the dominant member's field is the embedding ambient, and the identity $2X_{\text{dom}}/\Lambda_X^4 = (1 - y)/y$ makes the vertex factors equal to $y(s)$ itself: the response recovers $\mathcal{R} \sim y^3 \sim s^4$, the quartic deep-interior asymptote. For a comparable-mass wide binary the mutual field is the measured response rather than a fixed embedding—the two members' field gradients oppose and cancel at the pair's bridge region—so the vertex ambient saturates at the Galactic field and the transition profile is the flux-conserving $s^{4/3}$ law, the exponent the data select (Section 4; effective exponent $\simeq 1.3$ through the full forward model). This is a continuous Temporal-Topology

recovery rather than a thin-shell boundary—consistent with the rejection of the thin-shell morphology reported in Section 4. The shell radius is the derived law $R_{\text{shell}}(M) \equiv r_*(M) = \sqrt{GM/g_t}$, with $g_t = cH_0/(2|\beta_A|) \simeq 3.4 \times 10^{-10} \text{ m s}^{-2}$, not a per-channel fitted scale. On this single map every existing Solar-System constraint lies deep inside the solar shell, where the hierarchical response $\mathcal{R} \sim s^4$ applies: $\mathcal{R} \simeq 10^{-23}$ at the Cassini conjunction separation ($1.6 R_\odot$), $\simeq 3 \times 10^{-15}$ at 1 AU, and $\simeq 3 \times 10^{-11}$ at Saturn's orbit (9.5 AU)—the corresponding anomalous acceleration $\delta a \simeq 3.5 \times 10^{-15} \text{ m s}^{-2}$ sitting more than two orders of magnitude inside the Cassini-era ephemeris bound—while the Earth–Moon pair lies at $\simeq 2 \times 10^{-15}$ inside the nested terrestrial and solar shells. The pairwise response at the conjunction separation thus sits seventeen orders of magnitude below the already stringent single-body bound $S_\Sigma^{(\odot)} \lesssim 1.15 \times 10^{-5}$ (5.8×10^{-6} isotropic), whereas a wide-binary companion at $s \sim R_s$ samples the shell boundary itself. The fitted $R_s = 2,646 \text{ AU}$ is therefore a measurement of the solar-type screening-shell radius, not a new screening variable. The corpus's zero-parameter derived value, $R_{\text{shell}} \simeq 4.7 \text{ kAU}$ at the sample mean mass, sits a factor $\simeq 1.9$ above the fitted scale under the isolated-profile reading—but the isolated profile is not the embedded one. Solving the source perturbation on the Galactic nonlinear background ($u_0 = \sqrt{X_{\text{gal}}} \simeq 0.72 g_t$; Paper 0 step 32) shows the ambient stiffness $Z_{ij} = P_{,X}\delta_{ij} + 2P_{,XX} a_i a_j$ suppressing the pair's far-field response to $y_{\text{far}} \simeq 0.56$ and compressing the transition scale to $\simeq 0.6 r_*$. Folded through this paper's own sem-weighted single-scale estimator and mass distribution (step 017), the embedded prediction returns $R_s \simeq 3,300 \text{ AU}$ and a half-excess marker at $\simeq 1,715 \text{ AU}$ —the observed half-excess—a factor $\simeq 1.26$ residual now inside the systematic band rather than a factor- $\simeq 2$ contradiction. The ambient floor displaces the transition inward through the propagator itself, not through a tuned normalization (Smawfield 2025a, §7). At the directly measured ambient ($u_{\text{sun}} = 0.57$, above) with the ambient setting the vertex and the propagator together, the same estimator returns $\alpha = 0.354$ against the observed 0.366—the amplitude is reproduced at the percent level with no calibrated input—at $R_s \simeq 3,750 \text{ AU}$, a factor $\simeq 1.4$; the two ambients bracket the physical value because the ambient field gradient is itself a solved quantity in the theory. The environmental ordering constrains and, under the corpus's own ambient sectors, now checks quantitatively. The observed subsamples show the midplane more suppressed under either estimator — free- α fits give $(R_s, \alpha) = (2,821 \text{ AU}, 0.244)$ at $|Z| < 0.1 \text{ kpc}$ versus $(4,681 \text{ AU}, 0.401)$ at $|Z| > 0.15 \text{ kpc}$, both parameters falling toward the plane along the mechanism's own ambient trajectory; the shared- α joint fit reads the same suppression ordering through the amplitude–scale degeneracy as a larger midplane R_s . Inverting the free- α amplitudes on the fully-coupled sweep requires the effective ambient to be $\sim 36\%$ stronger at the midplane ($u_{\text{eff}} \simeq 0.71$ versus $\simeq 0.52$). A Galactic field-model check (step 017, supply-versus-demand block) tests two ambient estimators. The linear-gradient estimator $u = |a|/g_t$ supplies only $\sim 4\%$, and with the wrong sign for the amplitude ordering — the radial floor $u_r \simeq 0.57$ is common to both heights and the vertical component adds with $|Z|$ — but that estimator is the weak-field exterior reading: inside the nonlinear disk the local acceleration is not the ambient field state (the disk's field *value* sits at its density-set equilibrium, $u_{\text{min}} \propto \rho^{1/3}$ on the quartic branch, while the pair-relevant ambient is the shear variable, whose nonlinear response the k-branch balance resolves). The density-set ambient supplies the contrast quantitatively: the three-component Galactic model gives $\rho_{\text{mid}}/\rho_{\text{halo}} = 2.44$ between the strata's median heights, a 34.6% ambient contrast toward the midplane under $u_{\text{eff}} \propto \rho^{1/3}$ against 35.8% required, with the identified $u_{\text{eff}} \propto \rho^{1/2}$ k-branch law bracketing the same requirement at 56% under the median-density mapping — the two-bin contrast is an estimator-level check that brackets the two exponent channels, while the five-stratum inversion fixes $p \simeq 0.46\text{--}0.5$ (Section 6) and the ambient map is identified at coefficient level as the local k-branch law $u_{\text{amb}} = \sqrt{4\sqrt{2}\pi G\rho L/(Kg_t)}$ with coherence length $L \simeq 0.84 \text{ kpc}$ (Paper 0, step 67). The matched-covariate audit of step 018 finds that mass, separation, distance and metallicity matching absorbs only $\simeq 22\%$ of the contrast (and the mass covariate predicts the opposite sign), so ordinary demographic channels do not supply it either. The ordering is therefore consistent with the density-set ambient the corpus specifies for a nonlinear environment. Which projection enters the pair vertex — the former Gate-A question — is now resolved at response-function level: the embedded pair-scale solve on the constraint-slice background returns a response charge $D^{-1} \propto u_{\text{amb}}^{-2,0}$ with local background-density sensitivity $\lesssim 0.02\%$ across the full disk range, so the vertex samples the ambient boundary variable and the density ordering enters only through the ambient map (step 67, embedded-response block).

The shell reading also fixes the amplitude bookkeeping. With $\beta_A = -1$, a fully unscreened source would contribute $2\beta_A^2 g_N = 2g_N$ of extra centripetal acceleration—a plateau at $\alpha_{\text{sat}} = \sqrt{3} - 1 \simeq 0.73$. The observed plateau corresponds to $+0.87 g_N$, approximately 43% of the unscreened response: a partial recovery bounded by the Galactic ambient floor rather than a return to unsuppressed charge, with an implied per-body factor $S_{\text{body}} \simeq 0.66$ comparable to the corpus's ambient suppression at the solar circle evaluated on the measured field, $\mathcal{S}_\Sigma(X_{\text{sun}}) \simeq 0.75$ for $u_{\text{sun}} = a_\odot/g_t = 0.57$ ($a_\odot = v_c^2/R_0 \simeq 1.94 \times 10^{-10} \text{ m s}^{-2}$ with $v_c = 220 \text{ km s}^{-1}$, $R_0 = 8.1 \text{ kpc}$; the legacy corpus value $\mathcal{S}_\Sigma \simeq 0.66$ at $X_{\text{gal}} = 0.52$ was itself calibrated to this plateau under the no-propagator bookkeeping; step 017, consistent-vertex sweep). The same shell map closes the corpus's remaining constraints on this class of source. At Oort-cloud separations $s \sim 10^4\text{--}10^5 \text{ AU}$ the pairwise factor is $\gtrsim 0.95$: bodies there orbit in the recovered-shear regime, the domain analyzed by the companion outer-Solar-System boundary study (Smawfield 2026, Paper 36). Terrestrial clock networks, at 1 AU, sit deep inside the solar shell (solar contribution $\mathcal{S}_{\text{eff}} \sim 10^{-15}$) and measure the nested terrestrial well's own shear, consistent with the GNSS correlation phenomenology (Smawfield 2025b, 2025c). And the ambient-floor-limited response surviving beyond every stellar shell supplies the effective-gravity contribution invoked by the rotation-curve sector (Smawfield 2025g). The screening shell is therefore not an additional assumption but the spatial semantics of the single operator $\mathcal{S}_\Sigma(\mathcal{E})$ evaluated across separations: one map, bounded inside by Cassini, measured at its boundary by wide binaries, and read off asymptotically by the Galactic field.

2.2 The Screening Radius

R_T here denotes the projected separation at which the Temporal Shear suppression transitions from the screened to active-shear regime, related to but distinct from the GNSS correlation length λ_T (which characterises the radial relaxation scale within a gravitational well) and the geometric saturation radius $R_T(M, \rho_T, \epsilon_{\text{env}})$ derived in Paper 6 (which characterises the system-level boundary condition for a self-gravitating source of mass M). The wide-binary screening radius R_s introduced below in Equation (5) is the operative benchmark scale for this transition.

For a binary of total mass M , the locally active Temporal Shear $\Sigma_\mu \equiv \nabla_\mu \ln A(\phi)$ is suppressed continuously by the environmental and source state rather than gated by an on/off density threshold (Smawfield 2025a, Section 7). As the binary's internal Newtonian potential becomes shallow enough that the source-charge sector no longer fully suppresses the locally observable Temporal Shear, the residual gradient of $\ln A$ generates an additional kinematic enhancement above the Keplerian baseline. The transition is therefore characterized by a continuous recovery of the Temporal Shear, parameterized below by an effective screening radius R_s .

Within the Galaxy, however, the binary is not isolated. The TEP framework itself does not commit to a single microscopic screening mechanism: Vainshtein, Galileon, DBI, and symmetron mechanisms are treated as candidate completions of the conformal-sector screening (Smawfield 2025a, §A4, §7), with the defining ontology being the continuous suppression of the locally observable Temporal Shear $\Sigma_\mu \equiv \nabla_\mu \ln A(\phi)$ by source-charge sector, environmental state, and boundary conditions; the inverse-power chameleon branch is excluded under the TEP coupling convention, since $\beta_A = -1$ with $\phi > 0$ in potential wells admits no admissible minimum for $V \propto \phi^{-n}$ (Smawfield 2025a, constraint F1). The quantitative benchmark for the wide-binary transition scale is therefore an empirical environmental-density parameterization rather than a specific potential. Any admissible completion supplies an environment-dependent screened equilibrium—a strong effective scalar response that pins the ambient clock landscape in dense environments and suppresses the incremental shear, driving the locally observable $\Sigma_\mu \rightarrow 0$ continuously rather than via thin-shell matching. The surrounding halo and disk already drive ϕ toward this screened configuration, so the transition scale relevant to a wide binary is set not by ρ_T directly but by the residual local screening floor that the binary samples. This floor admits a closed-form parameterization that serves as a benchmark; admissible completions would yield qualitatively similar but quantitatively different forms.

$$\rho_{\text{floor}} = \epsilon_{\text{env}} \rho_T \quad (4)$$

$$R_s = \sqrt{\frac{GM}{g_t}} \propto M^{1/2}, \quad \rho_{\text{eff}}(R_s) \equiv \frac{3M}{4\pi R_s^3} = \rho_{\text{floor}} = \epsilon_{\text{env}} \rho_T \quad (5)$$

Here $\epsilon_{\text{env}} < 1$ is the pre-screening factor generated by the Galactic environment. The left-hand member is the canonical derived radius law of the nested operator (the flux-conserving shell radius; Section 2.1); the right-hand member is the equivalent enclosed-density diagnostic coordinate of the same transition—for the canonical law ρ_{eff} scales as $M^{-1/2}$ at fixed environment rather than as a fixed density—and it serves as a calibration benchmark rather than a unique prediction of the underlying theory. Two of the three ingredients— ρ_T and the field equation whose solution determines ϵ_{env} —are already derived in earlier work. The Temporal Topology saturation scale $\rho_T \approx 20 \text{ g/cm}^3$ is anchored by GNSS atomic-clock networks and supported by SPARC rotation-curve slopes and compact-object consistency tests (Smawfield 2025g). The density-dependent effective potential $V_{\text{eff}}(\phi; \rho)$ whose ground state sets the local screening floor follows from the TEP two-metric action together with the scalar self-interaction sector (Smawfield 2025a, §7). The character of the transition is further tested by comparing the phenomenological TEP recovery profile—the historical fitted form whose steepness is now derived from the master-action kinetic completion $P(X, \phi) = -V(\phi) + \frac{2k}{3\Lambda_X^2} X|X|^{1/2} + X|X|/\Lambda_X^4$ ($P_{,X} = k\sqrt{\xi} + 2\xi$, $k \simeq 16.03$; unit-term control $X - V + X|X|/\Lambda_X^4$)—to alternative transition morphologies. As shown in Section 4, both the thin-shell exterior profile itself and a sigmoid step-like proxy are rejected relative to the TEP exponential ($\Delta\chi^2 = +613.0$ and $+131.5$, respectively; Table 4.1). This rejection is naturally explained by the continuous Temporal Topology framework: the field gradient recovers smoothly, not via a step-function onset.

The canonical cross-scale anchor is now the master-action closure itself: the derived acceleration scale $g_t = cH_0/(2|\beta_A|) \simeq 3.4 \times 10^{-10} \text{ m/s}^2$ yields the radius law $R_s = \sqrt{GM/g_t} \simeq 4.7 \text{ kAU}$ at the sample mean mass, a factor $\simeq 1.8$ above the fitted R_s (Section 2.2; Smawfield 2025a, §7)—the physically expected direction, since the ambient Galactic field raises the environmental floor and displaces the transition inward. An older SPARC-extracted benchmark, $g_{\text{TEP}} \approx 5 \times 10^{-10} \text{ m/s}^2$ (Smawfield 2025g), returns 3,831 AU (factor 1.5), and a further external-field divisor $\eta = 2$ tightens that empirical normalization to 2,709 AU (3%)—but that tightening is a tuned factor rather than a derived one, and both empirical anchors are superseded as predictions by the canonical g_t . The residual is now quantified by a same-estimator check (step 017): folding the mass-convolved derived profile through this paper's own sem-weighted exponential estimator returns $\simeq 5.1 \text{ kAU}$ against the fitted $2,646 \pm 609 \text{ AU}$ (ratio $\simeq 1.9$), and the physical half-excess marker agrees ($\simeq 3.4 \text{ kAU}$ versus the observed $\simeq 1.7\text{--}1.9 \text{ kAU}$). Definitional and mass-normalization explanations are quantitatively insufficient — the canonical-fit scale reduction is $\lesssim 2\%$, $\sqrt{M}$ -weighting accounts for $\sim 4\%$, and mass convolution broadens the transition upward. The embedded-propagator check of step 017 then resolves most of the scale offset: solving the source perturbation on the Galactic nonlinear background ($u_0 = \sqrt{0.52}$) rather than on the isolated-vacuum

background compresses the transition, and the same estimator returns $\simeq 3.3$ kAU against the fitted $2,646 \pm 609$ AU (ratio $\simeq 1.26$, within the systematic band), with the physical half-excess marker coinciding with the observed value ($\simeq 1.7$ kAU on both sides). The mean-scale offset is therefore substantially a propagator effect of the ambient stiffness $Z_{ij} = P_{,X}\delta_{ij} + 2P_{,XX}a_i a_j$ — a derived consequence of the master-action kinetic operator acting in its own nonlinear ambient field — not a free environmental coefficient. The remaining offset (factor $\simeq 1.26$) collects the open residuals quantified by the corpus: the amplitude factorization, with the observed $\alpha_{\text{sat}} = 0.366$ sitting between the pure-propagator (0.455) and propagator-times-vertex-squared (0.24) readings, and the Gate-A single-configuration question; the older standalone divisor $\eta_{\text{env}} \simeq 3.7$ on g_t (implied $g_{\text{eff}} \simeq 1.3 \times 10^{-9} \text{ m s}^{-2}$) is its bookkeeping shorthand, not an additional mechanism. The wide-binary R_s is therefore not an unanchored binary-sector scale: its order of magnitude is fixed by the cross-scale architecture, and its residual scale is supplied by the same ambient-stiffness propagator rather than by an environmental coefficient introduced for the purpose.

The absolute value of ϵ_{env} at a given Galactic position remains calibration-dependent at this stage. A semantic point should be stated plainly: the quantity $\rho_{\text{eff}} = 3M/(4\pi R_s^3) \approx 9.2 \times 10^{-18} \text{ g/cm}^3$ evaluated at $M \approx 1.2 M_\odot$ is the binary's own enclosed mean density at the transition separation—some six orders of magnitude above the ambient ISM density—not an environmental density floor. Under the canonical operator the transition criterion is the ratio of the pair's internal field gradient to the ambient environmental shear floor; ρ_{eff} is an equivalent diagnostic coordinate of that criterion (for $R_s = \sqrt{GM/g_t}$ it scales as $M^{-1/2}$, not as a fixed density), while $\epsilon_{\text{env}} \approx 4.6 \times 10^{-19}$ is the numerical value of the benchmark parameterization $\rho_{\text{floor}} = \epsilon_{\text{env}} \rho_T$ rather than a physical density. The height dependence of ϵ_{env} is then semi-predictive within that parameterization: the screening radius is taken to scale as a power law in ambient baryonic density, $R_s \propto \rho_{\text{amb}}^p$, so the ratio of screening radii at two Galactic heights depends only on the ambient density ratio and the scaling exponent p . As shown in Section 5.2, the benchmark value $p = 1/2$ combined with a standard three-component Galactic density model reproduces the observed environmental R_s ratio within $\sim 1\sigma$. The TEP-native statement is qualitative—more strongly screening environmental states more strongly suppress the locally observable Temporal Shear—while the quantitative exponent p is an empirical property of the environmental response, not of a specified microscopic completion. (The inverse-power chameleon form that historically motivates writing $p = 1/(n+1)$ is excluded under the TEP coupling convention; Smawfield 2025a, F1.) Promoting the absolute ϵ_{env} to a first-principles prediction would require solving $\nabla^2 \phi = dV_{\text{eff}}/d\phi$ across the full three-dimensional baryonic density field, feasible with existing N-body–scalar-field codes (e.g., Llinares, Mota & Winther 2014). The large gap between ρ_T and ρ_{eff} does not imply two unrelated saturation scales; it reflects the many orders of magnitude of pre-screening already absorbed by the Galactic halo and disk before the binary's own potential is tested. The radius R_s is thus the screening-shell boundary introduced in Section 2.1: the surface on which the pair's enclosed mean density falls to the effective environmental floor, equivalently parameterized by the corpus's derived transition law $R_{\text{shell}}(M) = \sqrt{GM/g_t}$. The density and acceleration parameterizations of the shell boundary identify the same physical transition, and their distinct mass scalings are discriminated directly by the mass-ordered transition itself (Section 4), whose measured exponent is consistent with the derived $\sqrt{GM/g_t}$ law.

2.3 The Velocity Enhancement Profile

Within this framework, the kinematic profile is expected to satisfy three qualitative constraints: (i) $\tilde{v} \rightarrow 1$ in the fully screened limit $s \ll R_s$; (ii) a monotonic transition near $s \sim R_s$; and (iii) saturation at a bounded enhancement $\tilde{v} \rightarrow 1 + \alpha_{\text{sat}}$ for $s \gg R_s$. These follow directly from conformal-sector screening and do not depend on the detailed nonlinear field dynamics.

These qualitative features are not merely asserted; they follow from the continuous Temporal Topology of the conformal scalar field. In the TEP framework (Smawfield 2025a, §7), screening operates not via the thin-shell approximation but through the continuous spatial profile of $\ln A(\phi)$ (the Temporal Topology), governed by the non-linear superposition of conformal-factor gradients (the Temporal Shear, $\Sigma_\mu \equiv \nabla_\mu \ln A = \alpha(\phi) \nabla_\mu \phi$). In dense environments the source-charge sector and environmental state drive the locally observable shear toward zero, flattening the Temporal Topology. As that suppression weakens with the diluting ambient environment—rather than crossing a binary on/off density threshold—the Temporal Shear recovers continuously, producing a smooth transition from the screened to the active-shear regime. This Temporal-Topology framing is independent of the specific microscopic screening completion; the empirical density-scaling benchmark used in §2.2 is consistent with the same continuous-suppression ontology shared by the admissible completions (Vainshtein, Galileon, DBI, symmetron).

The classical thin-shell formalism (Khoury & Weltman 2004; Burrage & Sakstein 2018) provides useful context. For a screened mass M with screening radius R_s embedded in a background of effective scalar mass $m_{\text{bg}} = \sqrt{V''_{\text{eff}}(\phi_{\text{bg}})}$, the exterior scalar field profile obtained under thin-shell matching conditions is:

$$\phi(r) = \phi_{\text{bg}} - C \frac{R_s}{r} e^{-m_{\text{bg}}(r-R_s)}, \quad r > R_s \quad (6)$$

where C is determined by the thin-shell matching and the R_s/r prefactor arises from the thin-shell boundary. The corresponding unscreening fraction $f(s) = 1 - (R_s/s) e^{-m_{\text{bg}}(s-R_s)}$ produces onset–rise–saturation morphology with a step-like character at R_s .

However, the TEP Temporal Topology framework replaces the thin-shell approximation with a continuous field profile. Without a step-function matching condition, the R_s/s geometric prefactor—which encodes the thin-shell boundary—does not arise. Continuous Temporal Topology therefore requires a smooth, one-scale recovery from the screened state toward a finite asymptotic Temporal Shear. The canonical wide-binary realization of this relaxation is the pure exponential:

$$\tilde{v}(s) = 1 + \alpha_{\text{sat}}(1 - e^{-s/R_s}) \quad (7)$$

where R_s is the characteristic scale of the Temporal Topology transition and α_{sat} is the saturation amplitude set by the asymptotic Temporal Shear. This is adopted as the canonical fitting function throughout the analysis. Unlike the thin-shell formula, the pure exponential is the canonical one-scale realization of continuous Temporal Topology in the wide-binary sector, containing exactly the two physical quantities required by the TEP screening geometry: the recovery scale R_s and the asymptotic response α_{sat} . Equation (7) is the minimal smooth realization of the TEP onset–recovery–saturation morphology; the transition scale and environmental ordering, rather than the exact exponential shape alone, carry the physical inference. Mass and environmental variation broaden this canonical profile rather than changing its underlying onset–recovery–saturation structure, as demonstrated explicitly by the mass-convolved TEP model in Section 4.1. A literal first-principles prediction for the precise profile shape would require solving the full coupled system $\nabla^2\phi = dV_{\text{eff}}/d\phi$ in the realistic potential of each binary within its Galactic environment. What the data can test at this stage is whether the morphological class of profile predicted by continuous Temporal Topology—finite onset, smooth exponential rise, bounded saturation—is preferred over the alternatives (scale-free MOND, constant boost, flat Newtonian, or the sharper step-like transition of the thin-shell sigmoid). The sensitivity of the inferred transition scale to the choice of functional form is assessed explicitly in Section 4 by fitting sigmoid and double-exponential alternatives, and the resulting spread is absorbed into the systematic uncertainty budget. As shown there, all three transition models agree on a $\sim 2,000$ – $3,200$ AU onset scale, confirming that the inferred R_s is a robust feature of the data rather than an artifact of any particular parametric choice. Notably, the data reject both the stated thin-shell exterior profile ($\Delta\chi^2 = +613.0$ versus the TEP exponential) and the sigmoid step-like proxy ($\Delta\chi^2 = +131.5$). This rejection is naturally explained by the continuous Temporal Topology framework: the field gradient recovers smoothly, not via a step-function onset.

A natural question is what distinguishes TEP screening from a generic chameleon or symmetron model at the observational level, given that the qualitative profile shape (onset, rise, saturation) is common to all screening gravities. The answer lies in three features that go beyond morphology. First, TEP makes a quantitative cross-scale prediction: the same universal conformal sector and Temporal Topology saturation architecture that govern atomic and compact-object phenomenology (Smawfield 2025g) also set the Galactic screening floor. The wide-binary transition scale is therefore not theoretically isolated within the binary sector, but is anchored to independently measured TEP scales. Second, the environmental test (Section 5) probes the environmental dependence of the locally observable Temporal Shear generated by the conformal sector; in the density-scaling benchmark this dependence projects primarily onto ambient baryonic density. Third, continuous Temporal Topology predicts a smoother recovery than the thin-shell approximation of standard chameleon models; in the canonical wide-binary realization adopted here, that recovery is represented by the pure exponential $\tilde{v}(s) = 1 + \alpha_{\text{sat}}(1 - e^{-s/R_s})$ rather than the thin-shell form $f(s) = 1 - (R_s/s)e^{-m_{\text{bg}}(s-R_s)}$ with its geometric R_s/s prefactor. Notably, the data already favor this prediction: the stated thin-shell profile $f(s) = 1 - (R_s/s)e^{-m_{\text{bg}}(s-R_s)}$ is rejected at $\Delta\chi^2 = +613.0$ and the sigmoid step-like proxy at $\Delta\chi^2 = +131.5$, both relative to the TEP exponential (Table 4.1). Future observations could sharpen these distinctions: Gaia DR4 radial velocities would enable three-dimensional deprojection, testing whether the proper-time signature of TEP—in which the same Temporal Shear that produces the kinematic enhancement also rescales matter-frame clock rates as $d\tau/dt = A(\phi)$ —produces a different velocity-anisotropy pattern than a generic screened scalar model that does not make the TEP-specific identification of the conformal gradient with dynamical proper-time geometry. Measurements of the transition radius as a function of Galactocentric radius R_{gc} would further discriminate among candidate completions, since the density-scaling benchmark adopted here predicts a screening floor set by the local baryonic density rather than the total gravitational potential (Burrage & Sakstein 2018), whereas symmetron- or Vainshtein-style completions would imply different gradient-profile dependences—a pattern that maps directly onto the underlying Temporal Shear morphology rather than onto a uniform acceleration threshold.

3. Data and Methodology

3.1 Sample Selection

This study constructs a wide-binary sample from the Gaia DR3 astrometric catalog (Gaia Collaboration et al. 2023), following the pair-identification methodology and quality criteria of El-Badry et al. (2021). The reproducible pipeline then applies: a chance-alignment probability cut of less than 1%; parallax signal-to-noise > 20 and proper-motion signal-to-noise > 10 for both components; a strict RUWE < 1.2 cut; and a projected-separation window of $50 < s < 50,000$ AU. Together these filters suppress visual interlopers, noisy astrometric solutions, and unresolved hierarchical multiples before any kinematic inference is attempted. The resulting high-purity sample contains 341,315 systems.

The RUWE cut deserves particular comment, since hierarchical triples are the principal contaminant identified by Pittordis et al. (2025). Gaia's Renormalized Unit Weight Error measures goodness-of-fit to a single-star astrometric model; RUWE > 1.4 is the

standard flag for an unresolved companion. The pipeline adopts a stricter threshold of 1.2, removing 193,545 systems (36% of the post-parallax sample), of which 130,733 have $\text{RUWE} > 1.4$ for at least one component. If the pre-filter triple fraction is 15–20% (Pittordis et al. 2025), the number removed exceeds the expected triple count by 1.8–2.4 $\times$, indicating that the cut also discards genuinely poor astrometric solutions beyond triple contamination alone.

The residual triple fraction is bounded by RUWE detection incompleteness for long-period inner binaries ($P \gtrsim 3\text{--}5$ yr), which Gaia's ~ 34 -month DR3 baseline cannot resolve. For plausible detection efficiencies of 50–70%, the residual contamination is $\sim 5\text{--}10\%$, and the surviving triples are preferentially the least kinematically disruptive—long inner periods, small velocity perturbations. Their effect on the median bin statistic is therefore further suppressed relative to their fractional count.

3.2 Metallicity-Dependent Mass Estimation

A central systematic—identified in the recent literature and confirmed by the present audit—is the effect of metallicity on stellar mass estimation. Metal-poor stars, characteristic of the halo population, are more luminous and bluer than solar-metallicity disk stars of the same mass. A standard solar-metallicity Mass-Luminosity Relation (MLR) therefore overestimates halo masses and artificially suppresses the inferred anomaly.

To address this bias, the analysis implements a color-dependent MLR correction. The pipeline first defines a disk reference sample using stars with Galactocentric vertical height $|Z| < 100$ pc, then fits a polynomial ridge line to its Color-Magnitude Diagram (M_G versus $B_p - R_p$). For every star in the full catalog, the pipeline computes the color offset $\Delta C = (B_p - R_p)_{\text{obs}} - (B_p - R_p)_{\text{ref}}$.

The stellar masses are then calculated as:

$$M = M_{\text{solar}}(M_G) \times (1 + \beta_{\text{MLR}} \Delta C) \quad (8)$$

Here $M_{\text{solar}}(M_G)$ is the baseline solar-metallicity mass from the empirical main-sequence relations of Pecaut & Mamajek (2013, updated 2022), and $\beta_{\text{MLR}} \approx 1.5$ is a conservative color-mass coefficient. To avoid circularity, the pipeline calibrates β_{MLR} only from independent spectroscopic metallicities when a LAMOST or APOGEE cache is available; photometric $[\text{Fe}/\text{H}]$ proxies derived from ΔC are retained for diagnostics but are not regressed back onto ΔC . In the absence of an independent spectroscopic cache, the pipeline uses the external $\beta_{\text{MLR}} = 1.5 \pm 0.5$ prior and then stress-tests the environmental ordering with $\beta_{\text{MLR}} = 0, 1, 2$, and a quadratic correction. (This notation is kept distinct from the fundamental conformal coupling β in $A(\phi)$.) The correction lowers the inferred masses of the bluer halo population and thereby restores a more accurate Newtonian baseline for the kinematic analysis.

3.3 Kinematic Analysis

The pipeline calculates the projected relative tangential velocity from the Gaia proper-motion difference and system distance, then compares it with the Newtonian circular velocity $v_c(s) = \sqrt{GM_{\text{tot}}/s}$. The central observable is the dimensionless velocity ratio $\tilde{v} = v_{\text{tan}}/v_c$.

The analysis takes the median of $\tilde{v}$ in logarithmic separation bins and normalizes the resulting profile by the mean of the first five bin medians, corresponding approximately to the 50–270 AU screened core. That window lies deep inside the screened regime of all fitted models ($R_s > 2,000$ AU), so the baseline is not contaminated by the transition itself. The choice of baseline affects the apparent amplitude of the outer-bin enhancement: normalizing at 50–270 AU yields $\tilde{v}_{\text{out}}/\tilde{v}_{\text{in}} \approx 1.37$, whereas normalizing at $\sim 500\text{--}1,200$ AU (where the transition is already underway) yields ≈ 1.24 , and normalizing at $\sim 1,200\text{--}2,400$ AU yields ≈ 1.16 . The $\sim 20\%$ velocity boost reported in earlier studies (Chae 2023; Hernandez 2023) is consistent with normalization closer to the transition onset. The present analysis normalizes deeper into the screened core so that α_{sat} captures the full transition amplitude rather than the residual above an already-elevated baseline.

At the shortest separations ($s \lesssim 100$ AU), unresolved spectroscopic binaries within one or both components can inflate the photometric mass estimate—by blending the luminosity of a hidden companion—while leaving the proper-motion difference largely unaffected. This would deflate the inner-baseline $\tilde{v}$, biasing the normalization downward and inflating the apparent outer-bin enhancement. The $\text{RUWE} < 1.2$ cut removes many such systems, and averaging over five normalization bins (50–270 AU) dilutes any residual contamination; nonetheless, the effect contributes an unquantified systematic at the few-percent level on α_{sat} . A sensitivity test excluding Bin 1 ($s = 59$ AU, $N = 128$)—which has the largest error bar and the most negative residual (-0.10)—shifts R_s by less than 2% and α_{sat} by less than 0.5%, confirming that the fit is not driven by the sparsest inner bin.

Because Gaia provides only sky-plane proper motions, $\tilde{v}$ is a projected quantity. For a thermal eccentricity distribution ($f(e) = 2e$), the median projected velocity ratio is related to the three-dimensional ratio by a constant geometric factor that cancels in the normalized profile, provided the eccentricity distribution does not vary systematically with separation. Dynamical processing by the Galactic tide could introduce a weak separation dependence at wide s ; any such effect would alter α_{sat} at the percent level but would not shift R_s , which is determined by the separation at which the profile departs from unity.

Bin-level uncertainties are estimated as follows. Within each logarithmic separation bin, the standard error of the median is computed analytically as $\sigma_{\text{med}} = 1.253 \sigma_{\text{bin}} / \sqrt{N}$, where σ_{bin} is the intra-bin standard deviation of $\tilde{v}$ and N the bin count. The 68% confidence interval is obtained by drawing 1,000 Gaussian resamples of the bin median at this SEM (seed 314159 for reproducibility) and taking the 16th and 84th percentiles. This is a parametric procedure rather than a true nonparametric bootstrap of the individual $\tilde{v}$ values; for heavy-tailed distributions the analytic SEM of the median may underestimate the true scatter, which contributes to the elevated χ^2_ν discussed in Section 4. The canonical phenomenological TEP recovery profile $\tilde{v}(s) = 1 + \alpha_{\text{sat}}(1 - e^{-s/R_s})$ is then fit alongside sigmoid, thin-shell, and double-exponential alternatives (Table 4.1), residual randomness is checked with a Wald–Wolfowitz runs test, and a conservative uncertainty budget is constructed by combining formal fit errors with jackknife and model-choice systematics. This strategy avoids the model dependence of full three-dimensional deprojection while remaining robust to outliers.

4. Results: The Screening Transition

The central prediction of TEP is a distinct kinematic transition in wide binaries at the separation where the source-charge suppression of the locally observable Temporal Shear continuously gives way and the conformal-sector gradient $\nabla \ln A(\phi)$ becomes kinematically active above the Keplerian baseline.

Table 4.0: Discriminating model comparison. The claim of this paper is not merely a velocity excess at large separation, but a specific three-feature pattern—resolved transition scale, bounded saturation amplitude, and environmental ordering—that must be reproduced simultaneously.

Model	Free parameters	Fits transition?	Fits saturation?	Fits environmental ordering?	Status
Newtonian flat profile	low	no	no	no	rejected ($\Delta\chi^2 = +14,845$)
Constant velocity boost	low	partial	yes	no	rejected ($\Delta\chi^2 = +3,583$)
MOND (no EFE)	medium	partial	partial	weak	disfavored ($\Delta\chi^2 = +7,195$ to $+10,346$)
MOND + EFE	medium	partial	partial	environment mismatch	tested ($\Delta\chi^2 = +540$ best case)
Triple contamination	high	possible boost	morphology?	ordering?	tested (fails to match profile shape)
TEP screening	medium	yes	yes	yes	favored under TEP assumptions

The central risk is that unresolved multiplicity, photometric mass bias, or selection effects could mimic part of the velocity excess. The discriminating claim of this paper is therefore not merely an excess at large separation, but a transition morphology and environmental ordering that must be reproduced simultaneously. Table 4.0 frames the empirical burden: any viable alternative must match not one but three distinct features—the characteristic transition scale ($R_s \approx 2,600$ AU), the bounded saturation amplitude ($\alpha_{\text{sat}} \approx 0.37$), and the environmental ordering (denser environments $\rightarrow$ larger R_s)—none of which are guaranteed to co-occur under generic systematic explanations.

4.1 The Transition Scale

Applying the methodology described above to the highly purified Gaia DR3 sample yields a clear and interpretable velocity profile $\tilde{v}$ with three recognizable regimes:

- 1. The Screened Regime ($s < 500$ AU):** The velocity profile is consistent with pure Keplerian expectations ($\langle \tilde{v} \rangle \approx 1$), deep within the screened regime.
- 2. The Transition Regime ($s \sim 500 - 1,500$ AU):** The profile begins a resolved, statistically significant upward departure from the Keplerian baseline. By fitting the profile with the canonical TEP generalized transition function, the analysis yields a best-fit screening scale of $R_s = 2,646 \pm 182$ AU, where the quoted uncertainty is the formal statistical fit error.
- 3. The Large-Separation Regime ($s > 5,000$ AU):** The velocity profile approaches a broad plateau approximately 35% to 40% above the inner baseline ($\tilde{v}_{\text{out}}/\tilde{v}_{\text{in}} \approx 1.35$ to 1.40). This behavior is consistent with the fitted saturation amplitude $\alpha_{\text{sat}} = 0.366 \pm 0.012$ and indicates that the anomaly remains bounded rather than diverging with separation.

The canonical fit yields $\chi^2 = 86.3$ for 17 degrees of freedom (19 bins minus 2 parameters), giving a reduced $\chi^2_\nu = 5.1$. The elevated χ^2_ν reflects bin-to-bin scatter beyond the diagonal error model. A residual autocorrelation analysis identifies the source: the standardized residuals exhibit significant lag-1 autocorrelation ($\rho_1 = 0.49$, $z = 2.14$, Durbin–Watson = 0.97), indicating that adjacent bins tend to deviate in the same direction—consistent with spatially correlated substructure such as moving groups or

distance-correlated completeness variations. The Wald–Wolfowitz runs test ($p = 0.485$) does not flag this because it tests only the sign sequence, not the magnitude correlation.

To properly account for this correlated structure, two covariance-aware refits are performed. First, an AR(1) Generalized Least Squares (GLS) refit uses the covariance matrix $C_{ij} = f^2 \sigma_i \sigma_j \rho_1^{|i-j|}$, where $f = \sqrt{\chi_\nu^2}$ inflates the diagonal to match the observed scatter. This yields $R_s = 2,692 \pm 181$ AU and $\alpha_{\text{sat}} = 0.366 \pm 0.011$ ($\chi_{\nu, \text{AR}(1)}^2 = 0.97$). Second, a Gaussian Process (GP) regression with a squared-exponential kernel in $\log_{10}(s)$ space models the covariance without imposing a rigid parametric structure. The GP marginal likelihood optimizes two hyperparameters: signal amplitude $\sigma_f = 0.015$ and correlation length $\ell = 0.18$ dex. The GP-covariance refit yields $R_s = 2,823 \pm 229$ AU and $\alpha_{\text{sat}} = 0.373 \pm 0.010$ ($\chi_{\nu, \text{GP}}^2 = 1.01$), with a marginal log-likelihood improvement of 2.9 over AR(1) model. Both covariance models absorb the excess scatter, both preserve the model-comparison hierarchy, and neither shifts the parameters outside the existing systematic uncertainty budget (± 609 AU). Throughout the remainder of this paper, the diagonal-fit values ($R_s = 2,646 \pm 182$ AU, $\alpha_{\text{sat}} = 0.366 \pm 0.012$) are reported as the primary results, with the AR(1) and GP refits serving as covariance robustness checks.

Relative to a flat Newtonian profile the fitted screening curve improves the description by $\Delta\chi^2 = 14,845$, 1,073, or 1,346, and relative to a separation-independent constant boost by $\Delta\chi^2 = 3,583$, 284, or 451. Under all three error models the observed signal is not merely elevated in amplitude; it is organized in separation in the specific way expected for a screened transition. When jackknife stability and transition-shape freedom are treated conservatively as systematics, the total uncertainty broadens to ± 609 AU, but the finite-separation onset remains intact.

Expanded model comparison reinforces that interpretation. Table 4.1 summarizes fit statistics for all eleven models considered, including four MOND variants fit directly to the same binned profile using per-bin median masses. Among the smooth-transition alternatives, a sigmoid is decisively worse than the canonical phenomenological TEP recovery profile ($\Delta\chi^2 = +131.5$), and the thin-shell exterior profile $f(s) = 1 - (R_s/s) e^{-m_{\text{bg}}(s-R_s)}$ itself—the form stated in Section 2 rather than a step-like proxy—is rejected still more strongly ($\Delta\chi^2 = +613.0$): its sharp onset at R_s followed by the slow geometric (R_s/s) recovery cannot reproduce the observed smooth approach to the saturation plateau, and the fit drives m_{bg} to its lower bound in a failed attempt to soften the tail. Meanwhile a double-exponential fit—which adds a shape exponent—achieves a lower raw χ^2 ($\Delta\chi^2 = -33.9$) and is preferred by AIC. The data therefore contain some transition-shape information beyond what the two-parameter canonical model captures.

The physical origin of this extra shape information is identified by a mass-convolved TEP model (Table 4.1, row 4). The derived screening-radius law (Equation 5) gives $R_s \propto M^{1/2}$; the sample spans a broad mass range ($M \approx 0.1\text{--}5.7 M_\odot$, $\sigma_M/M \approx 0.42$), so each separation bin contains a distribution of screening radii rather than a single value. The mass-convolved model integrates over the per-bin mass distribution: $\tilde{v}(s) = 1 + \alpha_{\text{sat}} \langle 1 - e^{-s/R_s(M)} \rangle_M$, with $R_s(M) = R_s^{\text{ref}} (M/M_{\text{ref}})^{1/2}$. This model has the same two free parameters (α_{sat} , R_s^{ref}) as the single-scale exponential, yet it captures the bulk of the double-exponential's advantage ($\Delta\chi^2 = -25.9$ versus -33.9 ; 76% of the improvement) without any additional free parameter. The remaining gap is consistent with environmental broadening from the $|Z|$ distribution within the sample. The double-exponential's extra shape parameter is therefore not evidence against TEP; it is a phenomenological proxy for the mass-dependent screening that TEP qualitatively predicts. The self-screening test (Section 6.4) independently confirms the mass dependence: $\alpha_{\text{sat}}(M)$ follows an exponential screening form with $M_{\text{screen}} = 0.46 \pm 0.04 M_\odot$.

The canonical single-scale exponential is nonetheless retained as the primary model because it is the minimal function satisfying the three qualitative TEP constraints (Section 2.3) and because both the mass-convolved and double-exponential variants agree on the physically meaningful parameters: onset scale ($R_s^{\text{ref}} \approx 2,822$ and $3,176$ AU versus $2,646$ AU) and saturation amplitude ($\alpha_{\text{sat}} \approx 0.377$ and 0.389 versus 0.366). The AIC-preferred models sharpen the transition but do not relocate it, and the spread among them is absorbed into the systematic uncertainty budget (± 609 AU total).

The overfitting diagnosis proceeds as follows. With $\chi_\nu^2 = 5.1$ on 19 bins, the diagonal error model underestimates the true scatter because residual substructure introduces correlated fluctuations that a per-bin variance estimate does not capture. A third free parameter—the shape exponent in the double-exponential—has the flexibility to track those bin-to-bin fluctuations, reducing χ^2 by fitting noise rather than signal. The clearest test of this interpretation is the inflated-error scheme: when bin uncertainties are scaled up by $\sqrt{\chi_\nu^2}$ to honestly reflect the observed scatter, the double-exponential's advantage collapses from $\Delta\chi^2 = -33.9$ to only -7 (Table 4.1, final column). The mass-convolved TEP, which has no extra free parameter, retains $\Delta\chi^2 = -5$ under the same inflation—confirming that its improvement reflects real mass-dependent broadening rather than noise-fitting. Once the error budget is corrected, the extra shape parameter of the double-exponential buys little leverage on the physically meaningful scale beyond what the mass-convolved model already provides. The canonical two-parameter model is therefore the appropriate primary description: it captures the transition without absorbing sample-specific scatter into a nuisance shape parameter.

The MOND comparison is particularly informative. In the simplest treatment, the ν -function (Famaey & Binney 2005; Milgrom 1983) with per-bin median masses and a single free parameter a_0 recovers a characteristic acceleration scale of the right order, confirming that the MOND transition scale is present in the data. However, without the External Field Effect (EFE) both variants are catastrophically rejected because the ν -function predicts $\tilde{v} \propto s^{1/2}$ in the deep-MOND regime while the data saturate.

Incorporating the EFE via the angle-averaged QUMOND prescription (Milgrom 2010; Famaey & McGaugh 2012) substantially improves the MOND fits by introducing saturation at large separations. With g_e fixed at the solar-neighborhood value the simple ν -

function drops to $\chi^2 = 1,625$ ($\Delta\chi^2 = +1,539$ versus TEP). Even with g_e treated as a second free parameter, the best MOND+EFE fit (simple ν) gives $\chi^2 = 627$ ($\Delta\chi^2 = +540$ versus TEP). The failure remains threefold: the preferred acceleration scale is driven far away from the canonical value, the transition shape is too steep, and the EFE-limited plateau undershoots the observed saturation.

The data require a finite, saturating transition whose shape and amplitude match the TEP screening family. The spread among saturating transition models—including the mass-convolved TEP, which explains the bulk of the double-exponential's advantage as a physical mass-broadening effect—is absorbed into the systematic uncertainty budget rather than interpreted as evidence against screening itself.

Because $\chi_\nu^2 = 5.1$ indicates that the diagonal error model underestimates the true scatter, a conservative robustness check inflates each bin uncertainty by a factor of $\sqrt{\chi_\nu^2}$, which forces $\chi_\nu^2 = 1$ for the TEP fit by construction. Under this inflated-error scheme every $\Delta\chi^2$ in Table 4.1 scales down by the same factor. The key comparisons remain decisive, confirming that the model-comparison hierarchy is not an artifact of the error model.

Table 4.1: Model comparison summary for the 19-bin velocity profile. k is the number of free parameters; $\text{AIC} = \chi^2 + 2k$. The mass-convolved TEP model integrates over the per-bin mass distribution with $R_s(M) \propto M^{1/2}$ (Equation 5), using the same two free parameters as the single-scale exponential. The final column reports $\Delta\chi^2$ under inflated bin errors ($\sigma \rightarrow \sigma\sqrt{\chi_\nu^2}$, forcing $\chi_\nu^2 = 1$ for the TEP fit), providing a conservative lower bound on all model-comparison significances. MOND fits use per-bin median masses; EFE rows use the angle-averaged QUMOND prescription.

Model	k	Transition (AU)	Amplitude	χ^2	dof	χ_ν^2	$\Delta\chi^2$ vs TEP	AIC	$\Delta\chi^2$ (inflated)
Flat Newtonian	0	—	—	14,932	19	786	+14,845	14,932	+2,923
Constant boost	1	—	0.180	3,669	18	204	+3,583	3,671	+705
TEP exponential	2	$2,646 \pm 182$	0.366 ± 0.012	86.3	17	5.1	0	90.3	0
Mass-convolved TEP	2	2,822	0.377	60.4	17	3.6	-25.9	64.4	-5
Sigmoid	3	2,019	0.356	217.9	16	13.6	+131.5	223.9	+26
Thin-shell	3	653	0.336	699.3	16	43.7	+613.0	705.3	+121
Double-exponential	3	3,176	0.389	52.4	16	3.3	-33.9	58.4	-7
MOND standard ν	1	—	unsaturated	10,432	18	580	+10,346	10,434	+2,037
MOND simple ν	1	—	unsaturated	7,281	18	405	+7,195	7,283	+1,417
MOND simple ν + EFE	1	—	EFE-limited	1,625	18	90	+1,539	1,627	+303
MOND simple ν + EFE (free g_e)	2	—	EFE-limited	627	17	37	+540	631	+106

Supplemental controls show that this transition is not confined to a single demographic slice. When the sample is split by mass ratio and by primary mass, all four half-samples continue to prefer a finite screening transition over a constant boost, with $\Delta\chi^2$ values ranging from 1,033 to 2,234. The saturation amplitude α_{sat} varies systematically with primary mass: the high-primary-mass half yields $\alpha_{\text{sat}} = 0.237$, the low-primary-mass half $\alpha_{\text{sat}} = 0.509$. This monotonic progression—weaker self-screening in lower-mass binaries revealing a larger fraction of the unsuppressed conformal-sector response—is qualitatively consistent with the independently measured compact-object conformal response hierarchy from the millisecond-pulsar spin-down analysis (screened full response $\tilde{\kappa}_{\text{MSP}} \sim 2.9 \times 10^4$, Paper 10; the unsuppressed geometric factor $\sim 10^6$ is the dense-environment upper bound; Smawfield 2026a; see Section 6.4). A stricter radial-velocity consistency subset (6,117 systems with measured component radial velocities, small formal errors, and mutual consistency) yields $R_s = 7,709 \pm 3,222$ AU and $\Delta\chi^2 = 33.6$ relative to a constant boost. The central R_s is $\sim 3\times$ the full-sample value, but the dramatic reduction in sample size leaves the fit poorly constrained: the full-sample $R_s = 2,646$ AU lies within the 2σ interval (1.6σ). Even in this small subset the data still prefer a finite screening transition, preserving the qualitative morphology of the signal.

Direct null controls reinforce the same conclusion. After scrambling the observed $\tilde{v}$ values globally—and again within distance quartiles to preserve large-scale distance structure—none of 10,000 valid realizations reproduced the observed improvement of the screening profile over a flat Newtonian or constant-boost description ($p = 1.0 \times 10^{-4}$ and $p = 1.0 \times 10^{-4}$, respectively). The resolved transition is therefore difficult to attribute to trivial label assignment or bulk distance mixing.




Figure 4.1: Upper panel: the observed dimensionless velocity profile $\tilde{v}$ as a function of projected separation, shown with 68% parametric-resampling confidence intervals and the best-fit canonical TEP screening model ($R_s = 2,646 \pm 182$ AU statistical fit uncertainty). Lower panel: residuals relative to that canonical fit. The transition is strongly preferred over both a flat Newtonian profile and a separation-independent constant boost, while broader smooth-transition fits still preserve the same few-thousand-AU onset scale.

Table 4.2: Bin-level velocity profile data. s is the geometric-mean projected separation; N the number of systems; $\tilde{v}$ the normalized median velocity ratio; σ the parametric-resampling uncertainty of the normalized median; $\tilde{v}_{\text{model}}$ the canonical TEP fit value. All 19 bins are logarithmically spaced between 50 and 30,000 AU.

Bin	s (AU)	N	$\tilde{v}$	σ	$\tilde{v}_{\text{model}}$	Residual
1	59	128	0.9068	0.0795	1.0081	-0.1013
2	83	450	1.0428	0.0442	1.0113	+0.0316
3	116	1,254	0.9948	0.0235	1.0157	-0.0209
4	162	2,955	1.0336	0.0158	1.0218	+0.0118
5	227	6,194	1.0220	0.0103	1.0301	-0.0081
6	319	11,267	1.0270	0.0077	1.0415	-0.0145
7	446	17,802	1.0610	0.0063	1.0567	+0.0043
8	625	25,383	1.0930	0.0055	1.0769	+0.0161
9	875	32,639	1.1253	0.0052	1.1029	+0.0224
10	1,225	38,010	1.1410	0.0050	1.1355	+0.0055
11	1,715	40,003	1.1824	0.0050	1.1744	+0.0080
12	2,402	38,660	1.2153	0.0053	1.2181	-0.0029
13	3,363	33,943	1.2359	0.0058	1.2631	-0.0272

Bin	s (AU)	N	$\tilde{v}$	σ	$\tilde{v}_{\text{model}}$	Residual
14	4,709	27,628	1.2846	0.0066	1.3040	-0.0194
15	6,594	21,889	1.3291	0.0076	1.3354	-0.0064
16	9,233	16,342	1.3704	0.0088	1.3545	+0.0158
17	12,929	11,468	1.3942	0.0104	1.3629	+0.0312
18	18,105	7,733	1.3860	0.0120	1.3653	+0.0207
19	25,352	4,620	1.3499	0.0147	1.3657	-0.0158

4.2 Forward-Model Test of the Derived Operator

A stronger test than profile fitting asks whether the operator derived in Section 2.1—rather than a phenomenological transition family—reproduces the observed profile when propagated through the full forward model. The catalog-level forward model deprojects each of the 341,315 binaries' projected separation, assigns the per-binary shell radius from the derived law $r_* = \sqrt{GM/g_t}$ (with lognormal scatter), carries the Galactic ambient floor S_{env} (with environmental scatter), and includes unresolved-triple contamination with the vector-addition kinematics of Section 3. The binned median prediction is then compared against the 19-bin measurement for two classes of candidate transition profile: the phenomenological power-law family $[1 + (R_s/s)^k]^{-1}$ over $k = 0.8 \dots 4$ and the exact flux-conserving profile $y(s)$ of the resolved operator (the implicit solution of $y[1 + y^2(r_*/s)^4] = 1$).

The result is direct support for the derived transition law. The exact $y(s)$ profile—the comparable-mass saturated-ambient regime of the nested operator, $\mathcal{R} = S_{\text{env}}^2 y \sim s^{4/3}$ —achieves $\chi_{\text{red}}^2 \approx 8.4$ under the full forward model, decisively better than the steep hierarchical form ($k = 4$: $\chi_{\text{red}}^2 = 18.3$) and the steeper power laws ($k = 2.5, 3$: 15.3, 16.3), and comparable to the best unconstrained interpolation in the family ($k = 1.3$: $\chi_{\text{red}}^2 = 3.9$), whose fitted slope is itself the effective exponent of the derived profile through the transition region. The ambient-floor parameter is simultaneously recovered at $S_{\text{env}} \approx 0.45$, matching the squared Galactic vertex factor $q_{\text{GAL}}^2 \approx 0.43$ implied by the plateau amplitude (Section 2.1): the forward model independently returns the plateau's derived value rather than requiring it as input. The resolved operator therefore passes the catalog-level test on both axes it predicts—the $s^{4/3}$ transition steepness and the ambient-floor-limited saturation—while the residuals against the ideal power law reflect the forward model's simplified ambient-scatter prescription rather than a failure of the profile itself.

5. Environmental Modulation: The Discriminating Test

A defining prediction of TEP is environmental screening. Whereas MOND is driven primarily by the internal acceleration of the binary, TEP predicts that the screening radius R_s should depend on the ambient gravitational environment.

5.1 Galactocentric Stratification

Under TEP, binaries embedded in more strongly screening environmental states remain suppressed to larger separations than binaries in more weakly screening environments. In the density-scaling benchmark adopted here, proximity to the Galactic midplane provides the corresponding higher-density environmental projection. Systems close to the Galactic midplane should therefore show a later transition than systems at greater vertical height, where the ambient density is lower.

This prediction can be tested by stratifying the Gaia DR3 sample by vertical height above the Galactic plane ($|Z|$). The midplane subsample ($|Z| < 100$ pc) lies well within the thin-disk scale height (~ 300 pc), where baryonic density is highest. The high- $|Z|$ subsample ($|Z| > 150$ pc) samples the thick disk and disk-halo transition, where ambient stellar density is measurably lower. A 100–150 pc buffer excludes systems with ambiguous classification. Using metallicity-corrected masses and parametric-resampling uncertainty propagation, the analysis yields:

- **Midplane / High Density ($|Z| < 100$ pc):** $R_s = 7,131 \pm 1,341$ AU.
- **High- $|Z|$ / Low Density ($|Z| > 150$ pc):** $R_s = 4,662 \pm 196$ AU.

Because the two subsamples have different population mixes, allowing α_{sat} to float freely in each would introduce an amplitude-scale degeneracy that complicates the comparison of R_s . The saturation amplitude is therefore fixed at $\alpha_{\text{sat}} = 0.4$ for both subsamples, close to the full-sample best fit (0.366), so that R_s absorbs only the transition-scale information. Under this constraint, the high- $|Z|$ transition radius remains smaller than the midplane value. This is the central environmental signature required by TEP: in the lower-density environment, screening fails earlier.

When R_s and α_{sat} are fitted independently in each subsample, their strong covariance prevents either transition scale from being identified separately. The joint profile-likelihood model therefore estimates a common underlying asymptotic channel response while allowing the environmental recovery scales to differ. The shared amplitude is not merely a statistical convenience: under the TEP screening semantics the saturation amplitude is the environment-shared asymptotic Temporal-Shear response, bounded by the Galactic ambient floor (Section 2.1), while the environment modulates the *scale* at which that response recovers. A model that

permits different plateaus would test a different hypothesis—an intrinsically environment-dependent response amplitude—than the screening transition TEP predicts. The joint fit yields $R_s^{\text{midplane}} = 4,912 \pm 792$ AU and $R_s^{\text{high-}|Z|} = 3,447 \pm 366$ AU, with $\Delta\chi^2 = 60.3$ ($p = 8.1 \times 10^{-15}$, 1 degree of freedom). The ordering is preserved in 100% of parametric-resampling realizations. Because α_{sat} is profiled from the data rather than imposed externally, the joint fit eliminates the arbitrary fixed-amplitude choice. Figure 5.3 further shows that the apparent free-fit inversion lies along the R_s - α_{sat} degeneracy valley, whereas the global joint solution robustly preserves the TEP environmental ordering. The environmental inference is consequently based on the joint likelihood surface, not on the unstable marginal minima of two separately underconstrained fits.

This conclusion does not rest on the metallicity correction alone. When the analysis is restricted to the solar-track subsample—where an empirical mass-luminosity calibration can be used without any color-dependent correction—the same ordering is recovered: $R_s = 4,145 \pm 276$ AU in the high- $|Z|$ subsample and $R_s = 6,856 \pm 920$ AU in the midplane. Permutation tests yield $p \leq 10^{-4}$ (0 exceedances in 10,000 realizations) for the full-sample ordering and $p \leq 10^{-3}$ for the solar-track control. That replication under an independent calibration makes the environmental signal much harder to dismiss as a mass-model artifact. The solar-track selection also answers a subtler demographic concern: because $|Z|$ correlates with age and metallicity in the Galaxy, the full-sample $|Z|$ subsamples could differ in their inner-binary demographics. Selecting only stars on the zero-correction solar-metallicity sequence holds the metallicity axis of that correlation fixed by construction, so the surviving ordering at $p \leq 10^{-3}$ cannot be attributed to a metallicity-correlated inner-binary population. The residual uncovariate is the age-correlated inner-binary channel, addressed explicitly in Section 6.

Additional stratified controls confirm that the signal is not generated by a single distance shell, stellar subpopulation, or analysis choice. Splitting at the median distance yields a finite-transition preference in both halves ($R_s = 3,551 \pm 420$ AU near; $R_s = 10,307 \pm 7,050$ AU far), although the distant subset is much less precise. Parallel splits by mass ratio and primary mass also preserve the transition morphology across all four half-samples. A leave-one-bin-out test drops each of the 19 separation bins in turn: the environmental ordering is preserved in all 19/19 cases, confirming that no single bin drives the result. Matched-subsample controls further isolate the signal from demographic confounders: restricting both populations to overlapping distance support, overlapping color support, or a stricter astrometric quality cut ($\text{RUWE} < 1.1$) all preserve the ordering. Alternative binning schemes (12 and 18 bins in place of the fiducial 19) likewise leave the ordering intact. Across all 5 matched controls, the ordering $R_s(\text{midplane}) > R_s(\text{high-}|Z|)$ is preserved in every case.

A mass-resolved measurement supplies an orthogonal check on the transition's physicality. Stratifying the full RUWE-selected sample into mass terciles (median masses 0.74, 1.20, and 1.70 $M_\odot$) and applying the same shared- α_{sat} joint fit yields transition scales that rise monotonically with system mass: $R_s = 2,778, 2,966,$ and $4,848$ AU ($\alpha_{\text{sat}} = 0.327 \pm 0.062$ shared). A log-space power-law fit gives the mass-scaling exponent $n = 0.64$ in $R_s \propto M^n$; bootstrap resampling returns a median exponent of 0.64, with 86% of resamples positive and 66% strictly monotone. A transition localized by a gravity-like field scale must grow with system mass; a mass-invariant artifact (detector, calibration, or population effect) would not. The measured exponent is consistent with the derived radius law $R_s = \sqrt{GM/g_t} \propto M^{1/2}$ of the canonical operator (Smawfield 2025a), matching the $M^{1/2}$ scaling expected under MOND as well. Therefore, while this positive mass trend is compatible with a mass-dependent gravitational field scale rather than a mass-invariant artifact, the primary discriminators between TEP and MOND remain the transition shape (flux-conserving recovery versus a deep-MOND tail) and the environmental ordering.

A complementary covariate-hardened test removes the profile-fit layer entirely and asks directly whether residual $\tilde{v}$ varies with environment once the known observational confounds are held fixed. The raw pool-level correlation between $\tilde{v}$ and $|Z|$ is distance-confounded—high- $|Z|$ pairs are systematically more distant and carry a higher proper-motion noise floor—so separation, distance, noise-floor, and mass are fixed in fine decile cells and the residual $\tilde{v}$ is regressed on $|Z|$ within each separation window with heteroscedasticity-consistent errors. The environmental coupling is absent at tight separations ($t = +0.7$ at $s < 0.5$ kAU), rises steadily through the transition window ($t = +2.2, +2.8, +3.7$ at $s = 0.5\text{--}1.5, 1.5\text{--}3, 3\text{--}5$ kAU), and reaches a slope of $+0.025$ per 100 pc at $s = 5\text{--}9$ kAU with $t = +12.7$ ($+11.9$ after adding a metallicity covariate). The coupling is monotone in $|Z|$ within every separation bin, and a direct high- versus low- $|Z|$ contrast at $s > 3$ kAU gives $+0.043 \pm 0.006$ (6.9σ , positive in 100% of resamples). The separation localization is the discriminating morphology: a uniform distance- or noise-related artifact cannot concentrate the $|Z|$ response specifically in the wide-separation window where the Temporal-Shear response is predicted to act. The interpretation is bounded by two candidate astrophysical alternatives—thick-disk orbital-population differences and Galactic-tide survival—which are quantitatively constrained and partially closed by the controls below rather than left fully open (Section 6).

A further control frees the eccentricity prior itself. The residual profile structure previously attributed to response-family mismatch proves to be dominated by the assumed eccentricity law: when the eccentricity distribution is fitted as a free mixture of thermal, superthermal, and uniform components jointly with the ambient response parameter u_{eff} on the selection-safe window (common random numbers and the measured per-pair proper-motion-error noise kernel), every $|Z|$ stratum selects the uniform-dominated component ($w_{\text{uniform}} \geq 0.8$), the absolute profiles fit at $\chi^2/\text{dof} \approx 1\text{--}8$, and the ambient ordering remains monotone ($u_{\text{eff}} = 0.40 \rightarrow 0.15$ from midplane to high- $|Z|$). Under the data-preferred mixture the density-scaling exponent is $p \approx 0.7$, giving an honest exponent band $p \approx 0.5\text{--}0.7$ across eccentricity-law systematics, centered on the $\rho^{1/2}$ benchmark but not precision-pinned. The noise-aware Newtonian null remains rejected at $\Delta\chi^2 \approx 800\text{--}1,900$ per stratum under the same best-fit mixture. Two features constrain the alternatives above: the fitted eccentricity weights are $|Z|$ -independent, so a thick-disk orbital-population explanation would require an e-law composition that varies with height yet produces identical mixture weights—a requirement the data do not

support—and the Newtonian null is evaluated under the identical mixture and noise kernel, so its rejection cannot be attributed to the prior choice.

Galactic-tide survival is quantitatively bounded at the relevant separations. At solar-circle enclosed mass ($M_{\text{enc}} \approx 10^{11} M_{\odot}$) the Jacobi radius of a typical pair is $r_J \approx 250\text{--}280$ kAU; pairs in the response window $s = 5\text{--}9$ kAU sit at a median $s/r_J \approx 0.023$, and even in the widest sample bin ($s = 9\text{--}30$ kAU) only 3.5% exceed $s/r_J = 0.1$. Beyond ~ 30 kAU the tide does become the dominant control (median $s/r_J \approx 0.13$), which bounds the reach of any pair-intrinsic inference rather than challenging the response-window signal. Differential tidal truncation between midplane and high- $|Z|$ pairs therefore cannot manufacture an effect localized fortyfold inside the tidal radius. The remaining tide-related channel—preferential survival of diffuse high- $|Z|$ pairs through earlier evolutionary truncation—affects pair demographics rather than the measured $\tilde{v}$ ordering within surviving bound pairs.

The widest-separation behaviour, where the hard purity cut truncates the $\tilde{v}$ tail, is recovered by replacing the cut with bound-probability weighting ($w = 1 - R_{\text{chance}}$): under soft weighting the normalized profiles no longer decline but rise monotonically to the last bin in every stratum, and the absolute high- versus low- $|Z|$ median contrast grows from $+0.040 \pm 0.005$ (8.4σ) at $s = 0.5\text{--}2$ kAU to $+0.094 \pm 0.018$ (5.3σ) at $s = 15\text{--}30$ kAU. A CDF-level mixture deconvolution then isolates the bound component explicitly, using the clearly unbound population ($R_{\text{chance}} > 0.2$) as the interloper shape. Because R_{chance} is a prior that can underestimate the true interloper fraction at wide s —the $\tilde{v} > 1.5$ tail allows up to $\sim 30\%$ contamination in the widest decade—the deconvolution is evaluated under interloper fractions of 0.11 (the R_{chance} estimate), 0.20, and 0.30. The bound-component median at $s = 20\text{--}50$ kAU remains 0.86–1.11 across strata, elevated above the $s < 2$ kAU baseline of 0.64–0.69 under every fraction tested. The outer-bin downturn in the hard-cut profiles is therefore an artifact of the purity threshold itself, not a property of the bound population; the transition radii quoted earlier from full-range fits are correspondingly biased low for the unsaturated dilute strata, and the selection-safe statements are the monotone ordering and the bound-component elevation reported here.




Figure 5.1: Comparison of velocity profiles for binaries in the dense Galactic midplane ($|Z| < 100$ pc) and the more dilute high- $|Z|$ population ($|Z| > 150$ pc, sampling the thick disk and disk-halo transition). The more dilute population transitions at smaller separation than the midplane profile, and this ordering is supported by permutation significance in both the full and solar-track analyses.

5.2 Chameleon-Completion Benchmark

TEP itself does not commit to a specific microscopic screening mechanism (Smawfield 2025a, §A4, §7), so the TEP-native macroscopic prediction for environmental modulation is that stronger environmental/source screening delays Temporal-Shear recovery. To make this quantitative the screening radius is parameterized as an empirical power law in ambient baryonic density, $R_s \propto \rho_{\text{amb}}^p$ (the inverse-power chameleon form that motivates this parameterization— $V(\phi) \propto \phi^{-n}$ giving $p = 1/(n+1)$ under thin-shell matching, Ratra & Peebles 1988; Khoury & Weltman 2004—is excluded under the TEP coupling convention $\beta_A = -1$; Smawfield 2025a, F1). In the density-scaling benchmark this becomes the specific relation denser environment $\rightarrow$ larger R_s . The specific power-law exponent is an empirical property of the environmental response, not of a specified microscopic completion. The ratio of transition radii at two Galactic heights then depends only on the ambient density ratio and the scaling exponent:

$$\frac{R_s(Z_2)}{R_s(Z_1)} = \left(\frac{\rho_{\text{amb}}(Z_2)}{\rho_{\text{amb}}(Z_1)} \right)^p \quad (9)$$

Using a standard three-component Galactic baryonic density model—stellar thin disk ($\rho_0 = 0.040 M_\odot \text{pc}^{-3}$, $h = 300 \text{ pc}$), thick disk ($0.005 M_\odot \text{pc}^{-3}$, $h = 900 \text{ pc}$), and gas disk ($0.050 M_\odot \text{pc}^{-3}$, $h = 150 \text{ pc}$; McKee, Parravano & Hollenbach 2015; Bovy 2015)—the median heights of the two subsamples ($|Z| = 47 \text{ pc}$ and 248 pc) correspond to densities $\rho = 0.075$ and $0.031 M_\odot \text{pc}^{-3}$ respectively, giving a density ratio of 0.41. The benchmark exponent $p = 1/2$ then predicts $R_s(\text{high-}|Z|)/R_s(\text{midplane}) = 0.64$, compared to the observed joint-fit ratio of 0.702 ± 0.074 . Calibrating with the midplane joint-fit value ($R_s = 4,912 \text{ AU}$), the $p = 1/2$ scaling predicts $R_s(\text{high-}|Z|) = 3,144 \text{ AU}$, consistent with the observed $3,447 \pm 355 \text{ AU}$.

Inverting the scaling relation to infer p from the data yields $p = 0.40 \pm 0.14$ (parametric resampling), consistent with the benchmark value $p = 1/2$ within uncertainty. The large uncertainty reflects the modest lever arm between the two height bins; a finer five-bin $|Z|$ stratification (Section 6.4, Ninth) gives $p = 0.50 \pm 0.03$ under the anchored one-parameter estimator, but a free two-parameter fit to the same five radii returns $p = 0.36\text{--}0.39$ with a rescaled uncertainty of order ± 0.1 , so the honest data-side statement is $p \approx 0.4 \pm 0.1$ — consistent with both the $\rho^{1/2}$ benchmark and the $\rho^{1/3}$ ambient channel that the noise-corrected inversion independently returns (0.40 ± 0.05). The key point for TEP is not that any particular completion is required, but that within this tractable benchmark the observed R_s ratio follows a simple density scaling without ad hoc tuning. Given a Galactic baryonic density model and one calibration point, the density-scaling benchmark renders ϵ_{env} semi-predictive at arbitrary heights, and that prediction is consistent with the observed high- $|Z|$ transition radius.




Figure 5.2: Predicted screening radius R_s as a function of Galactic height $|Z|$ from the empirical density-scaling relation with $p = 1/2$, calibrated from the midplane joint-fit value. The two data points (midplane and high- $|Z|$ joint-fit values) are consistent with the benchmark prediction.




Figure 5.3: 2D likelihood surface for α_{sat} versus R_s in the midplane (solid blue) and high- $|Z|$ (dashed red) subsamples. Contours show 1 σ , 2 σ , and 3 σ confidence regions. The free- α_{sat} minima (filled markers) reverse the R_s ordering but lie along an elongated, highly correlated degeneracy valley. The joint-fit minima with shared α_{sat} (open markers) preserve the physical ordering $R_s^{\text{midplane}} > R_s^{\text{high-}|Z|}$, confirming that the unconstrained inversion is an artifact of the shallow likelihood surface along the degeneracy axis rather than a physical reversal.

6. Discussion: Resolving the Controversy

The wide-binary debate is often framed as a stark choice: either gravity is universally modified at low acceleration, or the observed signal is primarily artifactual. TEP offers a more specific interpretation—one that preserves the reality of the anomaly while explaining why the signal should appear with a finite transition scale and a measurable environmental dependence.

Screening projection notice. Screening in TEP is represented at theory level by the environmental operator $S_{\Sigma}(E)$. Quantities such as ρ_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.

6.1 Why a Scale-Free MOND Interpretation Is Incomplete

The anomalous velocity boost is detected with high significance. The more discriminating question is whether it is adequately described as a scale-free enhancement or whether it instead follows a resolved transition in separation. The fitted screening profile is favored over a flat Newtonian description by $\Delta\chi^2 = 14,845$ and over a constant boost by $\Delta\chi^2 = 3,583$. A dedicated Newtonian orbital forward model, conditioned on the observed projected-separation distribution and marginalized over line-of-sight geometry, orbital phase, and eccentricity, produces an outer normalized median $\tilde{v} \simeq 1.01$ for the thermal case versus the observed $\simeq 1.37$ ($\chi^2 \simeq 14,531$ against the observed profile). The anomaly is therefore not merely an overall low-acceleration excess; it has a finite, separation-structured form that arises naturally in TEP because the scalar field remains environmentally suppressed until the binary enters a sufficiently diffuse local regime.

The model comparison sharpens that conclusion. A sigmoid transition is much worse than the canonical phenomenological TEP recovery profile ($\Delta\chi^2 = +131.5$), and the thin-shell exterior profile of Section 2.2 is rejected more strongly still ($\Delta\chi^2 = +613.0$), while a more flexible double-exponential reduces the raw χ^2 but still places the onset at a few thousand AU. The double-exponential's advantage is not evidence against TEP; it is a phenomenological proxy for the mass-dependent screening that TEP qualitatively predicts. A mass-convolved TEP model that integrates over the per-bin mass distribution with the derived radius law $R_s(M) \propto M^{1/2}$ (Equation 5) captures 76% of the double-exponential's χ^2 improvement using the same two free parameters, with the remainder attributable to environmental broadening from the $|Z|$ distribution within the sample. The self-screening test (Section 6.4, Twelfth) independently confirms the mass dependence. The data require a finite transition more strongly than they require any

particular phenomenological sharpness, and the spread between functional forms is best treated as model-choice systematic uncertainty rather than as evidence for a scale-free MOND-like uplift.

A direct MOND comparison makes this concrete. Fitting MOND interpolating functions to the binned profile using per-bin median masses (Table 4.1), the simple ν -function without the External Field Effect (EFE) still fails catastrophically, while the angle-averaged EFE via QUMOND reduces the discrepancy without curing it. Even the most generous MOND+EFE variant (simple ν with free g_e) remains worse than the TEP exponential by $\Delta\chi^2 = +540$. The failure is threefold: the preferred acceleration scale is driven away from the canonical value, the transition shape is too steep, and the EFE-limited plateau undershoots the observed saturation.

The near-coincidence between a_0 and the TEP screening transition is not accidental. Smawfield (2025g) showed that the SPARC rotation-curve database yields a characteristic transition acceleration $g_{\text{TEP}} \approx 5 \times 10^{-10} \text{ m/s}^2$, within a factor of four of a_0 , while Smawfield (2025e) showed that TEP produces the cosmological transition scale $g_t = cH_0/(2|\beta_A|) \approx 3.4 \times 10^{-10} \text{ m/s}^2$ of the same order as a_0 (within a factor of three), though without independently deriving the exact MOND numerical value. The distinction between the two frameworks is therefore not merely one of scale but of morphology: TEP produces a gradual, bounded enhancement whose amplitude is set by the local scalar-field value, whereas acceleration-dependent interpolation—even with the Galactic external field included—cannot simultaneously reproduce the transition shape, onset scale, and saturation level.

6.2 The Saturation Scale

The canonical fit gives a transition separation of $R_s = 2,646 \pm 182 \text{ AU}$, where the quoted error is the formal statistical fit uncertainty; when jackknife stability and model-choice freedom are included conservatively, the total uncertainty broadens to $\pm 609 \text{ AU}$. For the sample median mass ($M \approx 1.2 M_\odot$), this transition scale corresponds to an effective binary density $\rho_{\text{eff}} \approx 9.2 \times 10^{-18} \text{ g/cm}^3$. Although the Temporal Topology saturation scale is far higher ($\rho_T \approx 20 \text{ g/cm}^3$; Smawfield 2025g), the binary sits deep inside the screened Galactic potential. In the notation of Section 2.2, $\rho_{\text{eff}} \simeq \rho_{\text{floor}} = \epsilon_{\text{env}} \rho_T$, with $\epsilon_{\text{env}} \sim 4.6 \times 10^{-19}$. As discussed there, ρ_T and the scalar-field equation that governs ϵ_{env} are both derived in earlier work; only the numerical evaluation of ϵ_{env} in the specific Galactic environment remains to be computed from first principles. The observed transition density marks the point where the binary's internal Newtonian potential becomes sub-dominant to the pre-screened ambient background. The genuinely predictive test is the independent cross-check that follows.

The characteristic screening scale $R_s \approx 2,646 \pm 182 \text{ AU}$ is not merely a curve-fitting parameter; it represents the first kinematic detection of the Galactic screening floor. The canonical anchor is now the derived one: the master-action closure's acceleration scale $g_t = cH_0/(2|\beta_A|) \simeq 3.4 \times 10^{-10} \text{ m/s}^2$ gives the radius law $R_s = \sqrt{GM/g_t} \simeq 4.7 \text{ kAU}$ at the sample mean mass on the isolated-vacuum background, a factor $\simeq 1.8$ above the fitted value—the expected direction, since the ambient Galactic floor displaces the transition inward (Section 2.2). Solving the source perturbation on the Galactic nonlinear background instead (the embedded-propagator check of step 017) compresses the same estimator to $\simeq 3.3 \text{ kAU}$, a factor $\simeq 1.26$ above the fitted value within the systematic band, with the half-excess marker coinciding with the observed $\simeq 1.7 \text{ kAU}$: the mean-scale offset is substantially the ambient-stiffness propagator effect, not a free environmental coefficient. The older empirical anchor provides a complementary cross-check: the SPARC-extracted $g_{\text{TEP}} \approx 5 \times 10^{-10} \text{ m/s}^2$ (Smawfield 2025g) predicts $R_s^{\text{pred}} \approx 3,929 \text{ AU}$ for the transition-bin median mass ($M \approx 1.3 M_\odot$; factor 1.48), tightening to a factor of 1.05 ($R_s^{\text{pred}} \approx 2,778 \text{ AU}$) once the Galactic external field is included as an additional screening floor ($\eta = 2$). The wide-binary anomaly is therefore quantitatively consistent with the cross-scale screening scale of the TEP conformal sector.

6.3 Sensitivity to the Mass-Luminosity Relation

The environmental ordering depends critically on the stellar mass estimates. If a solar-metallicity Mass-Luminosity Relation (MLR) is applied uniformly, the masses of metal-poor high- $|Z|$ stars are systematically overestimated, the Keplerian baseline $v_c = \sqrt{GM/s}$ is inflated, and the inferred velocity ratio $\tilde{v}$ in the high- $|Z|$ population is suppressed. The color-dependent MLR correction (Section 3.2) removes this bias by adjusting masses according to each star's offset from the disk color-magnitude ridge, but it now does so without circular beta calibration: spectroscopic metallicities are used if cached, otherwise a conservative external prior is adopted and stress-tested. The corrected analysis recovers the expected ordering: the high- $|Z|$ transition radius is smaller than the midplane value for both the full sample ($R_s = 4,662 \pm 196 \text{ AU}$ versus $7,131 \pm 1,341 \text{ AU}$) and the solar-track control ($R_s = 4,145 \pm 276 \text{ AU}$ versus $6,856 \pm 920 \text{ AU}$; $p \leq 10^{-4}$ for the full sample and $p \leq 10^{-3}$ for the solar track). That the solar-track control—which uses an empirical mass calibration with no color-dependent correction at all—independently recovers the same ordering makes the result difficult to dismiss as an MLR artifact.

Supplemental controls narrow the most common non-TEP interpretations still further. If the observed profile were driven mainly by one demographic subset, the transition would collapse under stratification by mass ratio, primary mass, or observational geometry. Instead, the signal persists across all four stellar-population half-samples—each remaining strongly preferred over a constant boost—and is still recovered in a stricter radial-velocity consistency subset of 6,117 systems, though with broader uncertainty owing to the much smaller sample size. Conversely, when catalog-level labels are scrambled, either globally or within distance quartiles, none of 10,000 valid realizations reproduces the observed screening preference ($p \leq 10^{-4}$). The data indicate that the Gaia DR3 anomaly is structured, population-robust, and difficult to reduce to a single calibration artifact.

6.4 Limitations and Outlook

Several caveats deserve explicit acknowledgment. First, the reduced $\chi^2_\nu = 5.1$ for the canonical diagonal fit indicates that the diagonal error model does not capture the full scatter. A residual autocorrelation analysis identifies significant lag-1 correlation ($\rho_1 = 0.491$, $z = 2.14$, Durbin–Watson = 0.97), consistent with spatially correlated substructure or distance-correlated selection effects. Two covariance-aware refits address this directly: an AR(1) GLS model ($\chi^2_\nu = 0.97$) and a Gaussian Process with a squared-exponential kernel in $\log_{10}(s)$ space ($\chi^2_\nu = 1.01$). The GP is preferred by marginal log-likelihood and finds short-range correlated scatter ($\sigma_f = 0.015$, $\ell = 0.18$ dex). Under both covariance models all model-comparison conclusions are preserved, and parameters remain within the existing systematic uncertainty budget.

A direct investigation of the physical origin of this scatter finds no compelling single-population explanation. The standardized residuals show only weak rank correlation with median heliocentric distance, distance spread, and $\tilde{v}$ kurtosis, while the lag-1 residual autocorrelation rises from $\rho_1 = -0.05$ in the nearest distance quartile to $\rho_1 = 0.56$ in the most distant quartile. The correlated structure therefore appears concentrated in the distant subsample, consistent with known Gaia DR3 astrometric systematics at larger heliocentric distances ($d > 250$ pc), where parallax zero-point corrections and scanning-law artifacts naturally introduce spatially correlated covariance, rather than a localized astrophysical contaminant.

Second, the null-control and environmental permutation tests use 10,000 realizations each, giving a p -value resolution floor of 10^{-4} . No realization in either ensemble reproduced the observed screening preference, so the significance statements are limited primarily by the permutation-grid resolution rather than by an assumed parametric error model.

Third, an injection-recovery test validates that the pipeline recovers a known screening signal and does not hallucinate one when none is present. Mock catalogs are now constructed from the step_012 projected Newtonian forward-model null (rather than from globally shuffled observed $\tilde{v}$ values), then injected with the observed TEP enhancement parameters ($R_s = 2,646$ AU, $\alpha_{\text{sat}} = 0.366$). This directly tests recovery against an orbital-population null that includes projection, phase-mixing, and eccentricity structure.

Fourth, dedicated triple-contamination forward models directly test the claim that unresolved hierarchies can mimic the observed profile, with the photocentric reflex velocity treated as a vector added to the orbital tangential velocity—the observed speed $|\mathbf{v}_{\text{orb}} + \mathbf{v}_{\text{phot}}|$ is positive-definite, so the kick can only elevate the median. The RUWE < 1.2 cut leaves residual triples with wide inner orbits whose kicks are individually large but diluted by median robustness: at 10% residual contamination the predicted outer-bin median reaches only $\tilde{v} \approx 1.05$ versus the observed $\tilde{v} \approx 1.366$, and the deficit shrinks smoothly with fraction rather than growing—closing only at $f_{\text{triple}} \approx 50\%$, an order of magnitude above the ~ 5 –10% residual fraction permitted by RUWE detection incompleteness. The Pittordis-faithful population model, which draws inner orbits and orientations from the Pittordis et al. (2025) distributions, performs still worse— $\tilde{v} \approx 1.02$ at 10% and only 1.11 even at 50% contamination—because the sky projection of the inner-binary reflex velocity further suppresses the observable kick. No physically plausible triple fraction closes the gap; the median statistic is too robust for residual hierarchies to reproduce the observed enhancement.

Fifth, the observable $\tilde{v}$ is a projected quantity. A Monte Carlo orbital simulation quantifies the residual eccentricity systematic on α_{sat} . Relative to the thermal case, the uniform distribution recovers about 14% higher α_{sat} and the most super-thermal case about -7% lower, while the realistic interval remains bounded between -5% and $+5\%$. The transition radius remains stable to about 7% across the full sweep, so the eccentricity systematic does not propagate materially into R_s or the model-comparison hierarchy.

Sixth, distance selection and mass calibration remain obvious places to look for failure modes. Both distance halves independently favor a finite transition, with the nearer subsample yielding $R_s = 3,551 \pm 420$ AU. Likewise, four alternative MLR prescriptions preserve the environmental ordering, with midplane R_s spanning 6,911 to 8,582 AU and high- $|Z|$ R_s spanning 4,156 to 7,215 AU. The result is therefore robust to the reasonable range of calibration choices already explored.

Seventh, the MOND comparison in Table 4.1 is already nontrivial, but it is not exhaustive. Marginalizing over full within-bin mass distributions or testing alternative EFE geometries may reduce the discrepancy somewhat. Even so, the present best MOND+EFE variant remains worse than TEP by $\Delta\chi^2 = +540$, so any viable rescue must overcome simultaneous failures in acceleration scale, transition shape, and outer-plateau amplitude.

Eighth, several quantities imported from other TEP papers—the Temporal Topology saturation scale $\rho_T \approx 20$ g/cm³, the characteristic acceleration $g_{\text{TEP}} \approx 5 \times 10^{-10}$ m/s² (Smafield 2025g), and the emergent MOND-scale acceleration $a_0 \sim cH_0$ (Smafield 2025e)—are drawn from preprints not yet independently peer-reviewed. The present analysis does not depend on the precise values of ρ_T or g_{TEP} for any primary result; those quantities enter only in the cross-scale consistency check of Section 6.2. Readers should regard that check as a promising but provisional link pending independent verification.

Ninth, the pre-screening factor ϵ_{env} is no longer fully post-hoc. Adopting the empirical density scaling $R_s \propto \rho_{\text{amb}}^p$ as a tractable benchmark (TEP itself does not commit to a specific microscopic mechanism; Smafield 2025a, §A4), the scaling relation predicts the two-bin environmental ratio within the quoted errors, and the inferred exponent from the two-bin fit is $p = 0.40 \pm 0.14$. A finer stratification into 5 equal-count $|Z|$ bins, each containing 68–263 systems, yields transition radii declining from $R_s = 8,670 \pm 1,988$ AU at median $|Z| = 22$ pc to $R_s = 4,326 \pm 252$ AU at median $|Z| = 347$ pc (a selection audit notes that these radii come from full-range fits whose outer bins above ~ 9 kAU carry the purity-cut truncation: restricted to the selection-insensitive window the dilute-strata profiles remain unsaturated at the window edge, so the quoted radii are effective values of the truncated profiles; the selection-safe statement is the monotone median- $\tilde{v}$ ordering reported in Section 4, which survives every threshold including none at all). A direct density-scaling fit to these five points—using the actual three-component Galactic density

model $\rho(Z)$ rather than a log-linear approximation—yields an exponent $p = 0.50 \pm 0.03$ under the anchored one-parameter estimator (amplitude fixed at the first bin). A free two-parameter fit to the same five radii returns $p = 0.36\text{--}0.39$, and the quoted ± 0.03 is the raw $\Delta\chi^2 = 1$ error of a fit with $\chi^2 = 10.7$ on 3 dof (rescaled ± 0.06); the honest data-side exponent is therefore $p \approx 0.4 \pm 0.1$, centered between the $\rho^{1/3}$ and $\rho^{1/2}$ ambient channels and matching the noise-corrected ambient inversion (0.40 ± 0.05). An independent inversion through the ambient response parameter u_{eff} , with the eccentricity law freed as a mixture (Section 5.1), returns $p \approx 0.7$ under the data-preferred uniform composition and an honest band $p \approx 0.5\text{--}0.7$ across eccentricity-law systematics; the two estimators jointly bracket the exponent in $0.4\text{--}0.7$, consistent with the $\rho^{1/2}$ benchmark. A complementary weighted linear regression of $\ln R_s$ versus $|Z|$ gives a slope $-0.00162 \pm 0.00036 \text{ pc}^{-1}$ ($t = -4.45$, significance 0.021), which is 97% of the $p = 1/2$ predicted slope (-0.00167 pc^{-1} for $h_{\text{eff}} = 300 \text{ pc}$). The fine- $|Z|$ stratification therefore not only confirms the downward trend but provides a precise, physically grounded measurement of the environmental density-scaling exponent, equal to the square-root scaling $p = 1/2$ within a few percent. This empirical scaling now has a mechanism-side check (step 017, supply-versus-demand): the free- α amplitudes invert to a required ambient contrast of $\simeq 36\%$ rising toward the midplane, and the corpus's density-set ambient supplies it — the disk density ratio between the strata's median heights is $\rho_{\text{mid}}/\rho_{\text{halo}} = 2.44$, giving 34.6% in the ambient variable under $u_{\text{eff}} \propto \rho^{1/3}$ against 35.8% required, with the now-identified $u_{\text{eff}} \propto \rho^{1/2}$ k-branch law (Paper 0, AUD-3 closure) bracketing the same requirement at 56% under the median-density mapping; the two-bin contrast is an estimator-level check that brackets both, while the five-stratum inversion fixes the exponent at $p \approx 0.46\text{--}0.5$. The linear-gradient estimator $u = |a|/g_t$ fails by construction ($\sim 4\%$ with the wrong sign: the radial floor is common to both heights and the vertical field adds with $|Z|$), but that estimator is the weak-field exterior reading — inside the nonlinear disk the field sits at its density-set equilibrium and the local acceleration is not the ambient state. The ordering therefore discriminates between ambient projections and selects the density-set one — the corpus's own in-medium response channel, now identified at coefficient level as the local k-branch law $u_{\text{amb}} = \sqrt{4\sqrt{2}\pi G \rho L / (K g_t)}$ with coherence length $L \simeq 0.84 \text{ kpc}$ — rather than a fitted covariate. A matched covariate audit (step 018) closes the demographic alternative on the observational side: nearest-neighbour matching of the disk and halo subsamples on total mass, projected separation, distance and $[\text{Fe}/\text{H}]$ reduces the ΔR_s contrast from $2\{, \}468 \text{ AU}$ to $1\{, \}935 \text{ AU}$ — only $\simeq 22\%$ absorbed — and the mass covariate runs the wrong way (the halo systems are the more massive, which naively predicts the larger halo R_s , opposite to observation); the ordering survives every matched experiment. The ambient projection entering the pair vertex is now derived rather than assumed: the embedded pair response on the constraint slice is independent of the local background density ($\lesssim 0.02\%$ across the disk range) and steep in the ambient boundary field (charge $\propto u_{\text{amb}}^{-2.0}$), so the vertex samples the landscape ambient map — identified as the local k-branch density response $\rho^{1/2}$ with the $L \simeq 0.84 \text{ kpc}$ coherence length reproducing all five strata within 10.1% (Paper 0, AUD-3 closure).

Tenth, the environmental comparison in Section 5.1 presents fixed- α_{sat} results as the primary protocol because the amplitude-scale degeneracy prevents clean isolation of R_s when α_{sat} floats independently in each subsample. A joint profile-likelihood fit eliminates this concern: both subsamples are modeled simultaneously with a single shared α_{sat} and separate transition radii, yielding $R_s = 4,912 \pm 792 \text{ AU}$ (midplane) versus $R_s = 3,447 \pm 366 \text{ AU}$ (high- $|Z|$) with $\alpha_{\text{sat}} = 0.332 \pm 0.020$. A likelihood ratio test gives $\Delta\chi^2 = 60.3$ ($p = 8.1 \times 10^{-15}$), and the ordering is preserved in 100% of parametric-resampling realizations.

A related estimator dependence deserves explicit statement: every stratified or subset estimate of R_s lies above the pooled headline value of 2,646 AU (midplane 7,131 AU and high- $|Z|$ 4,662 AU under the fixed- α_{sat} protocol; 4,912 and 3,447 AU under the shared- α_{sat} joint fit; $8,670 \rightarrow 4,326 \text{ AU}$ across the five $|Z|$ bins; 3,551 AU and 10,307 AU for the near and far distance halves; 7,709 AU for the radial-velocity consistency subset; and up to 3,743 AU across the normalization sweep). This is an estimator-offset phenomenon, not an internal contradiction: the pooled estimate is a free two-parameter fit to a single 19-bin profile that mixes all environments, whereas each stratified estimate applies a different protocol (fixed or shared α_{sat} , or a subsample-specific normalization window) to a homogeneous subpopulation. A pooled single-scale fit to a heterogeneous mixture is not obliged to return a value interior to the stratum range—a mixture of exponentials with different scales and amplitudes is not itself a single exponential, and the diffuse environmental tail contributes an early partial rise that pulls the pooled onset inward relative to any resolved stratum. The pooled value is retained as the headline scale for conservatism and continuity with the primary fit; it should be read as the minimum onset scale of the transition, while the environmentally resolved estimators return larger effective scales. Notably, the resolved estimates cluster nearer to the independent cross-scale prediction $R_s^{\text{pred}} \approx 3,929 \text{ AU}$ from g_{TEP} (Section 6.2) than the pooled value does, so resolving the estimator dependence moves the measurement toward, not away from, the TEP cross-scale anchor.

Eleventh, at the shortest separations unresolved spectroscopic binaries could in principle deflate the inner normalization and inflate the apparent outer enhancement. The normalization sensitivity sweep tests this directly by refitting the canonical model under every contiguous baseline window that remains plausibly within the screened core. Across those screened-core windows the recovered R_s ranges from 2,196 to 3,743 AU, with a maximum deviation of 41% above the fiducial value. The transition scale is therefore stable while the apparent saturation amplitude responds to the normalization, which is the opposite of the pattern expected from a baseline-deflation artifact. Crucially, while the apparent saturation amplitude α_{sat} modulates with the chosen baseline, the transition scale R_s remains bounded within the 2,196–3,743 AU range across the sweep. This demonstrates that the tested baseline-deflation variations cannot manufacture a fictitious onset scale.

Twelfth, the saturation amplitude α_{sat} varies systematically across the demographic half-samples in a pattern consistent with mass-dependent self-screening. The high-primary-mass subset yields $\alpha_{\text{sat}} = 0.237$, the full sample 0.366, and the low-primary-mass subset $\alpha_{\text{sat}} = 0.509$ ($\Delta\chi^2 = 2,234$ versus a constant boost for the strongest split). A quantitative self-screening model

$\alpha_{\text{sat}}(M) = \alpha_{0,\text{WB}} \exp(-M/M_{\text{screen}})$ fits these three data points with $\chi^2 = 1.25$ for one degree of freedom, yielding an illustrative unsuppressed wide-binary response amplitude $\alpha_{0,\text{WB}} = 1.74 \pm 0.26$ and a self-screening mass scale $M_{\text{screen}} = 0.46 \pm 0.04 M_{\odot}$. As a two-parameter fit to three points, this model is presented as illustrative rather than definitive evidence. The unbounded exponential extrapolation is not, however, the physically meaningful limit: the response amplitude is structurally bounded by $\alpha_{\text{sat}} = \sqrt{1 + 2\beta_A^2 \mathcal{S}_{\text{eff}}} - 1 \leq \sqrt{3} - 1$, and in the Galactic ambient the plateau is further capped by the ambient vertex product. Re-fitting the same three points with the bounded recovery form $\alpha_{\text{sat}}(M) = \sqrt{1 + 2\mathcal{S}_{\text{amb}} y(M)} - 1$, where $y(M) = [1 + (M/M_s)^p]^{-1}$ is the self-screening recovery fraction, interpolates the data exactly and yields $\alpha_{0,\text{WB}} = \sqrt{1 + 2\mathcal{S}_{\text{amb}}} - 1 = 0.65 \pm 0.09$ (Monte-Carlo propagation of the point errors; the saturated fit itself carries no degrees of freedom) with $M_s = 0.72 M_{\odot}$ and $p = 3.7$ —within the absolute ceiling and consistent with the ambient-floor-limited plateau. The apparent $\alpha_{0,\text{WB}} > 1$ value is therefore an artifact of the exponential ansatz extrapolating through the response ceiling; the bounded form demonstrates that the demographic trend requires no super-ceiling amplitude. Here $\alpha_{0,\text{WB}}$ is the asymptotic response amplitude of the wide-binary kinematic channel, not the microscopic conformal coupling β_A . Cassini, millisecond-pulsar timing, and wide binaries probe different environmentally projected responses of the same universal conformal sector through $\mathcal{S}_{\Sigma}(\mathcal{E})$: the Cassini PPN bound $\mathcal{S}_{\Sigma}^{(\odot)} \lesssim 1.15 \times 10^{-5}$ under the unified unstretched-space metric (5.8×10^{-6} isotropic) on the screened solar source charge (Smawfield 2025a, Section 7), the compact-object regime screened full response $\tilde{\kappa}_{\text{MSP}} \sim 2.9 \times 10^4$ extracted from millisecond-pulsar spin-down in dense globular clusters (Smawfield 2026a; the unsuppressed geometric factor $\sim 10^6$ characterizes the dense-environment upper bound), and the wide-binary response amplitude $\alpha_{0,\text{WB}}$ each sample a different environmental projection of the same underlying screening hierarchy. The relevant cross-scale check is qualitative: the same conformal-sector physics that screens the Temporal Shear in dense regimes and recovers it in low-density regimes underlies all three measurements. Notably, the self-screening model is itself a continuous exponential function of mass—there is no thin-shell step function or sharp screening boundary—consistent with the Temporal Topology framework in which self-screening operates via continuous flattening of the field profile rather than a thin-shell transition.

Thirteenth, the environmental subsamples are not exchangeable populations: Galactic $|Z|$ correlates with age, kinematic heating, and metallicity, so the midplane and high- $|Z|$ groups can differ in inner-binary demographics—period and mass-ratio distributions that modulate photocentric contamination independently of the RUWE screen. The metallicity axis of that correlation is controlled: the solar-track subsample restricts both $|Z|$ groups to the zero-correction solar-metallicity sequence, and the ordering survives there at $p \leq 10^{-3}$. The matched distance, color, and astrometric-quality controls further exclude the observable demographic axes. The surviving uncovariate is the age-correlated component of the inner-binary period distribution, which no catalog quantity proxies; it is bounded indirectly by the triple-contamination forward models (Fourth), which show that residual hierarchies elevate the outer median by at most ~ 0.05 – 0.08 at realistic fractions, and by the sign constraint that an age-correlated contamination gradient would need to elevate the high- $|Z|$ profile specifically at few-thousand-AU separations while leaving the midplane unaffected—a morphology no identified contaminant produces. A definitive control requires an inner-binary-resolved catalog (e.g., epoch astrometry or radial-velocity monitoring in DR4+); the channel is recorded as the leading open demographic caveat.

6.5 Systematic Controls Summary

The following table collects the robustness checks performed across the analysis pipeline. Each row tests whether the primary conclusions—a finite screening transition and environmental ordering—survive a specific perturbation to the data selection, error model, mass calibration, or binning procedure.

Control	What it tests	Result
Global scramble (10,000 realizations)	Can noise produce the transition?	$p \leq 10^{-4}$ (0 exceedances)
Distance-quartile scramble	Distance-correlated artifacts	$p \leq 10^{-4}$ (0 exceedances)
Injection-recovery (100 mocks)	Pipeline fidelity and false-positive rate	100% detection, 0 false positives in 100 null mocks
Solar-track control	MLR correction dependence	Ordering preserved
Quadratic, no-correction, $\beta = 1.0$, $\beta = 2.0$ MLR	MLR functional form	4/4 preserve ordering
Joint profile-likelihood fit	Fixed- α assumption	$\Delta\chi^2 = 60.3$, $p = 8.1 \times 10^{-15}$
α_{sat} sweep (0.30–0.45)	Sensitivity to fixed amplitude	5/5 preserve ordering

Leave-one-bin-out (19 bins)	Single-bin dominance	19/19 preserve ordering
Distance-matched subsamples	Distance distribution imbalance	Ordering preserved
Color-matched subsamples	Stellar population imbalance	Ordering preserved
Strict RUWE < 1.1	Triple contamination	Ordering preserved
Alternative binning (12, 18 bins)	Bin definition dependence	2/2 preserve ordering
AR(1) GLS refit	Bin-to-bin correlation	$\chi^2_\nu = 0.97$, conclusions preserved
GP covariance refit	Flexible covariance model	$\chi^2_\nu = 1.01$, conclusions preserved
Eccentricity sweep (5 distributions)	Projection systematic on α_{sat}	R_s stable $\lesssim 4\%$; α_{sat} bounded by the recovered sweep
Demographic half-samples ($4\times$)	Subpopulation dependence	All prefer finite transition
RV-consistency subset (6,117 systems)	Kinematic purity	Transition recovered
Density-scaling benchmark	ϵ_{env} post-hoc concern	$p = 1/2$ prediction within 1σ
Fine $ Z $ stratification (5 bins)	Density-gradient resolution	$p \approx 0.4 \pm 0.1$ (free-fit; anchored estimator 0.50 ± 0.03); slope significance 0.021 (weighted regression)
Triple forward model (6 fractions)	Can residual triples mimic signal?	Predicted outer $\tilde{v} \approx 1.05$ at 10%; gap closes only at implausible $f \approx 50\%$
Newtonian orbital forward model	Projection, phase, and eccentricity null	Thermal null outer $\tilde{v} \approx 1.01$ versus observed ≈ 1.37
Self-screening model	Mass-dependent α_{sat}	Exponential fit $\chi^2 = 1.25$ (1 dof)
Mass-stratified transition (3 mass bins)	Mass scaling of R_s	R_s rises with mass (2,778 $\rightarrow$ 4,848 AU); exponent $n = 0.64$, consistent with $M^{1/2}$
Age-correlated inner-binary demographics	Unproxied $ Z $ confound axis	Metallicity axis controlled (solar track $p \leq 10^{-3}$); age channel bounded by the corrected vector-addition triple models (realistic fractions supply $\lesssim 10\text{--}20\%$ of the outer enhancement; the Pittordis-faithful population cannot reproduce the profile even at 50% contamination)
Normalization sensitivity sweep	Baseline deflation by spectroscopic binaries	R_s stable within $\pm 41\%$ max deviation for screened-core windows
Spatial substructure identification	Physical origin of $\chi^2_\nu = 5.1$	Lag-1 ρ concentrated in distant quartile ($\rho_1 = 0.55$)
Pittordis-faithful triple model	Exact Pittordis et al. (2025) triple distributions	Predicted outer $\tilde{v} \approx 1.11$ even at 50%; cannot explain at any plausible fraction

No single control eliminates every conceivable systematic, but the breadth of this matrix—spanning data selection, error modeling, mass calibration, binning, projection, contamination, and theoretical prediction—makes it difficult to construct a single astrophysical or methodological scenario that simultaneously survives all tests while mimicking both the transition morphology and the environmental ordering.

6.6 Distinguishing TEP from Generic Scalar Screening

The density-scaling benchmark in Section 5.2 demonstrates that the observed environmental modulation follows a square-root power law in ambient baryonic density. The inverse-power chameleon form that historically motivates the parameterization $p = 1/(n+1)$ is excluded under the TEP coupling convention—with $\beta_A = -1$ and $\phi > 0$ near masses, $V \propto \phi^{-n}$ admits no admissible minimum (Smawfield 2025a, F1)—so the measured exponent is phenomenology rather than a claimed microscopic realization. A fair question is therefore: what distinguishes TEP from the scalar-screening family generally? TEP itself does not commit to a microscopic

completion (Smafield 2025a, §A4, §7); chameleon-type, Vainshtein, Galileon, DBI, and symmetron mechanisms remain candidates where they satisfy the F1–F6 constraints, with the defining ontology being continuous suppression of the Temporal Shear $\Sigma_\mu = \nabla_\mu \ln A$ by source-charge and environmental state. Four features are specific to TEP rather than generic to the broader scalar-screening family:

1. The cross-scale anchor. The transition acceleration $g_N(R_s) \approx 1.1 \times 10^{-9} \text{ m/s}^2$ at the observed screening radius is within a factor of ~ 2 of the independently measured $g_{\text{TEP}} \approx 5 \times 10^{-10} \text{ m/s}^2$ from SPARC rotation curves (Smafield 2025g), with no parameter adjustment. Generic stand-alone scalar-screening models have no independent cross-scale anchor for the wide-binary transition. In TEP, the absolute regime of the transition is independently anchored by the Temporal Topology saturation scale ρ_T and the SPARC-derived g_{TEP} ; the microscopic completion determines the detailed environmental mapping within that anchored regime.
2. The mass-dependent saturation amplitude. The monotonic progression $\alpha_{\text{sat}} = 0.237$ (high mass) $\rightarrow 0.366$ (full sample) $\rightarrow 0.509$ (low mass) is a natural consequence of TEP's conformal coupling, where self-screening by the binary's internal potential attenuates the velocity-profile amplitude. Standard chameleon models predict mass-dependent screening radii but not a specific amplitude hierarchy tied to an independently anchored cross-scale conformal-response architecture.
3. The emergent a_0 . TEP produces a cosmological acceleration scale $g_t = cH_0/(2|\beta_A|) \approx 3.4 \times 10^{-10} \text{ m/s}^2$ from the master-action closure (Smafield 2025e), within a factor of three of the empirical MOND value $a_0 \approx 1.2 \times 10^{-10} \text{ m/s}^2$. While TEP does not independently derive the exact MOND numerical value, the emergence of a cosmological acceleration scale of the same order as a_0 and g_{TEP} from the same scalar-field structure is not generic to stand-alone chameleon or symmetron models.
4. Continuous Temporal Topology and the proper-time signature. The standard chameleon thin-shell approximation (Khoury & Weltman 2004) concentrates the active scalar variation into a narrow shell and yields the corresponding thin-shell exterior response, producing a profile with a geometric R_s/s prefactor. TEP is not committed to thin-shell matching: its defining ontology is continuous Temporal Topology (Smafield 2025a, §7), where screening operates via the smooth spatial profile of $\ln A(\phi)$, the locally observable Temporal Shear $\Sigma_\mu = \nabla_\mu \ln A$ is suppressed continuously in dense environments rather than concentrated in a narrow shell, and the same conformal factor $A(\phi)$ that sources the kinematic enhancement also rescales matter-frame proper time as $d\tau/dt = A(\phi)$. TEP identifies the recovered conformal gradient specifically with dynamical proper-time geometry, tying the wide-binary kinematic response to the same Temporal-Shear sector used across the wider TEP corpus. Generic screened scalar models need not make this TEP-specific cross-sector identification. This predicts a smoother transition than thin-shell models; in the canonical wide-binary realization adopted here, that recovery is represented by the pure exponential rather than the thin-shell form $f(s) = 1 - (R_s/s) e^{-m_{\text{bg}}(s-R_s)}$. The data already favor this prediction: the thin-shell profile itself is rejected at $\Delta\chi^2 = +613.0$ relative to the TEP exponential, with the sigmoid step-like proxy likewise rejected at $\Delta\chi^2 = +131.5$ (Table 4.1). This morphological distinction separates the canonical TEP recovery profile from the standard thin-shell chameleon approximation tested here.

The TEP-native prediction is strictly macroscopic: more strongly screening environmental states more strongly suppress the locally observable Temporal Shear; in the density-scaling benchmark used here, this maps onto increasing ambient density. The measured exponent $p \approx 0.4 \pm 0.1$ under free fits ($p = 0.50 \pm 0.03$ anchored) is an empirical property of that environmental response. Alternative completions would manifest differently: Vainshtein derivative screening, for example, would map the Temporal Shear suppression to the Galactocentric potential gradient $\nabla\Phi$ rather than to local baryonic density ρ_{baryon} , producing a different $R_s(R_{\text{gc}})$ dependence; symmetron screening would introduce a density-dependent symmetry-breaking scale that yields a distinct environmental signature. TEP requires the geometric suppression of the Temporal Shear, not the specific microscopic Lagrangian that achieves it. The observational task is therefore not to confirm any particular screening mechanism per se, but to test whether the macroscopic suppression pattern matches the TEP prediction across environments.

These distinctions are currently anchored to quantities from the wider TEP series that are not yet independently peer-reviewed (Section 6.4, Eighth). The present five-bin $|Z|$ stratification yields a density-scaling exponent of $p \approx 0.4 \pm 0.1$ under free fits, consistent with the square-root benchmark; the most decisive future test is the predicted $R_s(R_{\text{gc}}, |Z|)$ map: within the density-scaling benchmark, this becomes a specific density-dependent gradient calibrated by the Temporal Topology saturation scale ρ_T and the cross-scale acceleration g_{TEP} , whereas a stand-alone screening model would require fitting both the density exponent and the response normalization as free parameters. Gaia DR4, with its expanded volume and improved astrometry, could enable the finer $|Z|$ stratification needed to distinguish these scenarios. An anisotropy test—comparing the screening transition along and perpendicular to the Galactic disk—would provide a further discriminator, since the TEP screening floor is controlled by the local environmental state $\mathcal{S}_\Sigma(\mathcal{E})$, rather than by a universal acceleration threshold; in the density-scaling benchmark used here, the dominant environmental projection is the local baryonic-density field. Table 6.3 summarizes the frozen Gaia DR3 empirical baselines alongside the prospectively specified, zero-per-system-parameter TEP predictions for Gaia DR4 prior to catalog release.

Table 6.3. Frozen Gaia DR3 baselines and prospectively specified Gaia DR4 predictions for wide-binary kinematics and epoch astrometry (

$$P_{iX} = \epsilon_{\text{reg}} + k\sqrt{\xi} + 2\xi, k = 16.03, \beta_A = -1).$$

Observable Channel	Frozen DR3 Baseline & Prospectively Specified TEP Prediction	Discriminating Null / Alternative
Asymptotic velocity saturation α_{sat} (clean solar-neighborhood sample, $s \in [5,000, 30,000]$ AU)	Model-derived $\alpha_{\text{sat}}^{\text{flux}} = 4\sqrt{2}/k = 0.353$ ($\alpha_{\text{sat}}^{\text{prop}} = 0.371$ under propagator+vertex readout); frozen DR3 baseline $\alpha_{\text{sat}} = 0.366 \pm 0.012$	Newtonian GR: $\alpha_{\text{sat}} = 0.000$; unscreened MOND: $\alpha_{\text{sat}} \rightarrow \infty$ ($(s/r_M)^{1/2}$ rise); unit-term $1 + 2\xi$: excluded at $\sim 30\sigma$

Observable Channel	Frozen DR3 Baseline & Prospectively Specified TEP Prediction	Discriminating Null / Alternative
Pooled all-environment onset scale R_s^{pool} and embedded-propagator prediction	Model-derived $R_s = \sqrt{GM/g_t} \simeq 4.7$ kAU isolated ($\simeq 3.3$ kAU embedded Galactic propagator, Step 017) with flux-conserving $s^{4/3}$ recovery; frozen DR3 pooled baseline $R_s^{\text{pool}} = 2,646 \pm 182$ AU (stat; ± 609 AU total, $g_N \simeq 1.1 \times 10^{-9} \text{ m s}^{-2}$)	Thin-shell chameleon $1 - (R_s/s)e^{-m_{\text{bg}}(s-R_s)}$ (already disfavored by $\Delta\chi^2 = +613.0$ in DR3)
Environment-resolved stratification $R_s(Z)$ across $ Z \in [0, 1,000]$ pc	Frozen DR3 joint-fit baseline $R_s^{\text{midplane}} = 4,912 \pm 792$ AU ($ Z < 100$ pc) vs. $R_s^{\text{high-} Z } = 3,447 \pm 366$ AU ($ Z > 150$ pc; 5-bin full-range $8,670 \rightarrow 4,326$ AU); prospective DR4 scaling $R_s(Z) = R_s(0) [\rho_{\text{bar}}(Z)/\rho_{\text{bar}}(0)]^p$ with $p \approx 0.4 \pm 0.1$ ($p = 0.50$ anchored benchmark predicting $R_s(\text{high-} Z)/R_s(\text{midplane}) = 0.64$ vs. DR3 0.702 ± 0.074)	Pure acceleration-threshold MOND with external field $g_{\text{ext}}(R_{\text{gc}})$: R_s invariant to local vertical gas/stellar density $\rho_{\text{bar}}(Z)$ at fixed R_{gc}
Epoch-astrometry + radial-velocity 3D deprojection (DR4 inner-binary rejection)	Survival of $\tilde{v}(s) = 1 + 0.353(1 - e^{-s/R_s})$ and monotone $ Z $ ordering after direct epoch-astrometric removal of unresolved hierarchical triples ($P_{\text{in}} < 10^3$ d)	Triple-contamination hypothesis: $\tilde{v} - 1$ collapses toward zero once astrometric wobble removes inner companions

7. Conclusion

The Gaia DR3 wide-binary population provides strong empirical support for the Temporal Shear recovery predicted by TEP. From a high-purity sample of 341,315 systems, the analysis identifies a characteristic transition radius $R_s = 2,646 \pm 182$ AU (statistical) from the canonical phenomenological TEP recovery profile fit. A conservative uncertainty budget—including jackknife stability and model-choice freedom—broadens the total uncertainty to ± 609 AU while preserving the same few-thousand-AU onset scale.

The transition corresponds to an effective local binary density $\rho_{\text{eff}} \approx 9.2 \times 10^{-18} \text{ g/cm}^3$, far below the Temporal Topology saturation scale ($\rho_T \approx 20 \text{ g/cm}^3$; Smawfield 2025g). As discussed in Section 6.2, this gap is expected: the conformal scalar field is already heavily screened by the Galactic halo and baryonic disk. The observed transition marks the separation at which the source-charge suppression of the locally observable Temporal Shear weakens enough that the conformal-sector gradient becomes kinematically active above the Keplerian baseline.

An independent cross-check reinforces this picture. The canonical acceleration scale of the master-action closure, $g_t = cH_0/(2|\beta_A|) \simeq 3.4 \times 10^{-10} \text{ m/s}^2$, yields the derived radius law $R_s = \sqrt{GM/g_t} \simeq 4.7$ kAU at the sample mean mass on the isolated-vacuum background—a factor $\simeq 1.8$ above the fitted value, the physically expected direction for an ambient Galactic field displacing the transition inward, with the embedded-propagator evaluation (step 017) compressing the same estimator to $\simeq 3.3$ kAU, a factor $\simeq 1.26$ within the systematic band and the half-excess marker coincident with the observation (Section 2.2)—while the empirical SPARC-extracted anchor $g_{\text{TEP}} \approx 5 \times 10^{-10} \text{ m/s}^2$ (Smawfield 2025g) independently predicts a screening scale of the same order, and the agreement tightens further once the Galactic external field is included. The fitted screening profile is favored over a constant boost by $\Delta\chi^2 = 3,583$ (Table 4.1), and alternative smooth-transition fits preserve the finite onset scale, confirming that the transition is not an artifact of the canonical functional form. Direct comparison with MOND interpolating functions—including the angle-averaged External Field Effect with per-bin median masses—shows that even the most generous MOND+EFE variant ($k = 2$) is rejected by $\Delta\chi^2 = +540$ relative to TEP, failing simultaneously in transition shape, onset scale, and saturation amplitude.

The environmental test provides independent corroboration. The inferred ordering between midplane and high- $|Z|$ subsamples is sensitive to the mass-luminosity relation: a uniform solar-metallicity MLR overestimates halo-star masses and can bias the comparison, whereas the color-dependent correction (Section 3.2) now uses an independent spectroscopic-or-prior guardrail and a solar-track calibration that both recover the TEP-predicted pattern. With those corrections in place, the lower-density high- $|Z|$ population enters the anomalous regime at smaller radius than the higher-density midplane, both in the full sample ($R_s = 4,662 \pm 196$ AU versus $7,131 \pm 1,341$ AU) and in the solar-track control ($R_s = 4,145 \pm 276$ AU versus $6,856 \pm 920$ AU; permutation $p \leq 10^{-4}$ for the full sample and $p \leq 10^{-3}$ for the solar track). A finer five-bin $|Z|$ stratification confirms the density-dependent scaling at significance 0.021: within the empirical density-scaling benchmark, the inferred exponent $p \approx 0.4 \pm 0.1$ under free fits is consistent with the square-root scaling $p = 1/2$ and with the $\rho^{1/3}$ ambient channel independently returned by the noise-corrected inversion.

Supplemental controls strengthen that interpretation. The transition remains present after stratifying by distance, mass ratio, and primary mass, with each half-sample favoring a finite screening profile over a constant boost. A stricter radial-velocity subset also preserves the effect, though with broader uncertainty. Catalog-level scrambling tests fail to reproduce the observed screening preference in 10,000 realizations ($p = 1.0 \times 10^{-4}$), an injection-recovery test confirms that the pipeline recovers a known TEP signal with 100% detection rate while returning 0 false positives in 100 null mocks when no signal is injected, and phase-mixed Newtonian orbital forward models remain near $\tilde{v} \simeq 1$ at large separation rather than reproducing the observed plateau. These controls do not generate the signal, but they make it appreciably harder to attribute the result to a narrow demographic selection, a catalog artifact, or a pipeline bias.

The wide-binary anomaly therefore emerges not as a generic failure of dark matter or standard inertia, but as a structured, environmentally modulated screening transition—one whose morphology, onset scale, and environmental ordering are quantitatively consistent with the conformal scalar field of TEP and are not reproduced by the Newtonian orbital-projection or MOND/EFE parameterizations tested here. The present analysis establishes the wide-binary regime as a new, independent probe of TEP screening, complementing the rotation-curve and compact-object evidence presented elsewhere in the series.

References

- Banik, I., Pittordis, C., Sutherland, W., Famaey, B., Ibata, R., Mieske, S., & Zhao, H. 2024, MNRAS, 527, 4573. *Strong constraints on the gravitational law from Gaia DR3 wide binaries.*
- Bovy, J. 2015, ApJS, 216, 29. *galpy: A Python Library for Galactic Dynamics.*
- Burnham, K. P. & Anderson, D. R. 2002, *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*, 2nd ed. (New York: Springer).
- Burrage, C. & Sakstein, J. 2018, Living Reviews in Relativity, 21, 1. *Tests of Chameleon Gravity.*
- Chae, K.-H. 2023, ApJ, 952, 128. *Breakdown of the Newton–Einstein Standard Gravity at Low Acceleration in Internal Dynamics of Wide Binary Stars.*
- Chae, K.-H. 2024a, ApJ, 960, 114. *Robust Evidence for the Breakdown of Standard Gravity at Low Acceleration from Statistically Pure Binaries Free of Hidden Companions.*
- Chae, K.-H. 2024b, ApJ, 972, 186. *Measurements of the Low-acceleration Gravitational Anomaly from the Normalized Velocity Profile of Gaia Wide Binary Stars and Statistical Testing of Newtonian and Milgromian Theories.*
- El-Badry, K., Rix, H.-W., & Heintz, T. M. 2021, MNRAS, 506, 2269. *A million binaries from Gaia eDR3: sample selection and validation of Gaia parallax uncertainties.*
- Famaey, B. & Binney, J. 2005, MNRAS, 363, 603. *Modified Newtonian dynamics in the Milky Way.*
- Famaey, B. & McGaugh, S. S. 2012, Living Reviews in Relativity, 15, 10. *Modified Newtonian Dynamics (MOND): Observational Phenomenology and Relativistic Extensions.*
- Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, A&A, 674, A1. *Gaia Data Release 3. Summary of the content and survey properties.*
- Hernandez, X. 2023, MNRAS, 525, 1401. *Internal kinematics of Gaia DR3 wide binaries: anomalous behaviour in the low acceleration regime.*
- Hernandez, X., Vereteletskyi, V., Nasser, L., & Aguayo-Ortiz, A. 2024, MNRAS, 528, 4720. *Statistical analysis of the gravitational anomaly in Gaia wide binaries.*
- Khoury, J. & Weltman, A. 2004, Physical Review Letters, 93, 171104. *Chameleon Fields: Awaiting Surprises for Tests of Gravity in Space.*
- Llinares, C., Mota, D. F., & Winther, H. A. 2014, A&A, 562, A78. *ISIS: a new N-body cosmological code with scalar fields based on RAMSES.*
- McKee, C. F., Parravano, A., & Hollenbach, D. J. 2015, ApJ, 814, 13. *Stars, Gas, and Star Formation in the Solar Neighborhood.*
- Milgrom, M. 1983, ApJ, 270, 365. *A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis.*
- Milgrom, M. 1994, Annals of Physics, 229, 384. *Dynamics with a Nonstandard Inertia-Acceleration Relation: An Alternative to Dark Matter in Galactic Systems.*
- Milgrom, M. 2010, MNRAS, 403, 886. *Quasi-linear formulation of MOND.*
- Pecaut, M. J. & Mamajek, E. E. 2013, ApJS, 208, 9. *Intrinsic Colors, Temperatures, and Bolometric Corrections of Pre-main-sequence Stars.* (Online table updated 2022.)
- Pittordis, C., Sutherland, W., & Shepherd, P. 2025, Open Journal of Astrophysics, 8. *Wide Binaries from Gaia DR3: testing GR vs MOND with realistic triple modelling.* DOI: 10.33232/001c.142887.
- Ratra, B. & Peebles, P. J. E. 1988, Phys. Rev. D, 37, 3406. *Cosmological Consequences of a Spontaneously Broken Scalar Field.*
- Raghavan, D., McAlister, H. A., Henry, T. J., et al. 2010, ApJS, 190, 1. *A Survey of Stellar Families: Multiplicity of Solar-type Stars.*
- Smawfield, M. L. (2025). *Temporal Equivalence Principle: Dynamic Time & Emergent Light Speed.* Preprint v0.16 (Jakarta). Zenodo. DOI: 10.5281/zenodo.16921911 (Paper 0)
- Smawfield, M. L. (2025). *Global Time Echoes: Distance-Structured Correlations in GNSS Clocks.* Preprint v0.28 (Jaipur). Zenodo. DOI: 10.5281/zenodo.17127229 (Paper 1)
- Smawfield, M. L. (2025). *Global Time Echoes: 25-Year Analysis of CODE Precise Clock Products.* Preprint v0.21 (Cairo). Zenodo. DOI: 10.5281/zenodo.17517141 (Paper 2)
- Smawfield, M. L. (2025). *Global Time Echoes: Raw RINEX Consistency Test.* Preprint v0.9 (Kathmandu). Zenodo. DOI: 10.5281/zenodo.17860166 (Paper 3)

Smaufield, M. L. (2025). *Temporal-Spatial Coupling in Gravitational Lensing: A Reinterpretation of Dark Matter Observations*. Preprint v0.9 (Tortola). Zenodo. DOI: 10.5281/zenodo.17982540 (Paper 4)

Smaufield, M. L. (2025). *Global Time Echoes: Empirical Synthesis*. Preprint v0.8 (Singapore). Zenodo. DOI: 10.5281/zenodo.18004832 (Paper 5)

Smaufield, M. L. (2025). *Temporal Topology Saturation Scale: Cross-Scale Consistency of ρ_T* . Preprint v0.9 (New Delhi). Zenodo. DOI: 10.5281/zenodo.18064365 (Paper 6)

Smaufield, M. L. (2025). *The Soliton Wake: Exploring RBH-1 as a Temporal Topology Candidate*. Preprint v0.5 (Blantyre). Zenodo. DOI: 10.5281/zenodo.18059250 (Paper 7)

Smaufield, M. L. (2025). *Global Time Echoes: Optical-Domain Consistency Test via Satellite Laser Ranging*. Preprint v0.7 (Mombasa). Zenodo. DOI: 10.5281/zenodo.18064581 (Paper 8)

Smaufield, M. L. (2025). *What Do Precision Tests of General Relativity Actually Measure?*. Preprint v0.9 (Istanbul). Zenodo. DOI: 10.5281/zenodo.18109760 (Paper 9)

Smaufield, M. L. (2026). *Temporal Equivalence Principle: Suppressed Density Scaling in Globular Cluster Pulsars*. Preprint v0.10 (Caracas). Zenodo. DOI: 10.5281/zenodo.18165798 (Paper 10)

Smaufield, M. L. (2026). *The Cepheid Bias: Resolving the Hubble Tension*. Preprint v0.12 (Kingston upon Hull). Zenodo. DOI: 10.5281/zenodo.18209702 (Paper 11)

Smaufield, M. L. (2026). *Temporal Equivalence Principle: A Unified Resolution to the JWST High-Redshift Anomalies*. Preprint v0.8 (Kos). Zenodo. DOI: 10.5281/zenodo.19000827 (Paper 12)

Smaufield, M. L. (2026). *Temporal Equivalence Principle: Temporal Shear Recovery in Gaia DR3 Wide Binaries*. Preprint v0.7 (Kilifi). Zenodo. DOI: 10.5281/zenodo.19102061 (Paper 13 — this work)

Smaufield, M. L. (2026). *Temporal Equivalence Principle: A Proper-Time Domain Boundary in the Outer Solar System*. Preprint v0.2 (Addis Ababa). Zenodo. DOI: 10.5281/zenodo.22858190 (Paper 36)

Data Availability & Reproducibility

This work follows open-science practices. All results are fully reproducible from raw data using the documented pipeline. All numerical results, figures, and statistics are generated by deterministic Python scripts processing real observational data from the Gaia DR3 archive, with manuscript claims injected directly from pipeline outputs.

Repository & Code

GitHub Repository: github.com/matthewsmaufield/TEP-WB

The repository contains a deterministic, version-controlled analysis pipeline for wide binary screening using 341,315 Gaia DR3 systems.

Repository Structure

```
TEP-WB/
├── data/                                # Gaia DR3 catalogs and processed samples
│   ├── processed/                       # Cleaned binary samples
│   └── raw/                             # Original Gaia data (auto-downloaded)
├── logs/                                # Execution logs
├── results/                             # Analytical outputs and figures
├── scripts/
│   ├── steps/                           # Sequential analysis pipeline
│   │   ├── step_000_catalog_ingestion.py
│   │   ├── step_001_sample_selection.py
│   │   ├── step_002_kinematic_analysis.py
│   │   ├── step_003_screening_test.py
│   │   ├── step_004_sample_characterization.py
│   │   ├── step_005_environment_test.py
│   │   ├── step_006_audit_analysis.py
│   │   ├── step_007_supplemental_controls.py
│   │   ├── step_008_injection_recovery.py
│   │   ├── step_009_advanced_diagnostics.py
│   │   └── step_010_referee_hardening.py
```

```
| | | | step_011_mond_comparison.py
| | | | step_012_newtonian_forward_model.py
| | | | step_013_claim_consistency_audit.py
| | | | step_014_shell_consistency.py
| | | | step_015_derived_operator_forward.py
| | | | run_all_steps.py          # Master pipeline runner
| | | | utils/                  # Shared analysis utilities
| | | |
| | | | site/
| | | | | components/          # HTML manuscript source
| | | | requirements.txt       # Python dependencies
```

Data Provenance

Data Source	Provider	Access Method	Download Size	DOI/URL
Gaia DR3	ESA/Gaia	Public archive	~10 GB (full archive)	Gaia DR3
Wide Binary Catalog	El-Badry et al./Gaia	Public catalog	~500 MB (341,315 binaries)	Via Gaia archive

Total Download Size: ~500 MB for the wide binary catalog (Gaia DR3 subset).
Note: The analysis downloads data automatically via `astroquery` Gaia archive queries during Step 000 execution.

Reproduction Instructions

Quick Start (Full Reproduction)

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

# 2. Install dependencies
pip install -r requirements.txt
npm install --prefix site

# 3. Run complete pipeline
python scripts/steps/run_all_steps.py

# 4. Build manuscript
npm run build:markdown --prefix site
```

System Requirements

Component	Minimum	Recommended	Tested On
CPU	8 cores	14+ cores	Apple M4 Pro (14-core)
RAM	16 GB	32 GB	24 GB (M4 Pro)
Storage	10 GB	20 GB	NVMe SSD
Runtime	~2-4 hours (full 14-step pipeline)		~3 hours (M4 Pro)

Detailed Pipeline Steps

The analysis pipeline consists of 13 deterministic steps. Each step produces JSON/CSV outputs and logs for full traceability:

Step 000: Catalog Ingestion

- step_000_catalog_ingestion.py** — Download the El-Badry et al. (2021) Gaia eDR3 wide-binary catalog from Zenodo, convert the FITS table to a pandas DataFrame, and compute foundational geometric quantities (projected separation, system distance, Galactocentric coordinates) plus a placeholder dimensionless velocity ratio.

Step 001: Sample Selection

- **step_001_sample_selection.py** — Apply the high-purity cuts: chance-alignment probability $< 1\%$, parallax SNR > 20 and proper-motion SNR > 10 for both components, RUWE < 1.2 to suppress unresolved hierarchical triples, and projected separation $50 < s < 50,000$ AU. Output: a clean parquet sample of 341,315 systems.

Step 002: Kinematic Analysis

- **step_002_kinematic_analysis.py** — Compute the projected relative tangential velocity from proper-motion differences, the Newtonian circular velocity $v_c = \sqrt{GM_{\text{tot}}/s}$, and the dimensionless ratio $\tilde{v} = v_{\text{tan}}/v_c$. Stellar masses use the Pecaut & Mamajek (2013, updated 2022) main-sequence relations together with a color-dependent metallicity mass-bias correction whose β_{MLR} calibration is restricted to independent spectroscopic metallicities when available and otherwise falls back to a conservative external prior. Coordinates are also transformed to Galactocentric (R, Z) for environmental stratification.

Step 003: Screening Test

- **step_003_screening_test.py** — Bin systems in $\log s$, take the median $\tilde{v}$ in each bin with parametric-resampling 68% intervals, and fit phenomenological screening models (TEP exponential, sigmoid, thin-shell, double-exponential, flat Newtonian, constant boost). Performs a Wald–Wolfowitz runs test, computes residual lag-1 autocorrelation and Durbin–Watson, executes covariance-aware refits (AR(1) GLS and a Gaussian Process in $\log_{10} s$), and runs jackknife bin-drop stability.

Step 004: Sample Characterization

- **step_004_sample_characterization.py** — Generate diagnostic plots characterizing the cleaned sample: distance, separation, and stellar-mass distributions, plus a sky map. Used as a sanity check that the high-purity sample lies in the local solar neighborhood and is adequately populated across separation bins.

Step 005: Environmental Test

- **step_005_environment_test.py** — Stratify the sample by Galactocentric height $|Z|$ into a midplane subsample ($|Z| < 100$ pc) and a high- $|Z|$ subsample ($|Z| > 150$ pc), fit the screening model with α_{sat} fixed to remove the amplitude–scale degeneracy, and compare the resulting R_s values. Includes a solar-track control (zero metallicity correction), a joint profile-likelihood fit with shared α_{sat} , permutation tests, and an empirical density-scaling benchmark that infers the exponent p from the observed R_s ratio (the scaling is treated here as phenomenology, not as a claimed microscopic completion; Smawfield 2025a, §A4).

Step 006: Audit Analysis

- **step_006_audit_analysis.py** — Compute the implied effective density at the fitted screening radius from the median total binary mass, and verify population demographics by comparing midplane and high- $|Z|$ color–magnitude diagrams to confirm that the high- $|Z|$ subsample is genuinely bluer/metal-poor (validating the metallicity mass-bias correction in step_002).

Step 007: Supplemental Controls

- **step_007_supplemental_controls.py** — Robustness checks on demographic subsets (RV-consistent subsample, distance halves, mass-ratio halves, primary-mass halves) and direct null controls via permutation scrambling of $\tilde{v}$ both globally and within distance quartiles. Reports effect sizes and permutation p-values for each control.

Step 008: Injection / Recovery

- **step_008_injection_recovery.py** — Build mock catalogs from the projected Newtonian forward-model null of step_012, inject a known TEP screening signal, and rerun the full analysis to measure recovery fidelity (bias and detection rate). Includes a null-injection ($\alpha = 0$) false-positive check, an R_s sweep, and an eccentricity-distribution sensitivity test.

Step 009: Advanced Diagnostics

- **step_009_advanced_diagnostics.py** — Three diagnostics: (i) a triple-contamination forward model that tests whether residual unresolved hierarchies can reproduce the observed outer $\tilde{v}$; (ii) a quantitative self-screening fit $\alpha_{\text{sat}}(M) = \alpha_0 \exp(-M/M_{\text{screen}})$ across the demographic mass splits; (iii) a fine $|Z|$ stratification into 5 equal-count height bins, refining the inferred density-scaling exponent p .

Step 010: Referee Hardening

- **step_010_referee_hardening.py** — Three referee-requested checks: (i) a normalization-window sensitivity matrix that sweeps the inner-baseline definition; (ii) a spatial-substructure cross-match of high-residual bins against sky position; (iii) a Pittordis-faithful triple forward model using the exact triple-population distributions of Pittordis et al. (2025).

Step 011: MOND Comparison

- **step_011_mond_comparison.py** — Compare the binned $\tilde{v}$ profile to MOND ν -functions (simple and standard) with per-bin median masses, both with and without the angle-averaged External Field Effect (QUMOND prescription). Reports χ^2 , AIC and $\Delta\chi^2$ relative to the TEP exponential fit. Also computes the a priori TEP R_s prediction from the SPARC-derived g_{TEP} .

Step 012: Newtonian Forward Model

- **step_012_newtonian_forward_model.py** — Construct projected Newtonian orbital null catalogs conditioned on the observed projected-separation and mass distributions, marginalized over line-of-sight geometry, orbital phase, and eccentricity distributions. Reports the same binned normalized profile used by the main screening test and quantifies whether a purely Newtonian orbital population can reproduce the observed outer plateau.

Step 013: Claim Consistency Audit

- **step_013_claim_consistency_audit.py** — Cross-check manuscript source and generated markdown against the latest pipeline outputs, including non-circular metallicity calibration, MOND comparison values, transition-model statistics, Newtonian forward-model results, and injection-recovery false-positive rates.

Step 014: Screening-Shell Consistency Map

- **step_014_shell_consistency.py** — Evaluate the resolved nested two-body screening operator $\mathcal{R}(s) = \mathcal{S}_\Sigma(X_{\text{env}})^2 y(s)$, where y is the exact flux-conserving profile $y[1 + y^2(r_*/r)^4] = 1$, $r_* = \sqrt{GM/g_t}$ is the derived shell radius with $g_t = cH_0/(2|\beta_A|)$, and X_{env} is the embedding ambient of the hierarchy level above the pair — the dominant member's field for hierarchical systems (giving $\mathcal{R} \rightarrow y^3 \sim s^4$ deep inside the shell) and the Galactic field for comparable-mass binaries (giving the $\sim s^{4/3}$ transition) — across the separation ladder (Cassini conjunction, 1 AU, Saturn, Earth–Moon, wide-binary transition, Oort cloud) and the environmental operator $\mathcal{S}_\Sigma(\mathcal{E}) = [1 + (g/g_t)^2]^{-1}$ at benchmark field strengths. Verifies that all Solar-System constraints sit deep inside the solar screening shell while the wide-binary companion samples the shell boundary, and computes the amplitude ledger relating the observed α_{sat} plateau to the ambient-floor-limited partial recovery.
- **step_015_derived_operator_forward.py** — Forward-model test of the derived operator against the real 341,315-binary catalog: each binary's projected separation is deprojected, the per-binary shell radius follows the derived $r_* = \sqrt{GM/g_t}$ law with lognormal scatter, the Galactic ambient floor carries environmental scatter, and unresolved-triple contamination uses the corrected vector-addition kinematics of steps 009/010. The binned median profile is compared against the 19-bin measurement across a family of candidate transition profiles — the power-law $[1 + (R_s/r)^k]^{-1}$ family over $k = 0.8 \dots 4$ and the exact flux-conserving profile $y(s)$ of the resolved operator — yielding χ_{red}^2 for each: the data select transition-region slopes near the derived $4/3$ (the exact y profile reaches $\chi_{\text{red}}^2 \approx 8.4$, versus 18.3 for the steep $k = 4$ form) and an ambient floor $S_{\text{env}} \approx 0.45$ consistent with the squared Galactic vertex factor $q_{\text{GAL}}^2 \approx 0.43$.

Each step produces JSON/CSV outputs with full metadata in `results/outputs/`, and execution logs are written to `logs/` with timestamps for complete traceability. Run all steps via: `python scripts/steps/run_all_steps.py`

Software Versions

- **Python** 3.11+ (tested with 3.13)
- **NumPy** 1.24+
- **SciPy** 1.10+
- **Pandas** 2.0+
- **Matplotlib** 3.7+
- **Astroquery** (Gaia data access)