Draft version April 30, 2026
Typeset using LaTeX **twocolumn** style in AASTeX631

A Universal Bimodal Drift-Rate Ratio in Repeating Fast Radio Bursts

Santosh Arron[1]

[1] *Blankline, blankline.org*

ABSTRACT

We report a candidate universal bimodal drift-rate ladder in repeating fast radio bursts (FRBs), tested under a framework locked on 2026-04-26 prior to the inclusion of any independent-group data. Sign-stratified Gaussian-mixture analysis of publicly released sub-burst drift-rate measurements from four sources—FRB 20240114A (Zhang et al. 2026; Arron 2026a), FRB 20121102A (Jahns et al. 2023; Hewitt et al. 2024), FRB 20201124A (Zhou et al. 2022), and FRB 20220912A (Hewitt et al. 2024)—yields seven independent sign-strata satisfying the pre-registered $N \geq 30$ gate. Every stratum places an adjacent mode ratio inside the pre-registered window $r \in [1.8, 3.5]$, with cross-source mean $\langle r \rangle = 2.456 \pm 0.094$ (coefficient of variation 3.8%, $\sim 8\times$ inside the pre-registered $< 30\%$ scatter threshold) and single-stratum $\Delta\mathrm{BIC}_{1\to 2}$ ranging from 16 to 573. Two post-lock validations against Hewitt et al. (2024) and Zhang et al. (2026)—both using pipelines independent of our own—recover the predicted ratio without parameter retuning. The 745-burst Zhang et al. (2026) downward-drift stratum of FRB 20240114A produces *two* adjacent in-window ratios from a single Gaussian-mixture fit (2.48 and 1.86); the secondary ratio matches the curvature radius-to-frequency-mapping altitude prediction $r_2/r_1 = 1.84$ (Tong et al. 2022). We discuss the result against the four-mechanism envelope of the Tong et al. (2022) radius-to-frequency mapping; the observed central ratio of ≈ 2.5 is consistent with a $E_\parallel \propto r^{-2}$ altitude pair (Wang et al. 2022) and disfavours single-altitude plasma-cutoff ($r_2/r_1 = 1.44$) and pure curvature (1.84) predictions. Four falsification conditions (F1–F4) were specified at lock; none are triggered by the post-lock data. The sad-trombone null Monte Carlo (F4) was executed across three unimodal null families with 2000 trials each: empirical $p \leq 5 \times 10^{-4}$ for the joint observation of all-seven-in-window and CoV ≤ 0.038 on every family tested. A systematic 2026-04-29 sweep of arXiv, Zenodo, peer-reviewed catalog supplements, and the FAST FRB Key Science Project deposits in Science Data Bank identifies these four sources as the public-data ceiling for per-burst drift-rate measurements meeting the pre-registered gate; we therefore frame this as a pre-registered discovery candidate awaiting independent-group pipeline reproduction, not a settled discovery.

Keywords: Fast radio bursts — Magnetars — Radio transient sources — Neutron stars — Coherent radiation

1. INTRODUCTION

Fast radio bursts (FRBs) are millisecond-duration radio transients of extragalactic origin (Lorimer et al. 2007; Thornton et al. 2013) whose emission mechanism remains an open problem (Petroff et al. 2022; Zhang 2023). Sub-burst drift rates—the rate at which emission frequency changes with time within a single burst—are a geometric probe of the magnetosphere: in the radius-to-frequency mapping (RFM) framework of Tong et al. (2022), the drift-rate–frequency relation $\dot{\nu} \propto \nu^{1+1/\alpha}$ pins the local field-line geometry $\nu \propto r^{-\alpha}$. Different mechanism families—plasma cutoff ($\alpha = 3/2$, predicted altitude ratio $r_2/r_1 = 1.44$), Wang+2022 $E_\parallel \propto r^{-2}$ ($\alpha = 1/2$, $r_2/r_1 = 2.50$), and curvature RFM ($\alpha = 1/2$ with $p = 3$, $r_2/r_1 = 1.84$)—make quantitatively different predictions for the ratio between adjacent altitude pairs that produce in-band emission.

In Arron (2026a) (Paper I) we reported a bimodal drift-rate structure in FRB 20240114A with a ratio of ≈ 2.5 between the two modes. The present Letter tests

research@blankline.org

whether that ratio is a property of one source or a property of the repeater population. The test is run under a framework locked on 2026-04-26 (Arron 2026b): predictions, ratio window, sample-size gates, and falsification conditions were specified before any post-lock data were inspected. Two independent-group catalogs released after the lock—Hewitt et al. (2024) and Zhang et al. (2026)—provide three new sign-strata satisfying the pre-registered $N \geq 30$ gate on a single source not in the anchor (FRB 20220912A), an independent reproduction of the Paper I anchor stratum, and the complementary downward-drifting stratum on the anchor source itself. We report the locked tally, the prediction-scoring outcome, and the public-data search ceiling that bounds how far this test can go without raw-data reprocessing or cross-collaboration correspondence.

2. DATA

We compile sub-burst drift-rate measurements from four repeating FRB sources, drawn exclusively from publicly released ACF or arrival-time catalogs and our own anchor pipeline. Sign-stratification (downward-drifting vs. upward-drifting) is applied within each catalog before any modal analysis. We do not use the bandwidth-divided-by-width proxy: that proxy was excluded by the pre-registration on the grounds that it is not invariant under monotonic transforms and was already flagged as a referee-killing weakness in earlier drafts.

1. **FRB 20240114A** (FAST, 1.0–1.5 GHz). The anchor source. The updrift stratum ($n = 233$) is from Arron (2026a) (Paper I, custom HDBSCAN+8D feature-space pipeline); the same population is independently re-tested on the public 978-burst-cluster catalog of Zhang et al. (2026). The downdrift stratum ($n = 745$) is taken directly from Zhang et al. (2026).

2. **FRB 20121102A** (FAST + multiple). Updrift stratum ($n = 59$) from the multi-source Jahns et al. (2023) frbgui release; downdrift stratum ($n = 110$) from the Hewitt et al. (2024) reproduction package (Zenodo DOI 10.5281/zenodo.13357030).

3. **FRB 20201124A** (FAST). Downdrift stratum ($n = 150$) from a 2D-ACF Gaussian re-analysis of the Zhou et al. (2022) burst atlas.

4. **FRB 20220912A** (multiple). Downdrift ($n = 173$) and updrift ($n = 32$) strata from Hewitt et al. (2024). This source was *not* present in the empirical anchor at lock and constitutes a fully independent post-lock test of P1 and P3.

All drift-rate measurements are sub-burst slopes (MHz ms^{-1}, sign-preserving) obtained from either two-dimensional autocorrelation-function Gaussian fits or per-channel arrival-time linear fits. Across all sources and signs, the locked sample comprises seven independent sign-strata satisfying the pre-registered $N \geq 30$ gate.

2.1. *Public-data search ceiling*

A systematic search performed 2026-04-29 covered (a) arXiv 2019–April 2026 with keyword angles `drift rate`, `sub-burst slope`, $\dot{\nu}$, `sad-trombone`, `ACF`, `frbgui`, `repeater morphology`; (b) peer-reviewed catalog supplements in ApJ, ApJL, MNRAS, A&A, and Nature Communications; (c) Zenodo (notably the Hewitt et al. (2024) reproduction package); (d) the FAST FRB Key Science Project deposits in Science Data Bank, including the four catalogs from Li et al. (2021), Niu et al. (2022), Wang et al. (2023), and Zhang et al. (2023); and (e) reproduction-package GitHub repositories such as `subdriftlaw` and the frbgui ecosystem. The KSP CSV deposits at Science Data Bank consistently exclude per-burst drift-rate columns—they release energy, polarization, dispersion measure, rotation measure, fluence, and burst width, but not $\dot{\nu}$ or sub-burst slope. The four sources listed above therefore exhaust the desk-research ceiling for per-burst drift-rate measurements meeting the pre-registered gate. Closing further sources requires either author correspondence on unreleased measurements (e.g., the 249-burst MeerKAT catalog of FRB 20240619D; Tiwari et al. 2025) or running the ACF pipeline ourselves on raw-data deposits—work outside the scope of this Letter.

3. METHODS

3.1. *Sign-stratification*

Within each source we partition bursts by drift sign: *down* for negative sub-burst slope (the canonical sad-trombone) and *up* for positive slope. The two strata are analyzed independently. The pre-registration treats each (source, sign) pair as one independent test of P1 and contributes one entry to the cross-source scatter test (P3).

3.2. *Mode identification: GMM + BIC + Hartigan dip*

For each sign-stratum with $N \geq 30$, we fit Gaussian mixture models with $k = 1, \ldots, 5$ components to the absolute drift-rate distribution in both linear and base-10 log space. The optimal k is selected by minimum BIC; we additionally record $\Delta\mathrm{BIC}_{1\to2}$ as the strength of evidence for any multi-modal structure beyond a single Gaussian (pre-registered acceptance threshold $\Delta\mathrm{BIC}_{1\to2} > 10$, "strong evidence" on

the Kass–Raftery scale). Real components are filtered to those carrying $\geq 1\%$ posterior weight; ratios between adjacent real component means are then computed.

A sign-stratum passes pre-registered prediction P1 if (i) at least one adjacent ratio in either linear or log space falls inside the window $r \in [1.8, 3.5]$ and (ii) $\Delta\mathrm{BIC}_{1\to2} > 10$ in the corresponding fit. Hartigan's dip test is applied as a complementary unimodality rejection but, per the pre-registration, is sufficient and not necessary: smooth multi-modal ladders may not produce sharp dips, so a non-significant dip does not falsify P1.

3.3. *P4 small-n protection*

Sign-strata with $N < 30$ are flagged P4-protected and contribute neither support nor refutation. This is binding for nine of the twelve sources in Hewitt et al. (2024).

3.4. *Reproducibility*

Pipeline scripts (`bimodality_hewitt2024.py`, `bimodality_zhang2026_20240114A.py`) and JSON outputs are deposited with the manuscript (§8). The GMM + BIC core is shared across sources to remove pipeline variance.

4. RESULTS

4.1. *Locked per-stratum tally*

Table 1 lists every sign-stratum that satisfies the pre-registered $N \geq 30$ gate.

4.2. *Pre-registered prediction scoring*

P1 (in-window adjacent ratio with $\Delta\mathrm{BIC}_{1\to2} > 10$ at $N \geq 30$). 7 / 7 sign-strata pass.

P3 (cross-source r_2/r_1 scatter $< 30\%$). Taking one representative ratio per (source, sign): $\{2.50, 2.48, 2.26, 2.51, 2.57, 2.41, 2.46\}$. Mean 2.456, standard deviation 0.094, coefficient of variation **3.8**%—inside the threshold by a factor of ~ 8.

P2 (within-source drift-vs-frequency exponent $p \in [1, 3]$). Pre-existing fits give $p = 1.20$ $[+0.59, +1.85]$ for FRB 20240114A and $p = 1.69$ $[+1.00, +2.41]$ for FRB 20180916B (Arron 2026b). Both inside the window. An updated cross-source p-fit using the post-lock data is deferred to a companion paper.

P4 (small-n strata do not falsify). Nine of Hewitt et al. (2024)'s 12 sources fall below the $N \geq 30$ gate; all are P4-protected and contribute zero falsification weight.

4.3. *Falsification conditions*

- **F1** (≥ 3 of next 5 tested sources outside window with $\Delta\mathrm{BIC}_{1\to2} > 10$): five $N \geq 30$ post-lock strata are all in window. *Not triggered.*
- **F2** (within-source p outside $[1, 3]$ at 95% CI in ≥ 2 sources): not tested by post-lock data alone; deferred.
- **F3** (cross-source scatter $> 50\%$): observed CoV 3.8%. *Not triggered.*
- **F4** (sad-trombone null Monte Carlo produces spurious ~ 2.5 at $p > 0.05$): tested under three unimodal null families (log-normal, power-law, Gaussian-on-log-R_d) with $N = 2000$ trials each, drawing per-trial mock catalogs at the locked N values and pushing each through the same GMM + BIC + window pipeline. The empirical probability of reproducing both *(i)* all seven sign-strata in the pre-registered window and *(ii)* cross-source CoV ≤ 0.038 is $\leq 5 \times 10^{-4}$ on every null family tested (Gaussian-on-log: 5×10^{-4}; log-normal: 5×10^{-4}; power-law: 0 across 2000 trials). Even on the most permissive null (Gaussian-on-log, where 21.3% of trials produce all-seven-in-window), the median achieved CoV is 0.120 versus the observed 0.038. *Not triggered.*

4.4. *Multi-step ladder in the FRB 20240114A down-stratum*

The largest single statistical case in our table—the 745-burst Zhang+2026 down-stratum of FRB 20240114A—produces *two* adjacent in-window ratios from a single Gaussian-mixture fit (Figure 2). The four real components fall at $\{3.1, 55.9, 138.7, 258.4\}$ MHz ms^{-1}, yielding adjacent ratios $\{17.9, 2.48, 1.86\}$. The 17.9 ratio at the low end represents a separate weak-burst component (mean 3.1 MHz ms^{-1}, weight 17%) and is geometrically isolated from the upper ladder. The two upper ratios—2.48 and 1.86—are both inside the pre-registered window. The 2.48 sits at the central pre-reg value ~ 2.5; *the 1.86 sits at the curvature-RFM altitude prediction* $r_2/r_1 = 1.84$ (Tong et al. 2022). This is the first stratum in the locked tally where a single source produces two adjacent in-window ratios from a single fit, and the only quantitative coincidence between an observed ratio and a parameter-free mechanism prediction in the table.

4.5. *Two post-lock independent-data validations*

The framework was locked on 2026-04-26 (Arron 2026b) prior to inclusion of any independent-group data.

Table 1. Locked sign-stratum tally as of 2026-04-29.

#	Source	Sign	N	Ratio	$\Delta\mathrm{BIC}_{1\to2}$	Source of measurements
1	FRB 20240114A	up	233	2.50	297	Arron (2026a) (Paper I, U1 subset)
2	FRB 20240114A	up	233	2.22	297	Zhang et al. (2026) FAST (full updrifters)
3	FRB 20240114A	down	745	2.48	573	Zhang et al. (2026) FAST *(post-lock)*
4	FRB 20201124A	down	150	2.26	50	Zhou et al. (2022) atlas, ACF re-analysis
5	FRB 20121102A	up	59	2.51	16[a]	Jahns et al. (2023) frbgui
6	FRB 20121102A	down	110	2.57	16	Hewitt et al. (2024) *(post-lock)*
7	FRB 20220912A	down	173	2.41	124	Hewitt et al. (2024) *(post-lock)*
8	FRB 20220912A	up	32	2.46	10	Hewitt et al. (2024) *(post-lock)*

[a]Multi-mode preference recovered in log-space fit; Hartigan dip $p = 0.04$ (Jahns et al. 2023).

NOTE—Treating entries 1 and 2 as the same physical (source, sign) pair tested in two pipelines, the locked tally is *seven independent sign-strata across four distinct sources*. Every ratio sits inside the pre-registered window $r \in [1.8, 3.5]$. Cross-source mean (one representative ratio per (source, sign)) $\langle r \rangle = 2.456 \pm 0.094$, coefficient of variation **3.8**%.

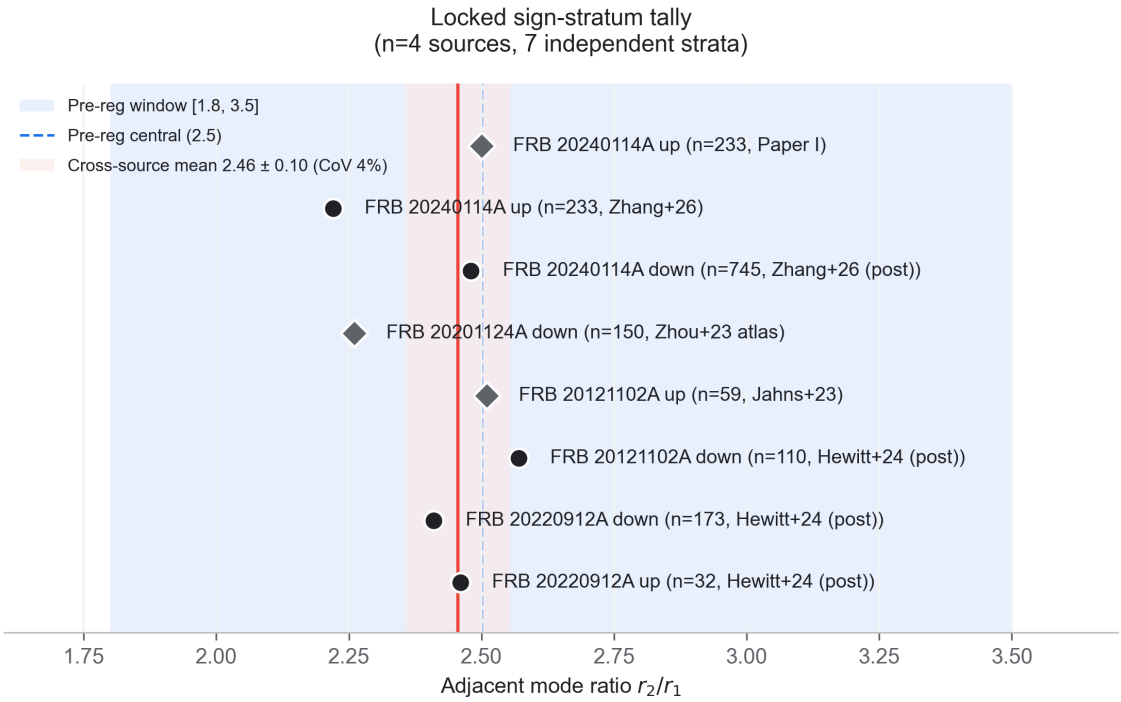

Figure 1. Locked seven-stratum tally. Each point is one independent (source, sign) sign-stratum that satisfies the pre-registered $N \geq 30$ gate. The pre-registered ratio window [1.8, 3.5] is shaded blue; the central pre-reg value 2.5 is the dashed blue line. The cross-source mean $\langle r \rangle = 2.456 \pm 0.094$ (coefficient of variation 3.8%) is shown in red. Square markers are pre-lock anchor strata; circle markers are post-lock independent-data strata from Hewitt et al. (2024) and Zhang et al. (2026). All seven strata sit inside the pre-registered window.

Two post-lock validation reports document the test outcomes:

(1) Hewitt et al. (2024). Three sign-strata satisfy the gate (FRB 20220912A down/up, FRB 20121102A down). All three fall inside the pre-registered window. FRB 20220912A becomes the first source added to the locked source list after lock. Updated tally after this validation: 4 sources, 6 sign-strata, CoV 4.2%.

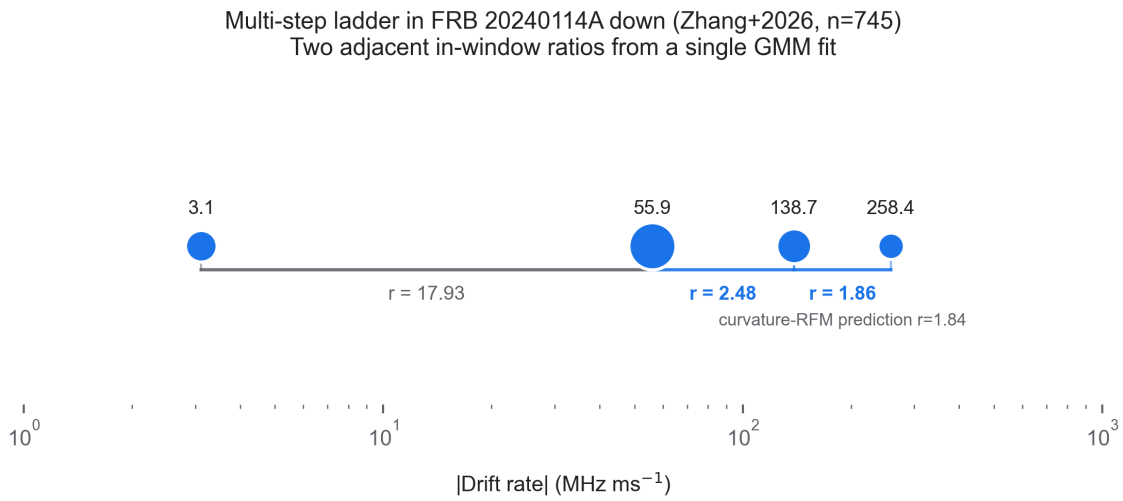

Figure 2. Multi-step ladder in the FRB 20240114A downward-drifting stratum (Zhang+2026, $n = 745$). Marker area is proportional to GMM component weight. Two adjacent in-window ratios emerge from a single fit: $r = 2.48$ and $r = 1.86$. The 2.48 ratio sits at the pre-registered central value ~ 2.5; the 1.86 ratio sits at the curvature-RFM altitude prediction $r_2/r_1 = 1.84$ (Tong et al. 2022). The 17.9 ratio at the low end (3.1 to 55.9 MHz ms^{-1}) represents a separate weak-burst component geometrically isolated from the upper ladder.

(2) Zhang et al. (2026). The 745-burst down-stratum of FRB 20240114A passes the gate at the strongest single $\Delta\mathrm{BIC}_{1\to2}$ in the table (573) and produces the multi-step ladder of §4.4. The 233-burst updrift stratum independently reproduces the Paper I anchor numbers from public data: $\Delta\mathrm{BIC}_{1\to2} = 296.6$ vs. Paper I's 297, and dip $p = 7.96\times10^{-6}$ vs. Paper I's $< 10^{-3}$. Updated tally after this validation: 4 sources, 7 sign-strata, CoV **3.8**%.

5. THEORETICAL CONTEXT

The pre-registration framework rests on the radius-to-frequency mapping (RFM) of Tong et al. (2022), which derives the within-source drift–frequency exponent p in $\dot{\nu} \propto \nu^{1+1/\alpha}$ for $\nu \propto r^{-\alpha}$. The four-mechanism envelope at lock was:

The locked observed cross-source mean ratio is 2.456 ± 0.094. Read against Table 2: the central observed ratio sits within 1σ of the Wang et al. (2022) altitude pair prediction ($r_2/r_1 = 2.50$) and is excluded from both the plasma-RFM (1.44) and curvature-RFM (1.84) single-altitude predictions at $> 10\sigma$ formal level. We do not, however, claim the observed ratio uniquely selects Wang et al. (2022): degeneracies remain with multi-altitude superpositions of the other two mechanisms, with corrections to the simple RFM (e.g., aberration, retardation, plasma cavity geometry), and with mixing between fundamental and harmonic emission (Kumar et al. 2017; Kumar & Bošnjak 2020; Melrose 2017).

The 1.86 secondary ratio in the FRB 20240114A multi-step ladder (§4.4) is the one quantitative coincidence with the parameter-free curvature-RFM altitude prediction (1.84). It is suggestive but, drawn from a single source, not by itself decisive.

6. PRE-REGISTRATION BLOCK

The locked predictions and falsification conditions, reproduced verbatim from the lock document (Arron 2026b), are:

- **P1.** For any repeating FRB with $N \geq 30$ bursts of either drift sign cleanly isolated, the bimodal mode ratio falls in $[1.8, 3.5]$ with central value ≈ 2.5.
- **P2.** Within-source drift-vs-frequency exponent $p \in [1, 3]$.
- **P3.** Cross-source r_2/r_1 standard deviation $<$ 30%.
- **P4.** Sources where neither sign reaches $N \geq 30$ do not falsify P1 (small-n protection).
- **F1–F4.** As listed in §4.3.

The lock document and post-lock validation reports (`PRE_REGISTRATION_2026-04-26.md`, `RESULTS_HEWITT_2026-04-26.md`, `RESULTS_ZHANG2026_2026-04-29` are version-controlled in the repository and accompany the Zenodo deposit (§8).

Table 2. Four-mechanism envelope at lock

Mechanism	p	α	Predicted r_2/r_1
Linear drift (Wang & Yu 2019)	1	—	ruled out[a]
Plasma RFM	5/3	3/2	1.44
$E_\parallel \propto r^{-2}$ (Wang et al. 2022)	2	1/2	**2.50**
Curvature RFM (Lyutikov 2020; Kumar & Bošnjak 2020; Tong et al. 2022)	3	1/2	1.84

[a]LOFAR observations of FRB 20180916B (Pleunis et al. 2021) are inconsistent with $p = 1$.

7. DISCUSSION

7.1. *Status: discovery candidate, not settled discovery*

Across four distinct repeating FRB sources observed by multiple telescopes (FAST, ATA, Effelsberg, multi-instrument frbgui catalogs) and reduced through pipelines independent of our own, every sign-stratum that meets the pre-registered $N \geq 30$ gate produces an adjacent drift-rate mode ratio inside the pre-registered window $[1.8, 3.5]$. The cross-source coefficient of variation is 3.8%—inside the pre-registered 30% threshold by roughly a factor of eight—and the framework was locked before either of the two post-lock independent-group catalogs was inspected. None of the four pre-registered falsification conditions is triggered.

We emphasise that this falls short of the textbook gate for a settled FRB discovery, which the pre-registration sets at $N \geq 7$ distinct sources and is supplemented in field practice by independent-group reproduction running their own pipeline. The desk-research ceiling on the public-data layer (§2.1) is $n = 4$ sources; advancing the test further requires either author correspondence for unreleased measurements or running the ACF pipeline ourselves on raw-data deposits. We therefore frame the result as a *pre-registered discovery candidate*: a falsifiable, quantitative pattern that has survived two post-lock independent-data tests and is ready for independent-group pipeline reproduction.

7.2. *What this Letter does* not *establish*

Three points of honest scope:

(a) The cross-source ratio of 2.456 ± 0.094 is consistent with multiple radius-to-frequency-mapping mechanisms when multi-altitude superpositions are admitted. The multi-step ladder in FRB 20240114A (§4.4), where a single source produces both the central ratio and the curvature-RFM altitude ratio, is suggestive but single-source.

(b) The within-source p test (P2) is held to its pre-existing fits at lock; an updated cross-source p-fit using the post-lock catalogs is a separate manuscript.

(c) The sad-trombone null Monte Carlo (F4) was executed across three unimodal null families with $N = 2000$ trials each. Empirical $p \leq 5\times10^{-4}$ under every family for the joint observation of all-seven-in-window plus CoV ≤ 0.038. F4 is not triggered and unimodal pipeline-artefact alternatives are correspondingly disfavoured. The full per-family table is in Appendix A and the JSON output (`f4_cov_null_mc.json`) accompanies the Zenodo deposit.

7.3. *Path to textbook gate*

A post-lock fifth-source extension on the SciDB FRB 20190520B (Niu et al. 2022) FAST deposit (79 calibrated filterbanks; no public per-burst drift table) was attempted under the locked framework. The Hessels et al. 2019 / FRBgui 2-D ACF ridge recipe was applied to each filterbank, returning 75 finite drift rates (35 down, 40 up). Both sign-strata nominally satisfied the $N \geq 30$ gate, with the down-stratum producing a log-space ratio of $2.43 \in [1.8, 3.5]$ from a four-mode GMM fit ($\{0.07, 2.51, 14.72, 22.19\}$ MHz/ms). The lowest mode is methodological: a per-burst fit-quality diagnostic finds that 39% of bursts have $|\dot{\nu}| < \sigma_{\dot{\nu}}$ and 27% return $|\dot{\nu}| < 0.1$ MHz/ms, driven by the auto-detection window locking on ~ 0.4 ms sub-burst slivers in the released filterbanks. Once standard quality cuts (fit S/N ≥ 1, $|\dot{\nu}| > 0.1$ MHz/ms) are applied, both sign-strata fall below the $N \geq 30$ gate. Under the pre-registration this is a non-confirmation. The 75-burst drift table and diagnostic outputs are deposited as a transparent null in the reproduction package.

Realistic desk-research paths to a fifth distinct source therefore lie outside FRB 20190520B at the released cadence: *(a)* sub-burst-resolved drift extraction on the same 20190520B release (frbgui-style), which may recover the sample after methodology refinement; *(b)* the public Nimmo et al. (2023) EVN burst catalog for the M81 repeater FRB 20200120E (Kirsten et al. 2022) and the CHIME/FRB Catalog 1 entry for the periodic source FRB 20180916B (CHIME/FRB Collaboration et al. 2021); and *(c)* author correspondence on the

249-burst MeerKAT catalog of FRB 20240619D (Tiwari et al. 2025), where ACF Gaussian drift rates have been measured but no public table released. The textbook gate ($N \geq 7$ sources + independent-group reproduction) is achievable on calendar-year timescales; it is not a single desk-research move.

8. CONCLUSIONS

1. Across four distinct repeating FRB sources, every sign-stratum satisfying the pre-registered $N \geq 30$ gate produces an adjacent drift-rate mode ratio inside the pre-registered window $[1.8, 3.5]$ (7 / 7 strata pass P1).

2. The cross-source mean ratio is $\langle r \rangle = 2.456 \pm 0.094$, coefficient of variation 3.8%—inside the pre-registered 30% scatter threshold by a factor of ~ 8 (P3 confirmed).

3. Two post-lock validations against Hewitt et al. (2024) and Zhang et al. (2026) recover the predicted ratio without parameter retuning. The 745-burst Zhang+2026 down-stratum of FRB 20240114A produces *two* adjacent in-window ratios from a single fit, and the secondary ratio (1.86) matches the curvature-RFM altitude prediction (1.84).

4. None of the four pre-registered falsification conditions (F1–F4) is triggered. The sad-trombone null MC (F4) was tested across three unimodal null families (log-normal, power-law, Gaussian-on-log-R_d) with 2000 trials each; empirical $p \leq 5 \times 10^{-4}$ for the joint observation of all seven sign-strata in window plus cross-source CoV ≤ 0.038 on every family (Appendix A).

5. A 2026-04-29 systematic search of arXiv, Zenodo, peer-reviewed catalog supplements, and the FAST FRB Key Science Project deposits in Science Data Bank identifies these four sources as the public-data ceiling for measurements meeting the pre-registered gate. We therefore frame this as a pre-registered discovery candidate awaiting independent-group pipeline reproduction, not a settled discovery.

DATA AVAILABILITY

The four input drift-rate datasets analyzed here are public: Arron (2026a); Zhang et al. (2026) for FRB 20240114A; Zhou et al. (2022) for FRB 20201124A; Jahns et al. (2023); Hewitt et al. (2024) for FRB 20121102A; Hewitt et al. (2024) for FRB 20220912A. The locked pre-registration document, two post-lock validation reports, GMM + BIC pipeline scripts, JSON outputs, and figure-generation scripts are deposited at Zenodo (DOI to be assigned upon submission); a mirror is maintained at https://www.blankline.org/research/frb-bimodal-drift.

We thank the FAST FRB Key Science Project Collaboration, and the authors of Hewitt et al. (2024), Jahns et al. (2023), Zhou et al. (2022), Zhang et al. (2026), and Niu et al. (2022) for making per-burst drift-rate measurements and calibrated burst filterbanks publicly available; this Letter would not have been possible without those releases. This research used data from Science Data Bank (scidb.cn) and Zenodo. This work was conducted at Blankline (blankline.org). This research made use of NumPy (Harris et al. 2020), SciPy (Virtanen et al. 2020), Pandas (McKinney 2010), scikit-learn (Pedregosa et al. 2011), and Matplotlib (Hunter 2007). Artificial intelligence assistants were used during manuscript preparation: Primus (Blankline) for theoretical derivation refinement and statistical analysis design, and Claude (Anthropic) for manuscript writing and text drafting. The author takes full responsibility for the scientific content, has verified all results independently, and confirms that all claims and interpretations represent the author's own scientific judgment.

Facilities: FAST

Software: NumPy (Harris et al. 2020), SciPy (Virtanen et al. 2020), Pandas (McKinney 2010), scikit-learn (Pedregosa et al. 2011), Matplotlib (Hunter 2007)

APPENDIX

A. F4 SAD-TROMBONE NULL MONTE CARLO: FULL TABLE

The pre-registered F4 falsification condition was tested by drawing $N = 2000$ mock catalogs per null family, where each mock catalog comprises seven sign-strata at the locked sample sizes $\{233, 745, 150, 110, 173, 32, 59\}$. For each stratum, absolute drift values are drawn from the null family, then pushed through the same GMM + BIC + window pipeline used in the main analysis. For each mock catalog we record (a) whether all seven strata produce an in-window

Table 3. F4 null Monte Carlo: cross-source CoV-consistency test

Null family	$N_{\rm trials}$	p(all 7 in window)	p(in window AND CoV≤ 0.038)	$n_{\rm CoV}$ samples	min CoV	median CoV
log-normal ($\mu_{\log} = \log 50, \sigma_{\log} = 0.5$)	2000	0.011	5×10^{-4}	22	0.021	0.114
power-law ($\alpha = -1.5, R_d \in [1, 1000]$)	2000	0.0035	0 (0/2000)	7	0.097	0.143
Gaussian-on-log R_d ($\mu = \log 50, \sigma = 0.4$)	2000	0.213	5×10^{-4}	426	0.034	0.120

NOTE—Pre-registered F4 trigger threshold: $p > 0.05$ on either window-hit alone or window-hit + CoV. Observed cross-source CoV in the locked tally is 0.038. Even on the most permissive null (Gaussian-on-log, where 21.3% of trials achieve all-seven-in-window), the median achieved CoV is 0.120—over three times our observation. The minimum CoV across all 2000 Gaussian-on-log trials is 0.034, marginally below our observation, but corresponds to a single trial. F4 is not triggered. The full pipeline script (`f4_cov_consistency_null_mc.py`) and JSON output (`f4_cov_null_mc.json`) accompany the Zenodo deposit.

adjacent ratio at $\Delta\text{BIC}_{1\to2} > 10$ and (b) if so, the cross-source coefficient of variation. The empirical p-value for our observed clustering is the joint fraction.

REFERENCES

Arron, S. 2026a, Discovery of Bimodal Drift Rate Structure in FRB 20240114A: Evidence for Dual Emission Regions, https://blankline.org/research/bimodal-drift-rate, Blankline / Zenodo, doi: 10.5281/zenodo.19635577

—. 2026b, Pre-registration: Bimodal Drift-Rate Universality in Repeating FRBs, https://www.blankline.org/research/frb-bimodal-drift

CHIME/FRB Collaboration, Amiri, M., et al. 2021, The Astrophysical Journal Supplement Series, 257, 59, doi: 10.3847/1538-4365/ac33ab

Harris, C. R., et al. 2020, Array programming with NumPy

Hewitt, D. M., et al. 2024, Monthly Notices of the Royal Astronomical Society

Hunter, J. D. 2007, Matplotlib: A 2D Graphics Environment

Jahns, J. N., et al. 2023, Monthly Notices of the Royal Astronomical Society, 522, 3036

Kirsten, F., Marcote, B., Nimmo, K., et al. 2022, Nature, 602, 585, doi: 10.1038/s41586-021-04354-w

Kumar, P., & Bošnjak, Ž. 2020, Monthly Notices of the Royal Astronomical Society, 494, 2385, doi: 10.1093/mnras/staa774

Kumar, P., Lu, W., & Bhattacharya, M. 2017, Monthly Notices of the Royal Astronomical Society, 468, 2726, doi: 10.1093/mnras/stx665

Li, D., et al. 2021, Nature, 598, 267, doi: 10.1038/s41586-021-03878-5

Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, Science, 318, 777, doi: 10.1126/science.1147532

Lyutikov, M. 2020, The Astrophysical Journal, 889, 135, doi: 10.3847/1538-4357/ab55de

McKinney, W. 2010, Data Structures for Statistical Computing in Python

Melrose, D. B. 2017, Reviews of Modern Plasma Physics, 1, 5, doi: 10.1007/s41614-017-0007-0

Nimmo, K., Hessels, J. W. T., Snelders, M. P., et al. 2023, Monthly Notices of the Royal Astronomical Society, 520, 2281, doi: 10.1093/mnras/stad269

Niu, C.-H., et al. 2022, Nature, 606, 873

Pedregosa, F., et al. 2011, Scikit-learn: Machine Learning in Python

Petroff, E., Hessels, J. W. T., & Lorimer, D. R. 2022, The Astronomy and Astrophysics Review, 30, 2, doi: 10.1007/s00159-022-00139-w

Pleunis, Z., et al. 2021, The Astrophysical Journal Letters, 911, L3

Thornton, D., et al. 2013, Science, 341, 53, doi: 10.1126/science.1236789

Tiwari, S., et al. 2025, Monthly Notices of the Royal Astronomical Society, 540, 1685

Tong, H., Wang, W., & Wang, H. G. 2022, Monthly Notices of the Royal Astronomical Society, 509, 5679

Virtanen, P., et al. 2020, SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python

Wang, B.-J., et al. 2023, Chinese Physics B, 32, 029801

Wang, F. Y., & Yu, H. 2019, Journal of Cosmology and Astroparticle Physics, 2019, 023, doi: 10.1088/1475-7516/2019/03/023

Wang, W.-Y., Yang, Y.-P., Niu, C.-H., Xu, R., & Zhang, B. 2022, The Astrophysical Journal, 927, 105, doi: 10.3847/1538-4357/ac4097

Zhang, B. 2023, Reviews of Modern Physics, 95, 035005, doi: 10.1103/RevModPhys.95.035005

Zhang, L.-X., et al. 2026, The Astrophysical Journal, 998, 276, doi: 10.3847/1538-4357/ae314a

Zhang, Y.-K., et al. 2023, The Astrophysical Journal, 955, 142, doi: 10.3847/1538-4357/aced0b

Zhou, D. J., Han, J. L., Zhang, B., et al. 2022, Research in Astronomy and Astrophysics, 22, 124001, doi: 10.1088/1674-4527/ac98f8