

A Gaia search for boost-holonomy relics: the full-sky co-spatial radial-velocity pair census

Jonathan Schemoul

LibertAI, France

js@libertai.io

October 27, 2026

Abstract

A time machine’s unmanufacturable organ is a boost-holonomy identification: a locally flat region whose covering space is quotiented by a Lorentz boost. If such a relic exists, it leaves an exact, massless, frequency-multiplicative signature. A photon crossing the seam once emerges with its frequency multiplied by $r = \cosh \Lambda - n_z \sinh \Lambda$, exactly $e^{\pm \Lambda}$ for axial photons, so one source appears as a co-spatial spectral double image at a fixed, source-independent ratio, with no lens mass and no magnification. For the small rapidities accessible to stellar radial velocities, $\Lambda \sim 10^{-4}$ – 10^{-3} , the wound image is a second spectrum at the same sky position with a radial-velocity offset $\Delta RV = c(e^{\pm \Lambda} - 1)$. We execute the resulting full-sky census on Gaia DR3: a TAP self-join of the 33,812,183 sources with published radial velocities, tiled by `source_id` into 768 level-3 HEALPix tiles with level-4/5 splits, producing 1,032 pair CSVs and 16,358 unordered raw pairs at separations $\leq 1''$. An astrometry match (parallax and proper motion consistent at 3σ) labels 7,217 pairs (44.1%); 1,177 of these have $|\Delta RV| > 3\sigma$, and 939 satisfy the candidate-zone cut $|\Delta RV| > 30 \text{ km s}^{-1}$ and $> 3\sigma$. A completeness audit changes the interpretation of these counts: 6,040 of the 7,217 labelled pairs (83.7%) have at least one missing astrometric component, because the cut treats a missing value as a zero difference. Requiring all six astrometric values leaves 1,177 matched pairs, 171 above 3σ , and 135 in the candidate zone. The permissive and strict 1,177 counts are different samples; their equality is a coincidence of the data. The matched-pair ΔRV distribution (permissive sample) has median 7.4, 90th percentile 41.6, 99th percentile 126.8, and maximum 489.5 km s^{-1} ; the strict distribution is median 7.1, 90% 38.4, 99% 130.3, maximum 472.8 km s^{-1} . The corresponding axial- Λ window is 1.0×10^{-4} to 1.6×10^{-3} , complementary to the gravitational-wave channel’s $0.1 < |\Lambda| < 3$. The shared-base family test returns a largest value of 173 of 939 permissive candidates (24 of 135 strict candidates) at the low- Λ threshold ($\Lambda = 1.29 \times 10^{-4}$, $\Delta RV = 38.7 \text{ km s}^{-1}$). This is not a cluster: the tolerance is a half-width $0.1b$, i.e. a full 7.74 km s^{-1} window at the maximum, the local density is 22.4 candidates per km s^{-1} , and the smooth expectation is 173.4 against 173 observed (strict: 24.0 against 24). We calibrate the statistic with a composition-matched null that holds the candidate count and sky positions fixed and resamples the $(\Lambda, \sigma_\Lambda)$ pairs with replacement (20,000 draws): the observed maximum is below the null median in both samples ($p = 0.74$; 90% range 164–197; strict $p = 0.97$, 24–34), so no excess survives. Smooth Gamma and log-KDE composition models give $p = 0.18$ and 0.12 (strict 0.81 and 0.76), the exact assignment permutation is invariant ($p = 1.00$), and a position-fixed pool draw gives $p = 1.00$; the null is calibrated on real non-candidate controls ($p = 0.63$ – 0.86). The test has little power: a sharp synthetic family needs about 40 edges at the observed low- Λ base (strict 20) or 200 (strict 40) at high Λ for 95% power, and about 200 edges (strict 100) when the family has the measured 20% radial-velocity error spread. A realistic few-pair family is invisible. We report no significant signal, no detection, and no exclusion. The candidate set is consistent with

an astrophysical background: co-moving binaries and multiples, line-of-sight superpositions, and composite spectra. The epoch test is implemented, but only its epoch-constancy proxy is weakly validated (1/300 trials at $p < 0.01$, with a wide binomial confidence interval); the epoch-spectrum/exact-copy test has never been run on a candidate pair. Gaia DR3 exposes mean RVS spectra and no epoch RVS spectra: 9 of the 1,132 candidate sources in the EM-2 snapshot (524 tiles, 9,415 pairs; not the final census) have a mean RVS, 0 have published epoch RVs, and 0 candidate pairs have both mean spectra. The full epoch-series test is not executable on DR3; the weak summary-statistic proxy is available for 112 candidate sources, of which 51 are RV-variable. The 1,898-source epoch sample has zero overlap with the pair sample: 0 of the 18,830 snapshot pair sources have epoch RVs. The next steps are targeted multi-epoch spectroscopy of the nine pairs with one RVS spectrum and a future Gaia epoch-spectrum release. We quote no EM density limit: the family statistic is calibrated against the background composition, but the census measures neither the selection efficiency nor the pocket multiplicity. For context only, the GW companion’s calibrated background analysis, on the corrected M_c/D_L scaling cut, gives $\mu < 274$ (206 model-dependent; the Gamma–Poisson model is not adequate at $C > 0.1$) at its operational $C > 0.1$ threshold, or $\mu < 4.50$ (4.95) at $C > 0.9$, for the GW channel; this is not an EM result. The survey is sensitive only to pc-scale star-populated pockets (one star in the entry shell needs a pocket radius ~ 2 pc; 173 pairs need ~ 10 pc), not to the canonical metre-scale relic, which is causally cloistered.

Keywords: relativity — time machines — surveys — methods: data analysis — stars: kinematics and dynamics — gravitational lensing

1 Introduction

1.1 Chronology protection and the unmanufacturable organ

Hawking’s chronology protection conjecture proposes that the laws of physics prevent the creation of closed timelike curves [1]. The proposed mechanism is semiclassical: on a chronology horizon, null geodesics winding around a closed null geodesic pile up, virtual particles acquire a boost holonomy at each circuit, and the renormalized stress-energy tensor diverges. Explicit computations supported the picture on Taub–NUT space, wormhole time machines, and generic compactly generated horizons, and the Kay–Radzikowski–Wald theorem put the compactly generated case on a rigorous footing [2, 3]. The theorem is nevertheless a statement about distributions on a fixed background, not about back-reaction, and it applies only to compactly generated horizons. Non-compactly generated machines—rotating cylinders, Krasnikov tubes, and the Ori–Soen family—are untouched by it [4].

The Ori program is the sharpest constructive challenge. Ori’s 2005 solution is a class of curved, asymptotically flat, topologically trivial vacuum spacetimes that develop closed timelike curves inside an empty toroidal region bounded by a non-compactly generated chronology horizon, with the surrounding matter satisfying the weak, dominant, and strong energy conditions [5]. The 2007 construction completes the initial data with a dust envelope [6], and the Xavier traveling-wave family gives a second explicit core [7]; the geometric and causal classification of these spacetimes has since been sharpened [10, 11]. What none of these constructions supplies is a classical manufacture of the identification itself: the boost holonomy is the machine’s unmanufacturable organ. One can write the metric, but one cannot build the quotient by assembling ordinary matter. Modified-gravity embeddings re-solve the same core without removing its chronology-violating regions [12]; they change the theory in which the core exists, not the fact that its holonomy is an identification rather than a field.

That obstruction is also an opportunity. If the identification can be carried by a locally flat

region—the $f = 0$ limit of the Ori core, or the $f \approx 0$ limit of the Xavier core—then a *bare relic* is a flat tube with a boost identification and nothing else. Such an object has no Newtonian mass, no local curvature, and no ordinary lens phenomenology, but it has an exact effect on any radiation that crosses its seam. The effect is multiplicative and source-independent: a frequency is multiplied by a fixed factor, and a whole spectrum is copied at that factor. This is the observational signature we search for. The question changes from “can the holonomy be manufactured?” to “is a relic already there?”.

The survey program has two channels. The gravitational-wave (GW) channel exploits the fact that a GW’s frequency and its inferred source-frame masses are proportional, so the holonomy becomes a common scaling of both masses, with the mass ratio and spins invariant, the same luminosity distance, and no magnification. The GW catalog search is complete and finds no significant signal at any stage; its decisive methodological lesson, measured in the companion paper [15], is that the shared-base family test is degenerate at current parameter-estimation precision. The GWTC-4.0 lensing and GWTC-5.0 methods papers supply the catalog and lensing context [16, 17]. The electromagnetic (EM) channel, reported here, targets the same identification with photons and stellar radial velocities. The two channels are independent hardware and independent systematics. A coincident candidate would be much harder to explain than either alone; neither channel has one.

The closest existing analyses are strong-lensing searches for multiple images, which test pair consistency in time delay, magnification, and mass ratio and, in some cases, use wave-optics interference [18, 19, 16]. The holonomy pair is different in kind: it is frequency-shifted and unmagnified. Waveform systematics and precessional modulation can mimic lensing support [20, 21], and the host galaxy affects microlensing searches [22]; these backgrounds do not produce a shared base across independent pairs, so the family test is their falsifier in principle. The state-choice loophole [9] is the other classical escape: a tuned vacuum can render the stress finite at a chronology horizon, but it does not manufacture the identification either. The survey therefore targets the identification itself, not a particular quantum state.

1.2 Why Gaia

Gaia is the only existing dataset in which the EM pair cut can be applied at full-sky scale. Its DR3 catalogue contains 33,812,183 sources with a published radial velocity [13], and the RVS spectrograph supplies both the radial velocities and, for a subset, mean spectra. The signature is chiefly spectroscopic: the wound image is a copy of the source spectrum at a fixed frequency ratio, and the observable is the radial-velocity difference between two spectra at the same position. Gaia’s all-sky astrometry supplies the co-spatial and co-moving cut: parallax and proper motion must agree at 3σ . A line-of-sight superposition does not normally satisfy that consistency test, although it can at low significance or when a component is missing. The same archive supplies the mean-spectrum products that test whether two spectra are copies rather than two different stars.

The EM channel probes a different rapidity range from the GW channel. For a radial-velocity difference ΔRV , the axial holonomy is $\Lambda_{\text{axial}} = \ln(1 + |\Delta RV|/c)$; with RVS velocities from tens to hundreds of km s^{-1} , the census is sensitive to $\Lambda \sim 10^{-4}$ – 10^{-3} , two to three orders of magnitude below the GW window $0.1 < |\Lambda| < 3$. The two channels are therefore complementary, not redundant. The price is that the EM signature is small: the wound image’s velocity offset is $\sim c\Lambda$, so the search must resolve a velocity difference at the same position, not a spatial double image.

1.3 What this paper reports

We report the complete EM-1 census and its science cut. The census is a full-sky TAP self-join of Gaia DR3 RV sources at $\leq 1''$, tiled by `source_id`; the cut applies a co-spatial astrometry match, a radial-velocity discrepancy cut, the axial- Λ map, and the shared-base family test. We find no significant signal. The census produces 16,358 raw pairs, of which the operational cut labels 7,217 astrometry-matched, 1,177 above 3σ in ΔRV , and 939 in the candidate zone. A completeness audit shows that most of the permissive matches lack one or more astrometric components; the strict complete-astrometry sample is 1,177 matched, 171 above 3σ , and 135 in the candidate zone. The family test returns a maximum of 173 of 939 permissive candidates (24 of 135 strict), but this is the smooth candidate density seen through a half-width $0.1b$ window, not a cluster or a shared base: the full window at the maximum is 7.74 km s^{-1} , the local density is 22.4 candidates per km s^{-1} , and the expectation is 173.4 against the 173 observed. A composition-matched null, which holds the candidate count and positions fixed and resamples the $(\Lambda, \sigma_\Lambda)$ pairs with replacement, puts the observed maximum below its median in both samples ($p = 0.74$; strict $p = 0.97$), so no excess survives. The statistic has little power: a sharp family needs about 40 edges at the observed base (strict 20) and about 200 at high Λ for 95% power. We claim no significant signal, no detection, and no exclusion. The candidate set is expected to be dominated by the astrophysical background that a spectral follow-up must vet.

The decisive test is the epoch-RVS test: a holonomy shift is constant across epochs and exactly multiplicative across spectral features, while a binary’s velocity difference varies with orbital phase. The test is implemented, but only its epoch-constancy proxy is weakly validated (1/300 trials at $p < 0.01$, with a wide confidence interval); the epoch-spectrum/exact-copy test has never been run on a candidate pair. Gaia DR3 exposes mean RVS spectra and no epoch RVS spectra. In the EM-2 snapshot (524 tiles, 9,415 pairs; not the final census), 9 of 1,132 candidate sources have a mean RVS, 0 have published epoch RVs, and 0 candidate pairs have both mean spectra. The full epoch-series test is therefore not executable on DR3 candidate data; the weak summary-statistic proxy is available for 112 candidate sources. The paper states that limitation plainly, reports the census and the absence of a significant signal, and lays out the targeted spectroscopy and the future epoch-spectrum release that would make the test decisive.

The paper is organized as follows. Section 2 derives the signature and defines the EM observable. Section 3 describes the data and the census. Section 4 defines the cuts and the family statistic. Section 5 presents the results. Section 6 discusses the astrophysical background. Section 7 reports the EM-2 epoch-spectrum test and the DR3 data limitation. Section 8 states the limits honestly. Section 9 lays out the survey strategy, and Section 10 concludes. Appendix A is the exact-number ledger, Appendix B the reproduction commands, Appendix C the boost algebra, and Appendix D the family statistic.

2 The boost-holonomy signature

2.1 The relic as a flat boost quotient

The Ori 2005 core is a vacuum plane-wave geometry with a compact z -direction. In the $f = 0$ limit the (T, z) block has Riemann tensor identically zero; the core “is just the Misner space” [5, 8]. The same limit occurs in the Xavier core, where Riemann is zero exactly [7]. The covering is flat, and the identification $(T, z) \rightarrow (T, z + L)$ is a Lorentz boost in the Ori metric convention. Writing the flat (T, z) block in Milne coordinates, $\tau = 2\sqrt{-T}$ and $\psi = z/2 + \frac{1}{2} \ln(-2T)$, brings it to $-dW^2 + dZ^2$;

the deck transformation $z \rightarrow z + L$ is then the boost $Z \rightarrow Z + L/2$, so

$$\Lambda = \frac{L}{2} \quad (1)$$

with z equal to twice the covering boost angle up to a logarithm. An exact symplectic covering, a photon-transport check, and the Misner gate $\sigma_n = 4|T_0|(\cosh(nL/2) - 1)$ confirm this dictionary to 10^{-14} ; we do not claim a general uniqueness proof. The spectral cuts below use $\Lambda = L/2$.

A bare relic is this identification carried by a locally flat tube $\rho < \rho_0$, with no surrounding machine. It has no local curvature and no Newtonian mass. Its global holonomy is not a field that can be screened or shielded: a path that crosses the seam once returns with a Lorentz-transformed momentum. The wall matching to the exterior is supplied by the wrap in the full machine; in the relic survey we work in the flat ($a \rightarrow 0$, $f \approx 0$) limit and treat the seam as sharp. The canonical design used for the numerical examples in the companion paper is $\rho_0 = 0.5$ m, $e = 0.1$, $a = 10^{-4}$, $\tau_0 = -1$ m², $L = 1$, hence $\Lambda = 0.5$. The EM census does not assume this canonical design; it scans the rapidity range that stellar radial velocities can resolve.

2.2 Photons: the co-spatial spectral double image

Consider a photon that crosses the seam once. Its four-momentum is acted on by the inverse boost. For a boost along z with rapidity Λ and a photon direction cosine n_z along the boost axis, the frequency ratio between the wound and direct images is

$$r \equiv \frac{\nu_{\text{wound}}}{\nu_{\text{direct}}} = \cosh \Lambda - n_z \sinh \Lambda. \quad (2)$$

Because $n_z \in [-1, 1]$, the ratio lies in the interval $(e^{-\Lambda}, e^{+\Lambda})$. For photons running along the boost axis the boost does not redirect the light and the ratio is exact:

$$r(n_z = +1) = e^{-\Lambda}, \quad r(n_z = -1) = e^{+\Lambda}. \quad (3)$$

The source therefore shows a *co-spatial spectral double image*: one position on the sky, two copies of one spectrum at a fixed ratio $e^{\pm\Lambda}$. There is no mass, no deflection, and no source dependence; for the axial channel the same Λ applies to every pair from the same pocket. For $\Lambda = L/2 = 0.5$ ($L = 1$) the ratios at $n_z = -1, -\frac{1}{2}, 0, +\frac{1}{2}, +1$ are

$$r = 1.649, 1.388, 1.128, 0.867, 0.607, \quad (4)$$

so the wound copy is blue ($e^{+\Lambda}$) for one winding and red ($e^{-\Lambda}$) for the other. The pair always contains the direct image and one wound image; the ratio is r , not r/r^{-1} .

Off-axis photons are additionally tilted. For a photon at angle θ from the boost axis the aberration changes the angle by approximately $\Lambda\theta$ for small Λ , and the maximum tilt (at $\theta = \pi/2$) is $\arctan(\sinh \Lambda)$: 27.5° for $\Lambda = 0.5$, but only 0.09° for $\Lambda = 1.6 \times 10^{-3}$. On-axis photons are not tilted at all. The observable is therefore chiefly spectroscopic, not astrometric, and the 1-arcsec co-spatial cut used below is itself an axial-channel selector: at the small Λ probed here, the wound image stays within $1''$ of the direct image only for $\theta \lesssim 1''/\Lambda \approx 0.3^\circ$ at $\Lambda = 10^{-3}$. The census is an axial survey. Off-axis pairs would appear at larger separations and are not in this sample.

A multi-wound exit gives ratios $e^{\pm n\Lambda}$: integer powers of one fixed base, still a fixed-family signature. The in-tube Misner-phase sweep sets the winding number, $\Delta\psi \approx 2w(1 + n_z)/(n_x\tau_c)$ and $n \approx 2w/(\Lambda\tau_c)$ per unit axial alignment, where τ_c is the pocket clock at the crossing and w the

chord half-width. There is also a time echo: the wound copy's extra light path is of order the boost displacement

$$\Delta t_n \simeq 2\tau_{\text{wall}} \sinh\left(\frac{n\Lambda}{2}\right). \quad (5)$$

For the canonical metre pocket the wall Rindler radius is $r = \sqrt{e\rho_0^2 - t_0} = 1.01$ m, the Milne covering radius is $\tau_{\text{wall}} = 2r = 2.02$ m, the circuit length is $\Lambda\tau_{\text{wall}} = 1.01$ m $= L\sqrt{e\rho_0^2 - t_0}$, and the echo is 3.4 ns per circuit. For a large pocket, τ_{wall} can be astronomical and the echo can be days. The EM census does not measure the echo; it measures the frequency ratio, which is the cleaner observable for stellar spectra.

2.3 The radial-velocity observable and the small- Λ window

A multiplicative frequency ratio shifts every spectral line by the same factor. For a line at frequency ν , the observed shift is $\Delta\nu/\nu = r - 1$; for small Λ the apparent radial-velocity difference between the direct and wound images is

$$\Delta\text{RV} = c(r - 1) \simeq c\Lambda, \quad (6)$$

with the sign set by the winding direction. Inverting the exact relation gives the axial rapidity

$$\Lambda_{\text{axial}} = \ln\left(1 + \frac{|\Delta\text{RV}|}{c}\right). \quad (7)$$

The direct image is the unshifted reference; the wound image is the one shifted by $e^{\pm\Lambda}$. The pair's velocity difference is therefore $c(e^{\pm\Lambda} - 1)$, not the larger $c(e^{+\Lambda} - e^{-\Lambda})$ that would apply if both images were wound in opposite directions. This convention is used throughout.

The census's candidate threshold is $|\Delta\text{RV}| > 30 \text{ km s}^{-1}$, and the largest matched-pair difference is 489.5 km s^{-1} . The realized axial window is therefore

$$\Lambda_{\text{axial}} \in [1.0 \times 10^{-4}, 1.6 \times 10^{-3}], \quad (8)$$

with the lower bound set by the 30 km s^{-1} threshold and the upper bound by the largest observed difference. The design sensitivity is wider: RVS can measure velocities up to of order 10^3 km s^{-1} , so the census is in principle sensitive to $\Lambda \lesssim 3 \times 10^{-3}$. No pair above 489.5 km s^{-1} survived the cuts. The window is two to three orders of magnitude below the GW channel's $0.1 < |\Lambda| < 3$, and there is an uncovered gap between $\sim 3 \times 10^{-3}$ and 0.1 that neither channel probes today. The EM channel is therefore a small- Λ companion, not a replacement.

The radial-velocity resolution is the enabling fact. RVS velocities are accurate at the few km s^{-1} level for bright stars, so a 30 km s^{-1} difference is many sigma; the threshold is set by the astrophysical background, not by the noise. At the low end of the window, a 10^{-4} rapidity is a 30 km s^{-1} shift; at the high end, a 1.6×10^{-3} rapidity is a 490 km s^{-1} shift. Both are within the RVS range, and both are far below the relativistic regime where the small-angle expansion would fail. The expansion error is $O(\Lambda^2) \sim 10^{-6}$ at the top of the window, negligible compared with the velocity errors.

2.4 What the signature is not: no lens, no mass

The relic has no Newtonian mass. The direct image's deflection, integrated through the exact metric, is independent of the coupling a across the whole design window $a \in (2 \times 10^{-17}, 0.1)$; for a small-drift test photon it is about 3×10^{-5} deg (`simulations/relic_signature.py`, Sec. 5), and

Pocket ρ_0	Shell stars
0.5 m	3.8×10^{-51}
1 AU	1.6×10^{-17}
1000 AU	1.6×10^{-8}

Table 1: Stars in the entry shell of a bare relic. The count is $0.004(\rho_0/\text{ly})^3$ for the shell volume $V \simeq \pi\rho_0^2\tau_{\text{wall}} \simeq \rho_0^3$ and a density of 0.004 stars/ly³; it is reproduced by `simulations/relic_signature.py`, Sec. 2. The cubic scaling means that a star-populated pocket needs $\rho_0 \sim 2$ pc for one star and ~ 10 pc for 173 pairs, not ~ 1000 AU. The OFF relic is invisible to a generic stellar survey; the channel opens only for pc-scale pockets or at the trigger epoch.

the e -coupling $\Gamma_{zz}^x = -ex$ shows that the effect is cone geometry, not wave and not mass. Any mass fitted to the pair’s astrometry is fake. A 27.5° separation at 1 kpc would require a lens mass of order 10^{45} kg (the working value 5×10^{45} kg = $2.4 \times 10^{15} M_\odot$ is a conservative order of magnitude; the geometric estimate $\theta_E^2 c^2 D/(4G)$ at $\theta_E \simeq 0.25$ rad gives 6×10^{44} kg). The discriminator that no mass lens can fake is the fixed frequency ratio $e^{\pm\Lambda}$, and its integer-power ladder $e^{n\Lambda}$.

A conventional lens produces multiple images with the same frequency (the source redshift is the same), a magnification ratio, and a time-delay geometry tied to the lens mass. The holonomy pair has a frequency ratio, no magnification, and a delay set by the pocket’s wall radius. A frequency-shifted pair with no magnification anomaly therefore does not match any conventional lens model. In the Gaia census the magnification and the time delay are not measured for stellar pairs; the operational stand-in is the co-spatial astrometry match (same parallax and proper motion), which a lensed pair of one source would also satisfy. The astrometry cut alone is therefore not a lens-versus-holonomy classifier. It is a background-selection stage; the fixed frequency ratio, measured spectroscopically, is the classifier. The census’s job is to produce the sample on which that classifier can be run.

2.5 Causal cloistering and the pc-scale window

Why is the survey not dominated by nearby relics? The pre-chronology-horizon pocket world-tube is the Rindler band $0 < \tau < \tau_{\text{wall}}$. A null chord from a generic background star crosses the tube’s spatial location at covering radius $\tau \sim \tau_{\text{now}}$ (the covering age), far outside the band, and never enters. Entry requires the chord to graze the birth-cone null surfaces $W = \pm Z$ inside the tube, a shell of thickness $\sim r_{\text{wall}}$ (the Milne radius $2r_{\text{wall}} = \tau_{\text{wall}}$ is the same order). With a volume $V \sim \pi\rho_0^2\tau_{\text{wall}}$ in the disk shadow, $\tau_{\text{wall}} = 2\sqrt{e\rho_0^2 + |t_0|} \simeq 0.316\rho_0$ for large ρ_0 , so $V \simeq \rho_0^3$, and a stellar density of 0.004 stars/ly³, the count is $0.004(\rho_0/\text{ly})^3$: Table 1.

A pre-horizon bare relic is thus causally cloistered: it hides its holonomy until the chronology horizon opens it. This is the covering-space form of the program’s topological-censorship statement. There are two windows in which light does enter. The first is *large pockets*, with τ_{wall} astronomically large, where the pair channel opens with the entry-shell cross-section. The second is *the trigger*: at $t > e\rho^2$ the z -circles turn timelike, winding is free of a phase budget, and the classical Misner pile-up opens—many images of one source at geometrically compressed frequencies crowding the horizon direction.

The event-rate accounting is sobering. The disk cross-section $\pi\rho_0^2$ at 1 kpc subtends 8×10^{-40} sr, so a metre-scale pocket covers about 10^{-31} Gaia-like stars even ignoring the entry shell; metre-scale relics are individually undetectable by stellar surveys. The channel becomes statistical only for pc-scale pockets: the entry-shell count is cubic in ρ_0 , so one star requires $\rho_0 \sim 2$ pc and 173 pairs

require $\rho_0 \sim 10$ pc, not $\gtrsim 1000$ AU. Even the wrapped 10^{22} – 10^{26} kg dust at 1 kpc has an Einstein angle of only 1.4×10^{-4} – 1.4×10^{-2} mas, far below the resolution of current optical astrometry. The wrapped machine is found by its pair, not by its lens. The Gaia census is therefore a *pc-scale-pocket* search. The family statistic is now calibrated against the background composition (Section 4.5), but the census still measures neither the selection efficiency nor the pocket multiplicity, so it does not bound the pc-scale population; it selects the sample on which a calibrated family or spectral test can be run. We return to this in Section 8.

2.6 False positives and the three discriminators

The EM pair channel has three known false-positive classes.

Real binaries and co-moving multiples. A bound or co-moving pair has consistent parallax and proper motion by construction, and its two components can have a large radial-velocity difference. A wide binary at 10^3 AU has an orbital velocity of order 1 km s^{-1} , so the high- ΔRV tail is not a population of wide binaries; it is compact or unresolved systems, higher-order multiples, and phase-space coincidences in clusters and streams. Whatever the dynamical origin, the velocity difference is orbital or systemic, not a holonomy, and it varies with epoch.

Line-of-sight superpositions. Two unrelated stars at the same sky position can have any ΔRV . The astrometry match rejects the clean cases (different distance and motion), but the permissive missing-value treatment and the finite error floors leave a residual. The residual is largest in the Galactic plane, where the stellar density is highest.

Composite spectra. At sub-arcsecond separations the RVS measurement can be a blend of two spectra. The catalogue may still report two radial velocities, but each is a pipeline product of a contaminated or deblended spectrum. A composite spectrum is not a copy of either star, and its velocity difference need not be dynamical.

The three discriminators are designed to kill all three. First, a holonomy pair is an *exact copy*: the two spectra have the same line strengths and ratios, shifted by one multiplicative factor. A Doppler shift is also multiplicative to first order, so the copy test, not the shift alone, is what separates a holonomy from two different stars. Second, the shift is *constant in time*: a binary’s ΔRV varies with orbital phase, a holonomy’s does not. Third, the shift is *shared across independent pairs* from the same pocket: the base e^Λ is one number, while binaries and superpositions have random shifts. The first is tested by the mean-spectrum log- λ ratio fit, the second by the epoch-RV test, and the third by the shared-base family statistic. Section 7 reports the first two; Section 4 reports the third. The hardest background is a binary of twin stars, where the copy test alone can pass; only the epoch test and the family base resolve it.

3 The census

3.1 Data

The input is the Gaia DR3 `gaia_source` table [13, 14]. The catalogue contains 33,812,183 sources with a published radial velocity. This is the RVS subset, not the full $\sim 2 \times 10^9$ -source catalogue: the census is an RV–RV survey, so both members of a pair must have a radial velocity. The RVS sample is bright and not volume-complete; the selection function is discussed in Section 8. The

same release contains 999,645 mean RVS spectra and 1,898 sources with an RV time series (43,298 epoch rows) in `vari_epoch_radial_velocity`; these products are used in Section 7.

The DR3 TAP schema contains no Gaia–Gaia neighbour table: all `*neighbour*` tables are external-catalogue crossmatches. The co-spatial pair search therefore requires a self-join of `gaia_source`. The self-join is the expensive part of the census: both sides must be spatially constrained, or the query hits the service’s job-timeout cap.

3.2 Tiling by `source_id`

A Gaia `source_id` encodes a level-12 HEALPix pixel and a counter, so a contiguous `source_id` range is a contiguous sky region. The census tiles the sky by `source_id` ranges,

$$\text{source_id} \in [t 2^{2k+47}, (t+1) 2^{2k+47} - 1], \quad (9)$$

where k is the tiling level and $t = 0, \dots, 12 \cdot 4^k - 1$. At level 3 there are 768 tiles, each of order 7° across, with a range width of 2^{53} . The level-3 tile is the largest unit that completes within the TAP sync cap in the dense plane; level 2 times out.

On a server job-timeout (HTTP 408), a tile is split into its four level- $(k+1)$ children, recursively to level 5. The parent is marked `SPLIT`; the children carry the pairs. The final cache contains 768 level-3 files, 224 level-4 files, and 40 level-5 files: 1,032 pair CSVs. Pairs straddling a tile boundary are lost, because both sides of the join must lie in the same `source_id` range. The lost strip is of order $1''$ on each side of a 7° tile, a fraction $\sim 10^{-4}$ of the area per boundary; this is negligible for a census and is documented as a selection caveat.

3.3 The self-join query

The query selects unordered pairs with both radial velocities present and a separation below $1''$ (0.000278 deg):

```
SELECT a.source_id AS aid, b.source_id AS bid,
       3600.0*DISTANCE(POINT('ICRS',a.ra,a.dec),
                      POINT('ICRS',b.ra,b.dec)) AS sep,
       a.ra, a.dec, b.ra, b.dec,
       a.parallax, b.parallax,
       a.parallax_error, b.parallax_error,
       a.pmra, b.pmra, a.pmra_error, b.pmra_error,
       a.pmdec, b.pmdec, a.pmdec_error, b.pmdec_error,
       a.radial_velocity, b.radial_velocity,
       a.radial_velocity_error, b.radial_velocity_error
FROM gaiadr3.gaia_source a
JOIN gaiadr3.gaia_source b ON a.source_id < b.source_id
WHERE a.radial_velocity IS NOT NULL
      AND b.radial_velocity IS NOT NULL
      AND a.source_id BETWEEN :lo AND :hi
      AND b.source_id BETWEEN :lo AND :hi
      AND DISTANCE(POINT('ICRS',a.ra,a.dec),
                   POINT('ICRS',b.ra,b.dec)) < 0.000278
```

The `a.source_id < b.source_id` condition makes each unordered pair appear once. The CSV columns are the pair identity, the separation, the two sky positions, the parallaxes and proper

motions with errors, and the two radial velocities with errors. The loader preserves `aid` and `bid` as exact strings: a Gaia `source_id` exceeds 2^{53} , so a `float64` conversion silently rounds it. The numeric cuts do not use the identifiers, but the follow-up does, and the exact strings are kept for that reason.

3.4 Execution

The sync service has a fixed latency of about 59 s per query, independent of tile size, and a hard 60 s job cap. The census therefore runs one tile per query, with two workers and a 90 s cooldown after any server 408. A query that stalls from congestion is backed off and retried; a query that hits the job cap is split. The final pass took 9,308 s (2.6 h) of wall clock after the earlier cached tiles; the full campaign spanned several hours across a resume. The 1,032 CSVs occupy 7.8 MB. The resumed driver’s own running total was 6,943 new pairs; the full cache, read by the cut, contains 16,358 raw pairs. The census is resumable: a tile with a CSV is skipped, a `SPLIT` parent is replaced by its children, and a max-level 408 is left for a future resume. No census process is running.

The engineering history is worth one sentence, because it explains the cache structure. An earlier scratch cache was lost, taking a preliminary cut with it; the census was relaunched against a persistent, gitignored cache at `data/gaia_tiles/`. The driver distinguishes a server 408 from congestion, never marks a tile permanently lost, and logs outside the lock (an earlier version deadlocked by re-acquiring the lock inside `log()`). The final cache is the one analysed here.

3.5 Chance expectation and the design estimate

The EM-1 plan estimated about 3,400 chance co-located pairs full-sky before the astrometry cut, scaled from a 0.5° probe field. The census found 16,358 raw pairs, a factor 4.8 higher. The plan anticipated that the Galactic plane would add pairs; the size of the excess shows that a single probe field is not a reliable normalizer for a full-sky pair census. The raw pair count is therefore not a background-subtracted measurement. The astrometry cut, not the raw count, is the meaningful selection, and the completeness audit in Section 4 is what makes the selection interpretable.

The 3,400 estimate is useful as a sanity check on the strict sample. The strict complete-astrometry cut retains 1,177 matched pairs, below the chance estimate, as expected if astrometry rejects most chance coincidences. The permissive cut retains 7,217, above the estimate, as expected if the missing-value treatment lets chance pairs through. This pattern is the quantitative reason to prefer the strict sample for population statements.

4 The cuts

4.1 The 1-arcsec co-spatial pair

The first cut is the pair itself: two Gaia sources with both radial velocities and a separation $\leq 1''$. At the small Λ probed here the axial wound image stays within $1''$ of the direct image only for a narrow cone around the boost axis, as discussed in Section 2; the cut is therefore an axial-channel selector. A relic pair that is off-axis by more than the cone is not in the sample. The $1''$ radius is also the scale at which the RVS measurement can be a blend; the composite-spectrum caveat applies to the inner part of the separation distribution. The *permissive* candidate separations have median $0.768''$, 90th percentile $0.949''$, and maximum $1.000''$, so most of the permissive sample is in the blended regime. The strict candidates have median $0.824''$, slightly larger than the full-pair

median of $0.817''$; the apparent concentration is a permissive-sample effect, not a property of the complete-astrometry sample.

The pair cut is not a lens cut. A conventional lensed pair of one source would also be co-spatial and co-moving; the census cannot separate it from a holonomy pair without the spectrum. The pair cut is the background-selection stage.

4.2 The astrometry match and its completeness

The astrometry match requires the two sources to have consistent parallax, proper motion in right ascension, and proper motion in declination at 3σ :

$$|\Delta\varpi| < 3\sigma_\varpi, \quad |\Delta\mu_{\alpha*}| < 3\sigma_{\mu_{\alpha*}}, \quad |\Delta\mu_\delta| < 3\sigma_{\mu_\delta}, \quad (10)$$

with error floors of 0.1 mas in parallax and 0.5 mas yr^{-1} in each proper-motion component. The floors keep the cut meaningful for faint stars with large formal errors. The cut is the operational statement of “same sky position, same distance, same motion”: a line-of-sight superposition of two unrelated stars does not normally satisfy it.

The implementation has an important permissiveness. The CSV loader maps an empty field to NaN, and the match sets a NaN difference to zero. A component with one *or both* values missing therefore matches *unconditionally*: the measured value is not compared with anything, and the pair passes that component regardless of its astrometry. The same applies to each proper-motion component. The label “astrometry-matched” is therefore an upper bound on the astrometrically confirmed population, not a clean sample. The strict column below removes the ambiguity by requiring all six values to be finite before the 3σ comparison.

Table 2 reports the cut flow in two versions. The *permissive* column is the cut as implemented (missing values treated as zero difference). The *strict* column requires all six astrometric values (two parallaxes, four proper-motion components) to be present before the 3σ match is applied. The difference is large: 6,040 of the 7,217 permissive matches (83.7%) have at least one missing astrometric component. The strict sample has 1,177 matched pairs, 171 above 3σ , and 135 in the candidate zone. The permissive sample has 939 candidates, of which 804 (85.6%) have at least one missing value and 135 have all six.

The equality of the permissive $> 3\sigma$ count (1,177) and the strict matched count (1,177) is a numerical coincidence that a reader should not misread. The overlap between the two sets is 171 pairs: 171 pairs are both complete-astrometry and above 3σ ; the remaining 1,006 permissive $> 3\sigma$ pairs have missing astrometry, and the remaining 1,006 strict matched pairs are within 3σ . The two 1,006s are also different sets. We report both samples in every result and recommend the strict sample for population statements.

The strict matched-pair ΔRV distribution is essentially the same as the permissive one (Table 2), so the background shape is robust to the completeness correction. What changes is the sample size and the interpretation: the permissive 939 is an upper bound on the candidate zone; the strict 135 is the complete-astrometry candidate zone.

4.3 The radial-velocity discrepancy and the candidate zone

The radial-velocity difference is $\Delta RV = v_A - v_B$, with the two catalogue errors added in quadrature. If one error is missing, the other is used alone; the 1.0 km s^{-1} floor applies only when *both* errors are missing. A pair is *RV-discrepant* when $|\Delta RV| > 3\sigma_{\Delta RV}$. The candidate zone requires both a significant and a large difference:

$$|\Delta RV| > 3\sigma_{\Delta RV} \quad \text{and} \quad |\Delta RV| > 30 \text{ km s}^{-1}. \quad (11)$$

Quantity	Permissive	Strict
Raw pairs analysed	16,358	16,358
Pairs with all six astrometric values	10,318	10,318
Astrometry-matched	7,217 (44.1%)	1,177 (11.4%)
Matched with $ \Delta RV > 3\sigma$	1,177	171
Candidate zone	939	135
ΔRV median (km s^{-1})	7.4	7.1
ΔRV 90th percentile (km s^{-1})	41.6	38.4
ΔRV 99th percentile (km s^{-1})	126.8	130.3
ΔRV maximum (km s^{-1})	489.5	472.8
Family maximum	173/939	24/135

Table 2: Cut flow. The permissive column sets a NaN astrometric difference to zero, so a component with one or both values missing matches unconditionally; the strict column requires all six astrometric values to be finite. The permissive 1,177 (matched, $> 3\sigma$) and the strict 1,177 (complete-astrometry matched) are different samples; the overlap is 171 pairs. The equality of the two 1,177 counts is a coincidence of the data, not a cross-identification. The strict matched fraction is relative to the 10,318 pairs with all six values present.

The 3σ condition removes pairs whose difference is consistent with the noise. The 30 km s^{-1} condition removes the ordinary co-moving and binary population at the high-density end of the distribution; the matched-pair 90th percentile is 41.6 km s^{-1} , so 30 km s^{-1} is in the upper tail but not an extreme cut. The zone is deliberately a background selection: a holonomy with $\Lambda = 10^{-4}$ has a 30 km s^{-1} offset at the threshold, and a holonomy with $\Lambda = 10^{-3}$ has a 300 km s^{-1} offset well inside the zone.

In the permissive sample, 1,177 of the 7,217 matched pairs are above 3σ , and 939 of those are also above 30 km s^{-1} . In the strict sample, 171 of 1,177 matched pairs are above 3σ , and 135 are in the candidate zone. The 939 is 13.0% of the permissive matched sample and 5.7% of all pairs; the 135 is 11.5% of the strict matched sample and 0.8% of all pairs. The difference between the two candidate fractions is the missing-astrometry contamination.

4.4 The axial- Λ map

For each candidate, the axial rapidity is $\Lambda_{\text{axial}} = \ln(1 + |\Delta RV|/c)$. The absolute value folds the blue and red signs; a pair’s sign is not used because the catalogue does not say which image is direct. The permissive candidate set spans 1.0006×10^{-4} to 1.6314×10^{-3} ; the strict set spans 1.0025×10^{-4} to 1.5760×10^{-3} . The window is 16 in dynamic range in the permissive sample. The largest difference is 489.5 km s^{-1} permissive (472.8 km s^{-1} strict); the smallest is the 30 km s^{-1} threshold.

Table 3 shows the candidate counts in ΔRV bins and the half-width of the 10%-of- Λ family window at the bin centre. The condition $|\Lambda_i - b| < 0.1b$ defines a half-width $0.1b$; the full window is twice that, $0.2b$. In ΔRV the half-width is proportional to Λ , so it is narrow at low ΔRV and wide at high ΔRV : 3.5 km s^{-1} at 35 km s^{-1} , 5.0 km s^{-1} at 50 km s^{-1} , 8.0 km s^{-1} at 80 km s^{-1} , 15 km s^{-1} at 150 km s^{-1} , and 35 km s^{-1} at 350 km s^{-1} ; the full window is double these values. This is the property that makes the family statistic CDF-dominated at the low end.

ΔRV bin (km s ⁻¹)	Candidates	10% half-window (km s ⁻¹)
30–40	224	3.5
40–60	282	5.0
60–100	262	8.0
100–200	153	15.0
200–500	18	35.0

Table 3: Permissive candidate counts in ΔRV bins and the half-width of the 10%-of- Λ shared-base window at the bin centre. The full window is twice the tabulated value. The window is proportional to Λ ; at the low- Λ threshold it is only a few km s⁻¹, comparable to the RVS error and the threshold spacing.

4.5 The shared-base family test

A holonomy family is a set of pairs from the same pocket that share one base e^Λ . The test scans each candidate’s Λ as a trial base and counts the candidates within a relative tolerance

$$|\Lambda_i - b| < 0.1 \max(b, 10^{-4}), \quad (12)$$

with a floor $b_{\min} = 10^{-4}$ to keep the window finite at the threshold. The reported statistic is the maximum count over all trial bases. The 10% tolerance is not matched to the RVS errors. At the low- Λ maximum the half-window is 3.9 km s⁻¹, below the typical 9.3 km s⁻¹ pair error (a relative error of about 24% at $\Lambda \sim 1.3 \times 10^{-4}$), so a family with the full measured error spread would spill outside the window. At $\Lambda \sim 1.5 \times 10^{-3}$ the half-window is about 45 km s⁻¹, roughly five times the pair error (relative error $\sim 2\%$). The statistic is therefore a coarse shared-base test: too narrow at low Λ and too wide at high Λ .

In the permissive sample the maximum is 173 of 939 candidates (18.4%), at $b = 1.29 \times 10^{-4}$, corresponding to $\Delta RV = 38.7$ km s⁻¹. The tolerance is a half-width $0.1b$: the full window there is $0.2b$, i.e. 7.74 km s⁻¹, not 3.9 km s⁻¹. The local candidate density is 22.4 candidates per km s⁻¹, so the smooth expectation in that full window is $22.4 \times 7.74 = 173.4$, and 173 are observed. The maximum is therefore the smooth density of candidates, not a surplus and not a cluster. Using the half-width as if it were the full window would halve the expectation and manufacture a factor-two surplus. In the strict sample the maximum is 24 of 135 (17.8%), at $b = 1.69 \times 10^{-4}$ ($\Delta RV = 50.6$ km s⁻¹; half-width 5.06 km s⁻¹, full window 10.11 km s⁻¹, local density 2.37 candidates per km s⁻¹, expectation 24.0, observed 24). The same low- Λ , smooth-density behaviour appears in both samples.

Composition-matched null. The maximum is a deterministic functional of the Λ multiset: any null that holds the candidate count and the Λ marginal exactly leaves it unchanged, so an assignment permutation gives $p = 1.00$ by construction. We therefore calibrate against a composition-matched bootstrap. The 939 candidates are held fixed with their sky positions and pair identities; their joint $(\Lambda, \sigma_\Lambda)$ pairs are resampled with replacement (20,000 draws), preserving the composition in expectation. The observed 173 is below the null median (179; 90% range 164–197; $p = 0.74$). In the strict sample the observed 24 is below the null median (28; 90% range 24–34; $p = 0.97$). Smooth Gamma and log-KDE composition models give $p = 0.18$ and 0.12 (strict 0.81 and 0.76); the most extreme is $p = 0.12$, and a four-look correction over two samples and two models gives 0.48, not significant. A position-fixed pool draw, which holds the pair positions fixed and replaces the candidate Λ composition with the full matched-pool composition, gives $p = 1.00$ (null median

409 permissive, 62 strict): the candidate cut removes the low- Λ pile, so this misspecified null does not manufacture an excess either. The null is calibrated on real controls that cannot contain a co-located family: matched pairs just below the 30 km s^{-1} threshold and candidate-zone pairs that fail the astrometry match give $p = 0.63\text{--}0.86$, consistent with a nominal 0.5.

The sky positions permit one further check. The 173 permissive members of the maximum window have a median nearest-neighbour separation of their pair midpoints of 5.31° , against 6.02° for position permutations of the same Λ multiset ($p = 0.05$); no member lies within 0.1° of another. The strict 24 members show no concentration (18.35° versus 15.08° , $p = 0.85$). The permissive value is marginal, is not a compact same-pocket signature, and does not survive a two-sample correction (0.11); it does not change the null.

Power. The test has little power. Injecting a sharp synthetic family (3% relative spread) at the observed low- Λ base, the global maximum reaches 95% power at about 40 edges permissive (30 edges give 0.74) and 20 edges strict (10 edges give 0.27); at high Λ (5×10^{-4} – 1.5×10^{-3}) about 200 edges are needed permissive and 40 strict, because the low- Λ background cluster remains the global maximum. With the measured 20% relative error spread, about 200 edges permissive (100 edges give 0.32) and 100 strict (50 edges give 0.88) are needed. A known base is much more sensitive: a local count at $\Lambda = 1.5 \times 10^{-3}$ detects five edges with full power, but the base is not known a priori. A realistic few-pair family is invisible.

This is the EM analogue of the GW channel’s family-test degeneracy. In the GW case the degeneracy is driven by per-event measurement noise: $\sigma_\Lambda = 0.08\text{--}0.5$ is comparable to the Λ structure the test must resolve. In the EM case the per-pair Λ error from RVS is small (a few km s^{-1} corresponds to $\sim 10^{-5}$ in Λ), but the test is driven by the steepness of the ΔRV distribution above the threshold and by the proportional window. Both are CDF-dominated in the sense that the maximum is set by the shape of the Λ distribution, not by a lattice. The EM statistic is not a detection. The calibrated null gives no excess, so the maximum is a background-density measurement, not a family. It is also not an exclusion: the null conditions on the observed composition, and the power is low, so a family already part of the Λ marginal or smaller than the thresholds above is not excluded. It is a calibrated null result for the sharp family test.

4.6 What the cut chain does not apply

The census does not apply the mean-spectrum ratio test to the 939 candidates, because in the EM-2 snapshot only 9 of the 1,132 candidate sources have a mean RVS and no candidate pair has both mean spectra. It does not apply the full epoch-series test, because no candidate pair has a published epoch series; the weak DPAC summary-statistic proxy is available for 112 candidate sources (51 RV-variable at $p < 0.01$, 1 pair with both members measured), and it is discussed in Section 7. It does not fit a lens mass, measure a magnification, or measure a time delay. The cut chain is therefore a *selection* chain: it produces a candidate-zone sample and a family statistic from catalogue radial velocities. The spectral and full-epoch tests are the next stage, and they are data-limited on DR3 (Section 7).

5 Results

5.1 Census yield and the matched fraction

The census reads 16,358 unordered raw RV–RV pairs at $\leq 1''$. The permissive astrometry match labels 7,217 (44.1%); the remaining 9,141 pairs fail at least one of the three 3σ conditions. Of

the 7,217, 1,177 are above 3σ in ΔRV , and 939 are in the candidate zone. The strict complete-astrometry sample is 1,177 matched pairs (11.4% of the 10,318 pairs with all six values), 171 above 3σ , and 135 in the candidate zone.

The permissive matched fraction is inflated by the missing-value treatment. The strict fraction, 11.4%, is the more meaningful co-moving fraction for the complete-astrometry subset. The raw matched fractions should not be compared across samples with different amounts of missing astrometry; the strict and permissive columns in Table 2 are the correct comparison.

The candidate counts are 939 permissive and 135 strict. The 939 is the number the cut script prints. It is *not* the candidate set used in the EM-2 snapshot: that snapshot used the earlier 566-candidate, 9,415-pair partial census. The 135 is the complete-astrometry subsample. We use “candidate” in the operational sense: a candidate-zone pair, not a relic candidate. The astrophysical background discussion in Section 6 explains why almost all of them are expected to be co-moving stellar systems or chance superpositions.

5.2 The ΔRV distribution

The permissive matched-pair ΔRV distribution has median 7.4 km s^{-1} , 90th percentile 41.6 km s^{-1} , 99th percentile 126.8 km s^{-1} , and maximum 489.5 km s^{-1} . The strict distribution is median 7.1 km s^{-1} , 90% 38.4 km s^{-1} , 99% 130.3 km s^{-1} , maximum 472.8 km s^{-1} . The two are consistent; the completeness correction does not change the shape. The distribution is what the design expected: a dense low- ΔRV population (co-moving pairs and binaries), a tail to tens of km s^{-1} , and a sparse high- ΔRV tail to $\sim 500 \text{ km s}^{-1}$. The 30 km s^{-1} candidate threshold is near the 80th–90th percentile, so the candidate zone is the upper tail of the matched population, not a separate peak.

The candidate ΔRV distribution is steeper than the matched distribution because the threshold removes the low end. The permissive candidates have 10th, 25th, 50th, 75th, 90th, and 99th percentiles of 34.5, 40.5, 56.7, 87.1, 118.7, and 272.2 km s^{-1} . The mode is just above the threshold; the tail extends to 489.5 km s^{-1} . This shape is the direct cause of the family statistic’s low- Λ maximum.

The permissive candidate separations have median $0.768''$, 90th percentile $0.949''$, and maximum $1.000''$; the full pair sample has median $0.817''$ and 90th percentile $0.965''$. The permissive candidates are slightly more concentrated at small separations, consistent with a sub-arcsecond blend or a physical association. The strict candidates do not share that concentration: their median is $0.824''$, slightly larger than the full-pair median, so the effect is confined to the permissive sample and is plausibly a missing-astrometry selection. The $1''$ cut is the hard boundary; no candidate is outside it by construction.

5.3 The family test

The permissive maximum is 173 of 939 (18.4%) at $\Lambda = 1.29 \times 10^{-4}$ ($\Delta RV = 38.7 \text{ km s}^{-1}$; half-width 3.87 km s^{-1} , full window 7.74 km s^{-1}). The local density is 22.4 candidates per km s^{-1} and the smooth expectation is 173.4, so the observed 173 is the density itself, not a surplus. The strict maximum is 24 of 135 (17.8%) at $\Lambda = 1.69 \times 10^{-4}$ ($\Delta RV = 50.6 \text{ km s}^{-1}$; half-width 5.06 km s^{-1} , full window 10.11 km s^{-1}), with local density 2.37 and expectation 24.0 against 24 observed. Both maxima are at the low- Λ end, where the candidate density is highest and the proportional window is narrowest. The maximum is a threshold/CDF effect, not a shared base and not a cluster.

A true family would have a different signature: a narrow peak at a common base, with a width set by the Λ errors, and the same base would appear in independent pairs. At $\Lambda \sim 10^{-3}$ the 10%

half-window is $\sim 50 \text{ km s}^{-1}$ (full window $\sim 100 \text{ km s}^{-1}$), much wider than the RVS errors, so a true family at the high end would not be hidden by the window; it would be hidden by the background in that window. At the low end the half-window is only a few km s^{-1} , below the typical 9.3 km s^{-1} pair error, so the maximum is dominated by the smooth density of candidates just above threshold and a family with the full error spread would partly spill out of the window. The composition-matched null (Section 4.5) puts the observed maximum below its median in both samples ($p = 0.74$ permissive, $p = 0.97$ strict), so the threshold pile-up is a calibrated background, not a family. The test still has little power: a sharp family needs about 40 edges at the observed low- Λ base (strict 20) and about 200 at high Λ for 95% power; with the measured 20% error spread about 200 (strict 100) are needed. A realistic few-pair family is not detectable, so we do not claim a family or an exclusion.

5.4 No significant signal

The census returns no significant signal, no detection, and no exclusion. The operational statement is:

At $\leq 1''$, with both Gaia DR3 radial velocities, the full-sky RV–RV pair census finds 16,358 raw pairs; the permissive cut labels 7,217 astrometry-matched, 1,177 above 3σ , and 939 in the candidate zone. The strict complete-astrometry cut labels 1,177 matched, 171 above 3σ , and 135 in the candidate zone. The shared-base family statistic is CDF-dominated at the probed $\Lambda \sim 10^{-4}$ – 10^{-3} ; its maximum (173/939 permissive, 24/135 strict) equals the smooth local density and is not a cluster. A composition-matched null puts the maximum below its median in both samples ($p = 0.74$ permissive, $p = 0.97$ strict), so no excess survives; the null is calibrated on real non-candidate controls ($p = 0.63$ – 0.86). The test has little power (about 40/20 edges at the observed base and 200/40 at high Λ for 95% power), so we draw no exclusion. No candidate has the spectral or epoch data needed to test the exact-copy and constant-shift discriminators. We claim no detection.

The candidate set is treated as a background sample. The candidate-zone sample is the background for the spectral follow-up; it is not a spectral exclusion on the 939, because the spectra are not available for both members.

5.5 Comparison with the preliminary cut

An earlier partial cut, on 1,475 pairs from the first ~ 108 tiles with the dense-plane tiles still pending, gave 777 astrometry-matched (52.7%), 119 above 3σ , and 92 candidates, with a family maximum of 25/92. The final census supersedes those numbers. The final matched fraction is lower (44.1%) because the added plane tiles contribute more chance pairs with missing astrometry; the final candidate fraction ($939/16,358 = 5.7\%$) is close to the preliminary $92/1,475 = 6.2\%$. The family lesson is unchanged: the preliminary $25/92 = 27.2\%$ and the final $173/939 = 18.4\%$ are both dominated by the low- Λ pile-up. The preliminary numbers are recorded here only to make the supersession explicit; all results in this paper use the full census.

6 The astrophysical background

6.1 Co-moving pairs

The astrometry match selects pairs with the same distance and the same motion. A bound binary, a higher-order multiple, or two members of a cluster or stream satisfy it naturally. The ΔRV distribution is the projected relative velocity of the population: a few km s^{-1} for wide binaries, tens of km s^{-1} for compact systems and cluster members, and a tail to hundreds of km s^{-1} for chance phase-space coincidences at the distance and motion resolution of the catalogue.

The wide-binary expectation is a useful sanity check. A $1 M_{\odot} + 1 M_{\odot}$ binary at 10^3 AU has an orbital velocity of order 1 km s^{-1} ; at 10^4 AU it is 0.3 km s^{-1} . The observed candidate tail at $30\text{--}490 \text{ km s}^{-1}$ is therefore not a wide-binary tail. It is either a compact or unresolved system (where the catalogue’s two entries may be a blend or a photocenter artefact), a higher-order multiple, or a phase-space coincidence in a dense region. The strict sample’s matched-pair distribution is the same, so the conclusion is not an artefact of the missing astrometry. The high- ΔRV tail is astrophysical, but its dynamical interpretation is not a simple wide binary.

This matters for the holonomy test. A binary’s ΔRV is orbital and varies with epoch; a holonomy’s is fixed. The epoch test is therefore the right discriminator, and the wide-binary population is not the hard background. The hard background is a compact or unresolved system whose spectra are similar enough to pass a single-epoch copy test, and the twin-binary case is the one that needs the family base as well as the epoch test.

6.2 Line-of-sight superpositions

A line-of-sight superposition of two unrelated stars at the same sky position can have any ΔRV . The astrometry match rejects the clean cases, where the two stars have different parallaxes or proper motions. The residual is the population that happens to have consistent astrometry within the errors, which is largest in the Galactic plane and at faint magnitudes where the errors are large. The permissive missing-value treatment adds to this residual: a pair with one measured astrometry and one missing is not a confirmed co-moving pair.

The design estimate of about 3,400 chance pairs is the scale of this background before the astrometry cut. The strict matched count of 1,177 is below it, consistent with astrometry removing most chance coincidences; the permissive matched count of 7,217 is above it, consistent with the missing-value treatment retaining chance pairs. This is the quantitative case for the strict sample.

6.3 Composite spectra

At sub-arcsecond separations the RVS measurement can be a blend. The catalogue may report two radial velocities, but each is a pipeline product of a contaminated or deblended spectrum. A composite spectrum is not an exact copy of either star, and its velocity difference need not be dynamical. The mean-spectrum ratio test is the clean route: it works on the spectrum directly, and it can reject a composite even when the catalogue reports two velocities. The 12 both-RVS pairs in the EM-2 snapshot (not the final census) are all non-candidates with small ΔRV ; they are the methods sample for that test (Section 7).

6.4 Why the family test is the discriminator, and why it is CDF-dominated

A holonomy family shares one base e^{Λ} across independent pairs. Binaries, superpositions, and composites do not: their shifts are random. The family test is therefore the only cut in the chain

that can kill the astrophysical background without a spectrum. That is why it is the centrepiece of the survey design.

At the current Λ dynamic range, however, the test is CDF-dominated. The candidate distribution is steepest just above the 30 km s^{-1} threshold, and the shared-base window is proportional to Λ , so the window is narrowest where the density is highest. The maximum is the low- Λ threshold pile-up. The same lesson was measured in the GW channel, where the family test is degenerate at current parameter-estimation precision; the EM case is a different mechanism with the same outcome. The family test is a valid *design* discriminator, but it is not, by itself, a detection statistic at this dynamic range. The composition-matched null (Section 4.5) calibrates the background it sees, and it finds no excess, but the power is low: a larger Λ dynamic range with a narrower relative window, or a known base, would be needed to make it decisive.

6.5 Multiplicity per pocket

The family test assumes that one pocket writes many pairs. For a metre-scale relic the entry shell contains 10^{-51} – 10^{-17} stars (Table 1), so the expected multiplicity is essentially zero; the family channel is not open. For a pc-scale pocket, the entry shell can be star-populated, and one pocket can write many axial pairs with the same base. The family test is therefore a pc-scale-pocket discriminator in principle. It would bound the pc-scale population that produces multiple axial pairs only if the statistic were calibrated *and* the selection efficiency and pocket multiplicity were modeled; the background is now calibrated, but the efficiency and multiplicity are not, so we draw no such bound. This is the same pc-scale caveat that applies to the pair channel itself.

7 EM-2 and the Gaia DR3 data limitation

Snapshot warning. Every EM-2 number in this section and in Table 4 is from a snapshot taken on 2026-10-26 19:00: 524 tile CSVs, 9,415 pairs, 566 candidates, and 1,132 unique candidate sources. The final census has 16,358 raw pairs and 939 candidates. We did not rerun EM-2 on the final cache; thus the 9/1,132, 12, 112, 51, and one-pair counts are snapshot counts, not final-census counts, and all fractions below inherit that limitation.

7.1 What DR3 exposes

The EM-2 test requires spectra at more than one epoch. Gaia DR3 provides mean RVS spectra for 999,645 sources, and an RV time series for 1,898 sources (43,298 epoch rows) in the `vari_epoch_radial_velocity` table. It does not provide epoch RVS spectra. The DataLink service exposes RVS (mean), XP_CONTINUOUS, XP_SAMPLED, MCMC_MSC, MCMC_GSPHOT, EPOCH_PHOTOMETRY, and ALL; it rejects EPOCH_SPECTRUM, RVS_SPECTRUM, EPOCH_RVS, EPOCH_RV, EPOCH_RADIAL_VELOCITY, RVS_EPOCH, SPECTRUM, and MEAN_RVS as unknown retrieval types. The TAP schema contains no RVS epoch-spectrum table; the only epoch-RV table is `vari_epoch_radial_velocity` and its `gaiafpr` copy. The DR3 release notes list mean RVS spectra and RV time series, not epoch spectra [14]. This is the data limitation: the full epoch-series test is not executable on DR3 candidate data. A weaker summary-statistic proxy is available for 112 candidate sources, but it tests constancy, not the exact-copy ratio.

The mean-spectrum product is a VOTable 1.4 BINARY2 table with 2,401 rows, fields `wavelength` (nm), `flux`, and `flux_error`, covering 846–870 nm at 0.01 nm sampling. It has no per-epoch axis. A source with a mean spectrum and a mean RV gives one spectrum and one velocity, not a time series. The epoch RV series, where it exists, gives velocities but not spectra;

it can test constancy but not the exact-copy ratio. The two products are complementary, and the candidate set has neither in pair-complete form. The 1,898-source epoch-RV sample has zero overlap with the pair sample: 0 of the 18,830 snapshot pair sources have epoch RVs. No epoch constancy test can therefore be run on a pair using the published series.

7.2 The test statistics

Epoch-RV constancy. For one source with epochs (t_i, v_i, σ_i) , form the inverse-variance weighted mean $\bar{v}$ and

$$\chi^2 = \sum_i \left(\frac{v_i - \bar{v}}{\sigma_i} \right)^2, \quad p_{\text{const}} = P(\chi_{N-1}^2 > \chi^2). \quad (13)$$

A holonomy pair requires both members to be consistent with a constant RV at every epoch; a binary requires at least one member to vary. For a pair define

$$R = \frac{A_A + A_B}{|\Delta\text{RV}|}, \quad (14)$$

where $A_i = \max(v_i) - \min(v_i)$ over the epoch series (or the DPAC `rv_amplitude_robust` when the series is unpublished) and ΔRV is the pair’s catalogue difference. The verdict is *binary-compatible* if either source has $p_{\text{const}} < 0.01$ or $R \geq 1$; a *constant-shift candidate* requires both sources to be constant and $R < 1$. The second branch is deliberately weak: a long-period binary whose baseline does not cover a full orbit can hide its variation. It is a screening certificate, not a detection.

For 112 candidate sources in the EM-2 snapshot the epoch series is unpublished but the DPAC variability summary is present; the script uses `rv_chisq_pvalue` and `rv_amplitude_robust` as the weaker version of the same test.

Mean-spectrum ratio. For two mean spectra, work in $x = \ln \lambda$. A multiplicative frequency ratio $r = \nu_B/\nu_A$ is a pure shift $d = \ln(\lambda_B/\lambda_A) = -\ln r$, so a holonomy pair must satisfy

$$y_B(x + d) \approx a y_A(x), \quad d \text{ independent of } x, \quad (15)$$

after continuum normalization. The script estimates d by cross-correlation over a narrow range around the catalogue-predicted shift, then estimates d independently in six log-wavelength segments. The discriminants are the residual $c(\exp d_{\text{best}} - \exp d_{\text{cat}})$, the Pearson correlation after alignment, and the robust segment-to-segment scatter $1.4826 \text{ MAD}(d_k)$ in km s^{-1} . A true holonomy pair is an exact scaled copy: high Pearson r , segment scatter at the estimator floor, and a residual consistent with zero. Two different stars are not.

7.3 Validation

Epoch controls. Control 6427507623204526848 has 61 epochs, an amplitude of 168.2 km s^{-1} , and $p_{\text{const}} \approx 0$: variable, as expected. Control 412675610446139520 has 62 epochs, an amplitude of 32.9 km s^{-1} , and $p_{\text{const}} = 3.8 \times 10^{-125}$: variable, as expected.

Synthetic null. A constant series with 20 epochs, $\sigma = 2 \text{ km s}^{-1}$, and 300 trials gives 1 false positive in 300 at $p < 0.01$ (0.003) and a median p of 0.458. This is a weak consistency check, not a calibration: the binomial confidence interval on $1/300$ is wide, and the Gaussian $\sigma = 2 \text{ km s}^{-1}$ error model is not the real RVS error distribution. We do not attach a calibrated false-positive rate to the constancy test.

Pair A	Pair B	sep (")	ΔRV_{cat}	ΔRV_{spec}	resid	r	seg. scatter
147586243206612224	147586247500686336	0.77	-1.67	+3.84	+5.51	0.35	10.9
1444367711752056960	1444367716047984128	0.89	-3.41	+0.21	+3.63	0.79	2.8
1577779192786031488	1577779192787232384	0.95	+1.00	-0.36	-1.36	0.69	3.1
1605975687444174720	1605975687445529472	0.87	+1.36	-0.36	-1.72	0.75	2.8
1618370615167692544	1618370619463553024	0.97	-2.05	-0.19	+1.86	0.88	0.6
2617971996444786688	2617971996445201024	0.78	+0.75	-7.11	-7.86	0.59	17.5
3487101333650032512	3487101337946708736	0.74	-8.73	+10.31	+19.04	0.48	15.4
3523275476657585408	3523275476658292608	0.92	-3.17	+0.19	+3.36	0.67	2.6
3534416827981511296	3534416827983300736	0.80	-4.19	+0.27	+4.46	0.73	4.9
3719080664859936896	3719080664860599424	0.87	-1.56	+0.36	+1.92	0.61	6.6
3771755793045566464	3771755793046208384	0.86	+5.90	-3.88	-9.78	0.65	4.6
3884641758516220544	3884641758517313024	0.83	-2.18	+4.00	+6.19	0.72	2.9

Table 4: The 12 both-RVS pairs in the EM-2 snapshot (524 tiles, 9,415 pairs; not the final 16,358-pair census). Full source identifiers are listed. ΔRV is $v_B - v_A$; the residual is spectral minus catalogue; the segment scatter is the robust six-segment scatter in km s^{-1} . The most copy-like pair (row 5) has scatter 0.63 km s^{-1} below the $\sim 1 \text{ km s}^{-1}$ estimator floor and residual within the 1.09 km s^{-1} validation error; it is excluded by $|\Delta RV| = 2.05 \text{ km s}^{-1} < 30 \text{ km s}^{-1}$, not by a demonstrated failure of the copy test.

Injection and recovery. Injecting a known multiplicative shift into a real candidate RVS spectrum (338067080711985408) and recovering it with the same estimator gives: $-100 \rightarrow -100.35$, $-30 \rightarrow -31.09$, $0 \rightarrow 0.00$, $+30 \rightarrow +31.09$, $+100 \rightarrow +100.34 \text{ km s}^{-1}$, with errors at or below 1.09 km s^{-1} . The parser and the multiplicative-shift estimator work at the $\sim 1 \text{ km s}^{-1}$ level on real RVS data. The failure of real pairs to look like copies is therefore a property of the spectra, not of the estimator.

The 12 both-RVS pairs. In the *snapshot* (524 tiles, 9,415 pairs), twelve pairs have both mean spectra. All are non-candidates: their catalogue $|\Delta RV|$ is $0.75\text{--}8.7 \text{ km s}^{-1}$, below the 30 km s^{-1} threshold. They are the methods sample for the ratio test, not holonomy candidates; a rerun on the final 16,358-pair cache could change the sample. Table 4 lists them with full identifiers. The sample is not a clean set of non-copies. Pearson is $0.35\text{--}0.88$, the segment scatter is $0.6\text{--}17.5 \text{ km s}^{-1}$ (the estimator floor is $\sim 1 \text{ km s}^{-1}$), and the spectral shift disagrees with the catalogue shift by up to 19 km s^{-1} for most pairs. But the most copy-like pair, 1618370615167692544/1618370619463553024, has segment scatter 0.63 km s^{-1} , below the $\sim 1 \text{ km s}^{-1}$ estimator floor, and residual $+1.86 \text{ km s}^{-1}$, within the 1.09 km s^{-1} injection-recovery error. It is excluded as a holonomy candidate by its catalogue $|\Delta RV| = 2.05 \text{ km s}^{-1} < 30 \text{ km s}^{-1}$, *not* by being demonstrably non-copy. The 12-pair sample therefore cannot support the statement that no pair is an exact copy; it is a low- ΔRV methods sample, not a holonomy measurement.

7.4 The candidate-level blocker

The EM-2 snapshot is the 566-candidate set available on 2026-10-26 19:00, with 1,132 unique candidate sources. It is not the final 939-candidate set. The binding limit is data overlap, not census size: the 1,898-source epoch-RV sample has zero overlap with the pair sample, because 0 of the 18,830 snapshot pair sources have epoch RVs. A larger census therefore cannot supply the epoch series that DR3 does not publish. The full epoch-series test is not executable; the weak summary-statistic proxy is available. The snapshot numbers are:

- candidate sources: 1,132;

Pair A	Pair B	sep (")	ΔRV (km s ⁻¹)	RVS	Astr.
1935640742676411136	1935640746975070592	0.97	+109.0	A	full
2240940147792823808	2240940147794672512	0.65	-63.7	B	partial
3021533647633420032	3021533651928202752	0.99	+61.6	A	partial
3031724887776817664	3031724887778377472	0.84	+42.2	B	full
338067080711985408	338067085006103936	0.83	+43.7	A	partial
3342499920918698368	3342499925211852800	0.94	-37.7	A	partial
3531676054731889280	3531676054733307136	0.87	-146.4	B	partial
617517066240659200	617517070535728128	0.70	+61.7	A	partial
896264154603714560	896264158897940352	0.87	-99.3	A	partial

Table 5: The nine candidate pairs with one mean RVS spectrum. ΔRV is $v_A - v_B$. “RVS” is the member with the published mean spectrum. “Astr.” is “full” if all six astrometric values are present for both members, “partial” if the counterpart is missing its parallax and proper motion. The two “full” pairs are the highest-priority follow-up targets.

- candidate sources with a mean RVS spectrum: 9;
- candidate sources with published epoch RVs: 0;
- candidate pairs with both mean RVS spectra: 0;
- candidate pairs with both epoch series: 0;
- candidate sources with RV variability summary statistics: 112;
- of those, RV-variable at $p < 0.01$: 51;
- candidate pairs with both summary statistics: 1;
- snapshot pair sources: 18,830;
- snapshot pair sources with epoch RVs: 0.

The 9 candidate sources with a mean RVS each have a counterpart with no RVS product. The counterpart’s DataLink list is empty or contains only MCMC products. The 9 pairs are the natural first targets for targeted multi-epoch spectroscopy (Table 5). The 12 both-RVS pairs are all non-candidates, as noted; they are the methods sample.

The one pair with both summary statistics is 2527814581292865024/2527814585588920704: separation 0.715", $\Delta RV = +110.0$ km s⁻¹, $p_{\text{const}}(A) = 0.0557$, $p_{\text{const}}(B) = 0.335$, $A_A = 18.7$ km s⁻¹, $A_B = 5.1$ km s⁻¹. Neither is formally variable at $p < 0.01$; $A_A + A_B = 23.7$ km s⁻¹ < 110 km s⁻¹, so the weak summary-statistic branch labels it a *constant-shift candidate*. This is not a detection: the p -values are not significant, the sources have only 8 and 3 transits, and no epoch series is available to exclude a long-period binary. Of the 112 candidate sources with summary statistics, 51 are RV-variable at $p < 0.01$; that is the expected binary background. The remaining 61 are constant under this weak test, but only one *pair* has both members measured.

The NSS cross-match is too sparse to serve as a binary veto: of the 1,132 candidate sources, 1 appears in `nss_two_body_orbit`, and 0 in `nss_non_linear_spectro` or `nss_acceleration_astro`. The `rv_chisq_pvalue/rv_amplitude_robust` sample is the better current handle.

The completeness audit matters for the follow-up. Seven of the nine pairs have a counterpart with no published parallax or proper motion; only 1935640742676411136/1935640746975070592 and 3031724887776817664/3031724887778377472 are complete. A constant RV offset in a partial

pair could be a chance superposition; the two complete pairs are the cleanest targets. The targeted campaign should obtain both spectra and counterpart astrometry for the partial pairs, not only the spectra.

7.5 Next steps

The decisive follow-up is targeted multi-epoch spectroscopy, not a larger Gaia census. The nine pairs with one RVS spectrum are the first targets; 2–4 epochs over months at RVS-like resolution would directly test the constant-offset hypothesis. The same campaign can cover a ranked subset of the candidates, starting with the 135 complete-astrometry candidates. For each pair, the campaign should:

1. obtain a spectrum of the counterpart (the member without a mean RVS), so the $\log\lambda$ ratio test can run on both members;
2. obtain 2–4 epochs per member, so the epoch-constancy test can distinguish a fixed holonomy shift from an orbital ΔRV ;
3. obtain or confirm counterpart astrometry for the seven partial pairs, so the co-spatial interpretation is not a missing-value artefact;
4. calibrate the exact-copy threshold on the 12 both-RVS pairs and on injected shifts in real spectra.

The archival route is a future Gaia epoch-spectrum release. DR3’s public DataLink does not contain epoch spectra; do not plan around `EPOCH_SPECTRUM` until the product list shows it. When such a release exists, the 135 complete-astrometry candidates and the 1,898-source epoch-RV sample are the natural starting points. The family test should be re-run on the spectra, with a calibrated null and a narrower relative window, once a candidate set with spectra exists.

8 Limits

8.1 What the non-detection establishes

The census establishes that the full-sky Gaia DR3 RV–RV pair sample, at $\leq 1''$ and with the cuts above, contains no significant shared-base signal at the probed $\Lambda \sim 10^{-4}$ – 10^{-3} . The family maximum is the smooth candidate density, not a cluster (Section 5), and the composition-matched null puts it below its median in both samples ($p = 0.74$ permissive, $p = 0.97$ strict; Section 4.5). The census also establishes that the raw pair yield is consistent with astrophysical co-moving pairs and chance superpositions, with no evidence for a holonomy population, and that the candidate set is available for follow-up, with 135 complete-astrometry pairs and 9 pairs with one mean RVS in the snapshot.

It does not establish a spectral exclusion on the 939 candidates, because the spectra are not available for both members. It does not establish a time-constancy exclusion, because no candidate pair has an epoch series. It does not establish a lens-mass exclusion, because the census does not fit lens masses. It is a calibrated null in a low-power pair-family statistic, not a statistical exclusion: the null conditions on the observed composition, and a family smaller than the measured detection thresholds is not excluded.

8.2 No density bound

No density bound is claimed. The conversion from a non-detection to a relic number density requires three ingredients that are not modeled: the selection efficiency ε (the fraction of relics that would produce a detectable axial pair in the RVS sample), the pair multiplicity k (the number of detectable pairs per pocket), and the pocket luminosity function or volume distribution. The census measures none of these. The cloistering calculation (Table 1) shows that the canonical metre-scale relic produces essentially zero pairs in a stellar survey; the channel is open only for pc-scale pockets or the trigger epoch. A density bound from this census would therefore apply to a pc-scale population with an assumed multiplicity and an assumed alignment distribution. We do not have those assumptions, and we do not quote a density.

A naive zero-observed Poisson reference would give $\mu < 2.3$ at 90% confidence, but we do not use it as an EM limit: it assumes that the 173 maximum is a background cluster and therefore has no data content beyond the zero count. For context only, the GW companion’s 5,000-permutation calibrated background on the corrected M_c/D_L scaling cut gives $\mu < 274$ at 90% CL (206 model-dependent; the Gamma–Poisson model is not adequate at $C > 0.1$), or $\mu < 4.50$ (4.95 conservatively) at $C > 0.9$. Those are limits on detectable families in the GW channel, not EM limits; the EM census now has a calibrated composition-matched background, but it still has no selection-efficiency or multiplicity model.

8.3 Systematics

Permissive astrometry. The missing-value treatment inflates the permissive matched and candidate counts. The strict sample is the reliable one. Any population statement should use the strict 135, not the permissive 939.

Error floors. The 0.1 mas parallax and 0.5 mas yr^{−1} proper-motion floors make the match permissive for faint stars. A stricter floor would reduce the matched sample; a larger floor would increase it. The floors are a design choice, not a measurement.

Tile boundaries. Pairs straddling a tile boundary are lost at the $\sim 10^{-4}$ level per boundary. This is negligible for the census totals but is a known incompleteness.

RVS blending. At sub-arcsecond separations the RVS measurement can be a blend. The permissive candidate separations are concentrated at 0.4–1.0'' (median 0.768''), so blending is a real systematic for the inner part of that sample; the strict median is 0.824''. The mean-spectrum test is the clean route.

Chance normalization. The design estimate of 3,400 chance pairs is 4.8 times below the raw yield. The probe field was not a representative normalizer. The astrometry cut, not the raw count, is the background separator.

Family tolerance. The 10% tolerance is a coarse choice: it is a half-width $0.1b$, so the full window is $0.2b$, and at the maximum it is 7.74 km s^{-1} , not 3.9 km s^{-1} . It is also not matched to the RVS errors: the half-window at the observed base is 3.9 km s^{-1} against a typical 9.3 km s^{-1} pair error, while at $\Lambda \sim 1.5 \times 10^{-3}$ it is about 45 km s^{-1} . The tolerance is now calibrated against the composition-matched null (Section 4.5); a narrower tolerance would reduce the low- Λ maximum

but would also reduce the sensitivity to a real family with the measured error spread. The tolerance should be chosen with the family error model before a family claim is made.

Snapshot versus final census. Every EM-2 number is from the 566-candidate, 9,415-pair snapshot, not the final 939-candidate, 16,358-pair census. The 9/1,132, 12, 112, 51, and one-pair counts are snapshot counts; a rerun on the final cache would change them. The binding limit is the zero overlap between the 1,898-source epoch-RV sample and the 18,830 snapshot pair sources (0 with epoch RVs), so the qualitative conclusion is unchanged.

8.4 Sensitivity

The census is sensitive to the 33,812,183 RV sources, not to the full Gaia catalogue. The RVS sample is bright and not volume-complete; the selection function is not modeled here. The $1''$ cut selects the axial channel, so the census is blind to off-axis pairs at larger separations. The Λ window is 10^{-4} – 1.6×10^{-3} ; the gap to the GW window is 3×10^{-3} –0.1. The survey is a pc-scale pocket search; the canonical metre-scale relic is causally cloistered. The trigger epoch, where the many-image pile-up opens, is a distinct search not covered here.

9 Survey strategy

9.1 Gaia DR4 and epoch spectra

The decisive archival route is a future Gaia epoch-spectrum release. DR3 exposes mean RVS spectra and RV time series but no epoch spectra; the DataLink product list and the TAP schema are explicit. A DR4 epoch-spectrum product would make the 135 complete-astrometry candidates and the 1,898-source epoch-RV sample testable without new telescopes. Until the product list shows EPOCH_SPECTRUM, the survey should not plan around it.

9.2 Targeted multi-epoch spectroscopy of the nine

The nine pairs with one mean RVS are the natural first targets. Two to four epochs over months at RVS-like resolution, covering 846–870 nm, would directly test the constant-offset hypothesis. The two complete pairs are the cleanest; the seven partial pairs need counterpart astrometry as well as spectra. The same campaign can cover a ranked subset of the 135 complete-astrometry candidates. The exact-copy threshold should be calibrated on the 12 both-RVS pairs and on injected shifts in real spectra, as in the EM-2 validation.

9.3 Rubin/LSST

Rubin/LSST is entering its time-domain survey and will produce large alert streams. The EM pair channel there is photometric: transient pairs, repeated flares at fixed positions with fixed flux ratios, and no microlensing light curve. Broadband photometry does not resolve the frequency ratio of a stellar spectrum, so LSST is a trigger and a selection tool, not a ratio test. Spectroscopic follow-up is required to apply the exact-copy and family discriminators. The same caveat applies to OGLE and the Gaia alert stream.

9.4 VLBI and pc-scale pockets

VLBI pair imaging targets the pc-scale window: two compact components with identical structure and no mass surface density. The angular resolution is well matched to the mas-scale separations that a pc-scale pocket can produce, and the pair morphology is a direct test of the no-lens property. The channel is limited by the number of sufficiently bright, compact sources behind such a pocket; it is a targeted follow-up, not a full-sky census.

9.5 A wider-separation census for off-axis pairs

The $1''$ cut selects the axial channel. At small Λ the off-axis tilt is $\sim \Lambda\theta$, so a wider-separation census ($1'$) would capture off-axis pairs at $\Lambda \sim 10^{-3}$ that the current census misses. The price is a chance background that grows roughly as the separation squared, and a fixed ratio that is broadened by the off-axis continuum $r = \cosh \Lambda - n_z \sinh \Lambda$. A wider census would need spectroscopic follow-up to control the background. It is a natural extension, not a substitute for the axial census.

9.6 Combined EM and GW

The EM and GW channels probe different Λ ranges and have independent systematics. A coincident candidate—a GW mass-shifted pair and an EM spectral pair with the same base—would be much harder to explain than either alone. No such candidate exists. The combined family test is a long-term goal: it requires a pocket that writes both channels, and the two data sets do not currently overlap in Λ . The value of the comparison is the independence of the two non-detections, not a joint detection.

10 Conclusion

The full-sky Gaia DR3 RV–RV co-spatial pair census is complete. It reads 16,358 raw pairs at $\leq 1''$ from 1,032 tile CSVs covering the 33,812,183 RV sources. The permissive cut labels 7,217 astrometry-matched, 1,177 above 3σ , and 939 in the candidate zone; the strict complete-astrometry cut labels 1,177 matched, 171 above 3σ , and 135 in the candidate zone. The matched-pair ΔRV distribution has median 7.4 km s^{-1} , 90th percentile 41.6 km s^{-1} , 99th percentile 126.8 km s^{-1} , and maximum 489.5 km s^{-1} , mapping to an axial- Λ window of 1.0×10^{-4} – 1.6×10^{-3} . The family statistic reaches 173 of 939 permissive candidates (24 of 135 strict), but this is the smooth density seen through a full 7.74 km s^{-1} window (22.4 candidates per km s^{-1} , expected 173.4, observed 173), not a cluster. The composition-matched null puts the maximum below its median in both samples ($p = 0.74$ permissive, $p = 0.97$ strict), and real non-candidate controls give $p = 0.63$ – 0.86 ; no excess survives. The statistic is CDF-dominated, as in the GW channel, and has little power: a sharp family needs about 40 edges at the observed low- Λ base (strict 20) and about 200 at high Λ for 95% power; with the measured 20% error spread about 200 (strict 100) are needed. We claim no significant signal, no detection, and no exclusion. The candidate set is consistent with an astrophysical background: co-moving binaries and multiples, line-of-sight superpositions, and composite spectra.

The epoch test is implemented, but only its epoch-constancy proxy is weakly validated (1/300 trials, wide confidence interval); the epoch-spectrum/exact-copy test has never been run on a candidate pair. Gaia DR3 exposes mean RVS spectra and no epoch RVS spectra. In the EM-2 snapshot (524 tiles, 9,415 pairs), 9 of 1,132 candidate sources have a mean RVS, 0 have published epoch RVs, and 0 candidate pairs have both mean spectra; the 1,898-source epoch sample has zero

overlap with the 18,830 snapshot pair sources. The full epoch-series test is therefore not executable on DR3 candidate data; the weak summary-statistic proxy is available for 112 candidate sources. The next step is targeted multi-epoch spectroscopy of the nine pairs with one RVS spectrum—the two complete-astrometry pairs first—and, archivally, a future Gaia epoch-spectrum release. We quote no density bound. The survey is sensitive only to pc-scale star-populated pockets (one star needs $\rho_0 \sim 2$ pc; 173 pairs need ~ 10 pc); the canonical metre-scale relic remains causally cloistered.

A Exact-number ledger

Table 6 lists the principal data numbers quoted in the text; it is not an exhaustive list of every number in the paper. Theory quantities derived in Section 2 and Appendix C, and the GW-companion context values, are not itemized here. “Cut” means the printed output of `simulations/relic_survey_gaia_cut.py` on the full cache; “EM-2” means `simulations/relic_survey_gaia_em2.py` and its snapshot report; “audit” means `simulations/gaia_strict_audit.py`, a direct re-read of the same cut CSV cache using the cut’s own column definitions. The audit does not change any number the cut prints; it reproduces the strict sample and adds the completeness breakdown. “W74” means `simulations/relic_survey_gaia_family_null.py` and its result JSON under `data/` (gitignored).

Table 6: Exact-number ledger.

Quantity	Value	Provenance
Gaia DR3 RV sources	33,812,183	TAP count; EM-1 note
Census tiles	768 L3 + 224 L4 + 40 L5	cache inventory
Pair CSVs	1,032	cache inventory
Raw pairs analysed	16,358	cut
Pairs with all six astrometric values	10,318	audit
Permissive astrometry-matched	7,217 (44.1%)	cut
Permissive matched, any missing astrometry	6,040 (83.7%)	audit
Permissive matched, $> 3\sigma$	1,177	cut
Permissive candidates	939	cut
Permissive candidates, all six values	135	audit
Permissive candidates, any missing	804 (85.6%)	audit
Strict matched	1,177 (11.4% of 10,318)	audit
Strict matched, $> 3\sigma$	171	audit
Strict candidates	135	audit
Overlap permissive $> 3\sigma$ / strict matched	171	audit
ΔRV median / 90% / 99% / max (permissive)	7.4 / 41.6 / 126.8 / 489.5 km s ⁻¹	cut
ΔRV median / 90% / 99% / max (strict)	7.1 / 38.4 / 130.3 / 472.8 km s ⁻¹	audit
Λ window (permissive)	1.0×10^{-4} – 1.6×10^{-3}	cut, Eq. (7)
Family maximum (permissive)	173/939 (18.4%)	cut
Family maximum (strict)	24/135 (17.8%)	audit
Family max location (permissive)	$\Lambda = 1.29 \times 10^{-4}$, $\Delta RV = 38.7$ km s ⁻¹	audit
Family max location (strict)	$\Lambda = 1.69 \times 10^{-4}$, $\Delta RV = 50.6$ km s ⁻¹	audit
Family 10% half-window at max	3.87 km s ⁻¹ / 5.06 km s ⁻¹	audit

Quantity	Value	Provenance
Family full window at max	7.74 km s ⁻¹ / 10.11 km s ⁻¹	audit
Local candidate density at max	22.4 / 2.37 per km s ⁻¹	audit
Expected count in full window at max	173.4 / 24.0	audit
Median σ_Λ (permissive / strict)	3.12×10^{-5} / 3.24×10^{-5}	W74
Family null p (permissive / strict)	0.74 / 0.97	W74
Family null median (permissive / strict)	179 / 28	W74
Family null 90% range (permissive / strict)	164–197 / 24–34	W74
Smooth Gamma / log-KDE p	0.18 / 0.12; 0.81 / 0.76	W74
Assignment permutation p	1.00 (invariant)	W74
Position-fixed pool draw p	1.00 (median 409 / 62)	W74
Sky association p (permissive / strict)	0.05 / 0.85	W74
Non-candidate control p	0.63–0.86	W74
Power 95%, sharp family at observed base	~40 / 20 edges	W74
Power 95%, sharp family at high Λ	~200 / 40 edges	W74
Power 95%, 20% family spread at observed base	~200 / 100 edges	W74
Candidate sep. median (permissive / strict)	0.768'' / 0.824''	audit
Full-pair sep. median	0.817''	audit
Chance design estimate	~ 3,400 full-sky	EM-1 plan
Raw/design ratio	4.8	arithmetic
Mean RVS spectra in DR3	999,645	EM-2 TAP
Epoch-RV sources / rows	1,898 / 43,298	EM-2 TAP
EM-2 snapshot size	524 tiles, 9,415 pairs, 566 candidates	EM-2 report (snapshot)
Candidate sources (EM-2 snapshot)	1,132	EM-2 report (snapshot)
Candidate sources with mean RVS	9	EM-2 report (snapshot)
Candidate sources with epoch RVs	0	EM-2 report (snapshot)
Candidate pairs with both mean spectra	0	EM-2 report (snapshot)
Candidate sources with stats	112	EM-2 report (snapshot)
Candidate sources RV-variable at $p < 0.01$	51	EM-2 report (snapshot)
Candidate pairs with both stats	1	EM-2 report (snapshot)
Snapshot pair sources	18,830	EM-2 report (snapshot)
Snapshot pair sources with epoch RVs	0	EM-2 report (snapshot)
Synthetic null FP at $p < 0.01$	0.003 (1/300)	EM-2 report (snapshot)
Synthetic null median p	0.458	EM-2 report (snapshot)
Both-RVS pairs (snapshot)	12	EM-2 report (snapshot)
Both-RVS ΔRV range	0.75–8.7 km s ⁻¹	EM-2 report (snapshot)
Both-RVS Pearson range	0.35–0.88	EM-2 report (snapshot)
Both-RVS segment scatter	0.6–17.5 km s ⁻¹	EM-2 report (snapshot)
Both-RVS residual range	–9.78 to +19.04 km s ⁻¹	EM-2 report (snapshot)
Injection recovery error	≤ 1.09 km s ⁻¹	EM-2 report (snapshot)
One-pair summary ΔRV	+110.0 km s ⁻¹	EM-2 report (snapshot)
Direct-image deflection	3×10^{-5} deg	<code>relic_signature.py</code>
Fake lens mass (order of magnitude)	~ 10 ⁴⁵ kg	geometric estimate; text
GW companion family bound ($C > 0.1$)	$\mu < 274$ (206 model-dependent)	companion [15]
GW companion family bound ($C > 0.9$)	$\mu < 4.50$ (4.95)	companion [15]

B Reproduction

The census is reproducible from the Gaia archive with the scripts in `simulations/`. The pair CSVs are derived from cache under `data/gaia_tiles/` (gitignored, not redistributed); the archive is the source of truth.

Data availability. The Gaia DR3 data are public at the ESA archive; the derived tile and EM-2 spectral caches are gitignored and not redistributed; the analysis code is available from the project repository under `simulations/`, and the exact commands are below.

```
cd "$(git rev-parse --show-toplevel)"

# EM-1: full-sky RV-RV census (resumable; ~2.6 h final pass)
.venv/bin/python simulations/relic_survey_gaia.py \
  --cache-dir data/gaia_tiles --workers 2 --level 3

# EM-1: science cut (reads all L*_T*.csv)
.venv/bin/python simulations/relic_survey_gaia_cut.py \
  --cache-dir data/gaia_tiles

# EM-1: strict-astrometry audit (reproduces the strict column)
.venv/bin/python simulations/gaia_strict_audit.py \
  --cache-dir data/gaia_tiles

# EM-1: composition-matched family null (W74)
.venv/bin/python simulations/relic_survey_gaia_family_null.py \
  --cache-dir data/gaia_tiles

# EM-2: epoch-RV / mean-spectrum test (snapshot cache; cached after first run)
.venv/bin/python simulations/relic_survey_gaia_em2.py \
  --limit 12 --controls 2
```

The driver accepts `--level`, `--workers`, `--max-tiles`, `--dry-run`, and `--cache-dir`; it honours the `GAIA_CACHE_DIR` environment variable. The cut accepts `--cache-dir`. The EM-2 script accepts `--limit`, `--controls`, and `--dry-run`. The TAP service is <https://gea.esac.esa.int/tap-server/tap/sync> with `REQUEST=doQuery`, `LANG=ADQL`, `FORMAT=csv`. The sync cap is 60 s; the fixed latency is about 59 s; a 408 triggers a split, not a permanent loss.

Two implementation cautions are worth recording. First, the cut loader keeps `source_id` as a string because the identifiers exceed 2^{53} ; a `float64` conversion silently rounds them. Second, the `TAP_SCHEMA` tables listing must be filtered locally; a server-side `WHERE` on `TAP_SCHEMA.tables` fails. Neither affects the numeric cuts, but both matter for reproducibility.

C The boost algebra

A boost of rapidity Λ along z acts on a null four-momentum $k^\mu = (\nu, \nu n_x, \nu n_y, \nu n_z)$ as

$$\begin{pmatrix} \nu' \\ \nu' n'_z \end{pmatrix} = \begin{pmatrix} \cosh \Lambda & -\sinh \Lambda \\ -\sinh \Lambda & \cosh \Lambda \end{pmatrix} \begin{pmatrix} \nu \\ \nu n_z \end{pmatrix}, \quad (16)$$

with the transverse components unchanged. The frequency ratio is

$$r = \frac{\nu'}{\nu} = \cosh \Lambda - n_z \sinh \Lambda. \quad (17)$$

For $n_z = +1$, $r = e^{-\Lambda}$; for $n_z = -1$, $r = e^{+\Lambda}$. The transformed direction cosine is

$$n'_z = \frac{n_z \cosh \Lambda - \sinh \Lambda}{\cosh \Lambda - n_z \sinh \Lambda}, \quad (18)$$

which gives the aberration. For a photon at angle θ from the axis, $n_z = \cos \theta$; expanding for small Λ gives $\theta' \simeq \theta(1 - \Lambda)$ and a tilt $\Delta\theta \simeq \Lambda\theta$; the maximum at $\theta = \pi/2$ is $\arctan(\sinh \Lambda)$. The axial channel has $\theta = 0$ and no tilt.

The radial-velocity observable follows from $\nu' \simeq \nu(1 + \Lambda)$ for small Λ : a frequency ratio r is a velocity offset $c(r - 1) \simeq c\Lambda$. The exact inversion is $\Lambda = \ln(1 + \Delta RV/c)$ for one shifted image and one direct image. The pair's

two images are the direct image and the wound image, not two wound images of opposite sign; the factor is therefore $r - 1$, not $r - r^{-1}$.

The echo follows from the boost displacement. A circuit of covering length $\Lambda\tau_{\text{wall}}$ has a time delay $\Delta t \simeq 2\tau_{\text{wall}} \sinh(n\Lambda/2)$ for winding n ; for the canonical metre pocket, $\tau_{\text{wall}} = 2.02$ m and the $n = 1$ echo is 3.4 ns. The EM census does not measure this delay; it measures the frequency ratio.

D The family statistic

For candidates $i = 1, \dots, N$ with $\Lambda_i = \ln(1 + |\Delta\text{RV}_i|/c)$, the shared-base statistic at trial base b is

$$S(b) = \sum_{i=1}^N \mathbf{1}[|\Lambda_i - b| < 0.1 \max(b, 10^{-4})], \quad (19)$$

and the reported statistic is $\max_b S(b)$. The tolerance is a half-width proportional to b , so the half-width in ΔRV is $0.1cb$ and the full window is $0.2cb$: the half-width is 3.0 km s^{-1} at $b = 10^{-4}$, 3.9 km s^{-1} at the observed permissive maximum (full window 7.7 km s^{-1}), and 48 km s^{-1} at $b = 1.6 \times 10^{-3}$ (full window 96 km s^{-1}). The candidate density is highest just above the 30 km s^{-1} threshold, so the maximum of $S(b)$ occurs at the smallest b where the density is high, even though the window is narrowest there. At the permissive maximum the density is 22.4 candidates per km s^{-1} and the full window contains $22.4 \times 7.74 = 173.4$ by smooth interpolation, matching the observed 173. This is the CDF domination: the shape of the Λ distribution, not a lattice, sets the maximum.

A true family would produce a narrow peak at a common b , with a width set by the Λ errors, and the same b would appear in independent pairs. At high b the 10% half-window is much wider than the errors, so a true high- Λ family would not be hidden by the window; it would be hidden by the background in that window. At low b the half-window (3.9 km s^{-1} at the observed permissive base) is below the typical 9.3 km s^{-1} pair error, so a family with the full error spread would partly spill out of the window and the maximum is the smooth-density value. The statistic is a coarse shared-base test, not a matched filter.

The composition-matched null of Section 4.5 is implemented in `simulations/relic_survey_gaia_family_null.py`. It holds the candidate count and sky positions fixed and resamples the $(\Lambda, \sigma_\Lambda)$ pairs with replacement (20,000 draws). The observed permissive maximum 173 is below the null median (179; 90% range 164–197; $p = 0.74$); the strict maximum 24 is below the null median (28; 90% range 24–34; $p = 0.97$). The exact assignment permutation leaves S invariant ($p = 1.00$), smooth Gamma/log-KDE models give $p = 0.18/0.12$ (strict 0.81/0.76), and real non-candidate controls give $p = 0.63\text{--}0.86$. The null is calibrated, but the power is low: a sharp family needs about 40 (strict 20) edges at the observed base or about 200 (strict 40) at high Λ for 95% power; with the measured 20% error spread about 200 (strict 100) are needed. A larger Λ dynamic range, a narrower relative window matched to the error model, or a known base would be needed for a decisive test.

References

- [1] S. W. Hawking, “Chronology protection conjecture,” *Phys. Rev. D* **46**, 603 (1992).
- [2] B. S. Kay, M. J. Radzikowski, and R. M. Wald, “Quantum field theory on space-times with a compactly generated Cauchy horizon,” *Commun. Math. Phys.* **183**, 533 (1997); arXiv:gr-qc/9603012.
- [3] G. Klinkhammer, “Vacuum polarization of scalar and spinor fields near closed null geodesics,” *Phys. Rev. D* **46**, 3388 (1992).
- [4] M. Visser, “The quantum physics of chronology protection,” in *The Future of Theoretical Physics and Cosmology* (Cambridge University Press, Cambridge, 2003); arXiv:gr-qc/0204022.
- [5] A. Ori, “A class of time-machine solutions with a compact vacuum core,” *Phys. Rev. Lett.* **95**, 021101 (2005); arXiv:gr-qc/0503077.
- [6] A. Ori, “Formation of closed timelike curves in a composite vacuum/dust asymptotically flat spacetime,” *Phys. Rev. D* **76**, 044002 (2007); arXiv:gr-qc/0701024.
- [7] S. Xavier, “Closed Timelike Curves from a Vacuum Traveling Wave,” arXiv:2607.00788 (2026).
- [8] C. W. Misner, “Taub-NUT space as a counterexample to almost anything,” in *Lectures in Applied Mathematics*, Vol. 8 (American Mathematical Society, Providence, 1967), pp. 160–169.
- [9] L.-X. Li and J. R. Gott III, “Self-consistent vacuum for Misner space and the chronology protection conjecture,” *Phys. Rev. Lett.* **80**, 2980 (1998); arXiv:gr-qc/9711074.
- [10] J. Dietz, A. Dirmeier, and M. Scherfner, “Geometric Analysis of Ori-type Spacetimes,” arXiv:1201.1929 (2012).

- [11] N. E. Rieger, “Causal classification of pathological Misner-type spacetimes,” *Annals of Physics* **489**, 170455 (2026); arXiv:2501.00841. doi:10.1016/j.aop.2026.170455.
- [12] F. Ahmed, A. Al-Badawi, and I. Sakallı, “Are Petrov type-N and D spacetimes admitting CTCs valid in $f(R, L_m, \Phi, X)$ gravity?,” *Eur. Phys. J. C* **86**, 792 (2026); arXiv:2605.26696. doi:10.1140/epjc/s10052-026-16044-w.
- [13] Gaia Collaboration, A. Vallenari *et al.*, “Gaia Data Release 3. Summary of the content and survey properties,” *Astron. Astrophys.* **674**, A1 (2023); arXiv:2208.00211.
- [14] Gaia Collaboration, “Gaia Data Release 3 documentation,” <https://www.cosmos.esa.int/web/gaia/dr3> (accessed 2026-10-26).
- [15] J. Schemoul, “A search for boost-holonomy relics in the gravitational-wave catalog,” companion GW survey, in preparation (2026).
- [16] The LIGO Scientific Collaboration, the Virgo Collaboration, and the KAGRA Collaboration, “GWTC-4.0: Searches for Gravitational-Wave Lensing Signatures,” arXiv:2512.16347 (2025).
- [17] The LIGO Scientific Collaboration, the Virgo Collaboration, and the KAGRA Collaboration, “GWTC-5.0: Methods for Identifying and Characterizing Gravitational-wave Transients,” arXiv:2605.27224 (2026).
- [18] X. Shan, B. Hu, X. Chen, and R.-G. Cai, “An interference-based method for the detection of strongly lensed gravitational waves,” *Nature Astronomy* **9**, 916 (2025); arXiv:2301.06117.
- [19] X. Shan, Y.-F. Wang, S. Mao, H. Yang, and M. Zumalacárregui, “Gravitational lensing of gravitational waves by galaxy clusters,” arXiv:2609.16891 (2026).
- [20] D. Hegde and J. Janquart, “On the interplay between waveform systematics and lensing signatures in gravitational-wave signals,” arXiv:2609.03935 (2026).
- [21] T. N. Nguyen-Vo, T. Singh, B. McKallip, M. Kesden, and L. King, “Distinguishing lensing and precessional modulation in binary black-hole inspiral waveforms,” arXiv:2609.06347 (2026).
- [22] U. Deka, A. Ganguly, S. Bhattacharyya, and P. Ajith, “Impact of the host galaxy on the search for gravitational-wave microlensing by compact objects,” arXiv:2609.16990 (2026).