

Redundant Record Consensus in Planck PR3 CMB Component Maps

A T16 empirical demonstration: an aggregation diagnostic on public component-separated maps, re-executed end-to-end from the primary archives at conversion

Jeremy C. Jones

HoldingLight LLC

ORCID: 0009-0007-2515-3774

contact@universalcollapse.com

v1.0 · 2026-07-15

Abstract

Component-separated cosmic microwave background (CMB) maps are reconstructed from the same underlying multi-frequency sky observations by different algorithms. Planck Public Release 3 provides four widely used CMB temperature products — Commander, NILC, SEVEM, and SMICA — including half-mission splits and full-mission maps. This paper applies a minimal, operational record-consensus diagnostic: treat each published map variant as a redundant record channel of the same underlying sky signal under the shared PR3 common intensity mask, and measure how quickly disagreement shrinks when channels are averaged. Using masked-pixel sampling at $N_{\text{side}} = 2048$ (seeded; $n = 10^6$), we compute pairwise correlations and RMS discrepancies among fragments and a convergence curve $D(k)$ — the RMS residual of k -map subset averages against the full-set mean. Two fragment sets are analyzed: eight half-mission fragments (primary) and the four full-mission products (robustness).

Results: across the eight half-mission fragments, the mean off-diagonal correlation is $\bar{r} = 0.9756756$; across the four full-mission products, $\bar{r} = 0.9951544$. $D(k)$ decreases rapidly with k and is stable across sampling size (a 50,000-pixel smoke run reproduces $\bar{r}$ to $|\Delta\bar{r}| \approx 1.0 \times 10^{-5}$) and random seeds (span 2.0×10^{-4} across three seeds at $n = 250,000$). The interpretive boundary is stated up front and carried throughout: this is a **record-consensus demonstration, not an independence audit**. Because the reference in $D(k)$ is the full-set mean and the fragments are highly correlated, the descriptive effective-channel count $k_{\text{eff}} = k / (1 + (k - 1)\bar{r})$ evaluates to ≈ 1.02 for the eight half-mission fragments and ≈ 1.00 for the four full-mission products: the ensemble behaves as approximately one effective independent channel, and the observed convergence quantifies consensus among highly redundant records rather than the accumulation of independent evidence. The audited-independence upgrade is specified in the companion methods note (Jones 2026a).

At T16 conversion (2026-07), the robustness configuration was re-executed end-to-end from the primary public archives: the four full-mission maps and the common mask (≈ 7 GB) were re-downloaded from IRSA, the pixel sample was regenerated with the package's own seeded sampling function, and the analysis was re-run. The published values reproduce to every reported digit — $\bar{r} = 0.9951544$; off-diagonal correlation range 0.9924926–0.9979500; RMS range $6.8637 \times 10^{-6} - 1.3182 \times 10^{-5}$ K_CMB (6.86–13.18 μK) — establishing that the result is recoverable in a fresh environment from public materials and the deposited package alone.

Keywords: cosmic microwave background; Planck; component separation; records; consensus; redundancy; aggregation; robustness; replication; reproducibility.

1. Introduction

Modern CMB analyses begin from time-ordered detector data and produce calibrated frequency maps; the final CMB-only temperature map is then estimated by separating the CMB signal from Galactic and extragalactic foregrounds and residual instrumental systematics. Because component separation is an inverse problem, multiple algorithms are developed, compared, and released side by side. Planck PR3 provides four prominent pipelines — Commander, NILC, SEVEM, and SMICA — each producing a full-sky CMB map and, for several products, internal half-mission splits (Planck Collaboration 2020). The presence of multiple public products creates the opportunity for a simple but informative check: to what extent do independently produced reconstructions converge, and how does convergence improve when they are aggregated?

This paper offers a compact diagnostic aimed at readability and auditability rather than novelty in cosmological parameter inference. We operationalize record consensus as cross-product agreement under a common mask and measure how disagreement decreases as additional record channels are averaged. Strong convergence across algorithmically diverse reconstructions is a useful sanity check for downstream analyses; systematic divergence would localize failure modes — mask dependence, residual foreground leakage, pipeline-specific artifacts — and motivate targeted investigation.

UCT positioning (scope note). Within Universal Collapse Theory, cross-channel convergence is an instance of redundant-record stabilization under shared constraints: the standards layer treats records as the durable products of constraint-guided resolution (Jones 2026d), and TN-S₁ derives the exponential consensus bound that motivates aggregation diagnostics of this kind (Jones 2026c). The present analysis is fully interpretable without that framework, as a straightforward reproducibility and robustness check on released CMB maps. One boundary is stated at the outset and carried throughout: this is a consensus demonstration among highly correlated records, **not** an independence audit — the distinction is quantified in §5 and the audited upgrade is specified in the companion methods note (Jones 2026a).

This paper joins the T16 empirical corpus as a physics-wing entry alongside Entropy as Record, Rice Hysteresis, COGITATE iEEG, and the AI-substrate demonstration (Jones 2026b).

2. Pre-Specified Design and Registration Mechanism

2.1 Research questions, metrics, and models

The locked protocol pre-specified three research questions: **RQ1**, do Planck PR3 component-separated CMB temperature maps exhibit redundant-record consensus under aggregation; **RQ2**, what is the effective fragment count k_{eff} given inter-fragment correlation; **RQ3**, is disagreement reduction dominated by shared structure and systematics (high correlation) or by independent residuals. Metrics were fixed in advance: pairwise RMS disagreement $D_{ij} = \text{RMS}(m_i - m_j)$; pairwise correlation corr_{ij} ; the summary redundancy metric $\bar{r}$ (mean off-diagonal correlation); the aggregation curve $D(k) = E_S[\text{RMS}(m_S - m_{\text{all}})]$ over random size- k subsets against the full-set mean; and the descriptive effective count $k_{\text{eff}} = k / (1 + (k - 1)\bar{r})$, used as an interpretive axis rather than a physical model. An exponential-with-floor curve summary is fitted to $D(k)$ as a descriptive shape aid only. No conclusion in

this paper depends on that fit; the licensing quantities are $\bar{r}$, the pairwise tables, $D(k)$, and the reproducibility property, with k_{eff} carried as descriptive (§5).

2.2 Registration mechanism, stated plainly

Registration for this study is the weakest in the T16 family to date, and the paper says so rather than dressing it. The analysis plan — locked data products, fixed preprocessing, fixed metrics, no-post-hoc-substitution rule — is carried in the versioned reproducibility package and in the protocol section of the internal Registered Analysis Report ("registered" in the project's locked-source sense, not a third-party filing). An OSF preregistration template was drafted but not filed; the package's deviations log exists and contains zero entries. What this study offers in place of a stronger registration lock is a stronger reproduction property: the entire analysis is recomputable by any party from public archives using the package's own scripts and seeds — and §6 reports exactly that re-execution, performed at conversion in a fresh environment, reproducing every published digit. For this diagnostic class — public inputs, deterministic seeded sampling, simple fixed metrics — end-to-end public reproducibility is the operative guarantee.

3. Data Products, Fragments, and Methods

3.1 Data products and fragment sets

All inputs are Planck PR3 public products in HEALPix format at $N_{\text{side}} = 2048$ (Górski et al. 2005): the Stokes-I temperature field (I_{STOKES}) from each component-separated map, under the PR3 common intensity mask ($\text{TMASK} > 0$), retrieved from the IRSA Planck archive. Two fragment sets are analyzed. The **half-mission (HM) primary set** ($n_{\text{maps}} = 8$) comprises Commander, NILC, SEVEM, and SMICA $\times$ HM1/HM2 — within-pipeline splits over disjoint time segments. The **core-4 robustness set** ($n_{\text{maps}} = 4$) comprises the four full-mission products — cross-algorithm comparison on the full mission data. The two sets provide complementary views: within-method splits versus cross-method reconstructions.

3.2 Sampling, preprocessing, and computation

A fixed-size seeded sample of masked-pixel draws (default $n_{\text{sample}} = 10^6$; seed 7) is generated by batched rejection against the mask. Sampling is with replacement, so a small proportion of pixel indices can recur; at seed 7, 12,366 of the one million draws (1.24%) are repeated indices, as is directly checkable from the deposited index vector. The identical generated index vector is applied to every map. Maps are used exactly as released — thermodynamic CMB temperature units (K_{CMB}) and each product's native beam and smoothing, with no re-smoothing or beam harmonization applied. Each map's sampled Stokes-I vector is mean-centered (removal of the sample monopole; no dipole treatment beyond what the released products carry), and masked pixels are excluded by construction: only $\text{TMASK} > 0$ pixels enter the sample. Pairwise correlations use the Pearson coefficient on the centered vectors; $D(k)$ is estimated by Monte Carlo over 800 random subsets per k with a fresh seeded generator. The masked-in fraction of the common mask is 77.9% of the 50,331,648 pixels at this resolution. Because the sampling generator is deterministic given the seed and the mask, the exact pixel sample — and therefore every downstream number — is recomputable by construction; §6 exploits this.

4. Results

4.1 Completed runs

Run (output folder)	n_maps	n_sample	seed	$\bar{r}$	Role
A — HM default (outputs/)	8	1,000,000	7	0.9756756	primary
A' — HM smoke (outputs_smoke/)	8	50,000	7	0.9756654	stability
B — Core-4 full (outputs_core4/)	4	1,000,000	7	0.9951544	robustness
B (patched fit routine)	4	1,000,000	7	0.9951544	identical (§7)

Table 1. Completed runs of the reproducibility package (2026-02-18). The patched run differs only in a numeric-stabilization fix to the descriptive fit routine; all primary consensus metrics are byte-identical (§7).

Off-diagonal agreement ranges: **A** — correlation $0.9507176\text{--}0.9965699$, RMS $8.9534 \times 10^{-6} - 3.4284 \times 10^{-5}$; **A'** — correlation $0.9504169\text{--}0.9965646$, RMS $8.9885 \times 10^{-6} - 3.4427 \times 10^{-5}$; **B** — correlation $0.9924926\text{--}0.9979500$, RMS $6.8637 \times 10^{-6} - 1.3182 \times 10^{-5}$. RMS values are in K_CMB; equivalently, A: $8.95\text{--}34.28\ \mu\text{K}$, A': $8.99\text{--}34.43\ \mu\text{K}$, B: $6.86\text{--}13.18\ \mu\text{K}$. The smoke run reproduces the primary $\bar{r}$ to $|\Delta\bar{r}| \approx 1.0 \times 10^{-5}$ at one-twentieth the sample, and seed sensitivity at $n = 250,000$ (seeds 7, 13, 101) spans only 2.0×10^{-4} in $\bar{r}$. The half-mission fragments agree strongly; the four full-mission products agree more tightly still, consistent with full-mission maps being less noisy and more mutually aligned.

4.2 The core-4 correlation structure

The full pairwise matrix for the robustness set — which the original documents reported only as ranges — is shown in Figure 1, from the conversion re-execution of §6 (whose range endpoints match the archived values exactly). The weakest pair is Commander–NILC (0.99249) and the tightest NILC–SMICA (0.99795), consistent with Commander's parametric pixel-based approach sitting slightly apart from the ILC-family reconstructions. All pairs exceed 0.992 under the common mask.




Figure 1. Pairwise Pearson correlation among the four full-mission products on the shared 10^6 -pixel masked sample (conversion re-execution, §6). Range endpoints match the archived run exactly.

4.3 Convergence under aggregation

The consensus curve $D(k)$ declines rapidly with k in both fragment sets and is stable across sample size and seed. Figure 2 shows the core-4 curve from the conversion re-execution: subset means converge on the full-set mean from $\approx 6.2 \mu\text{K}$ RMS at $k = 1$ to $\approx 2.1 \mu\text{K}$ at $k = 3$, with $D(4)$ identically zero by construction (the size-4 subset is the full set). Aggregation suppresses fragment-specific residuals, and the dominant sky structure is shared across reconstructions — the operational content of the record-consensus claim. Because the core-4 subset space is tiny (4, 6, 4, and 1 subsets at $k = 1 \dots 4$), the Monte Carlo estimate was checked against exact subset enumeration at conversion: the two agree to within $6.2 \times 10^{-8} \text{ K}$ ($\approx 0.06 \mu\text{K}$) at every k .




Figure 2. Core-4 consensus curve $D(k)$ (conversion re-execution): RMS residual of k -map subset averages against the full-set mean, over 800 Monte Carlo subsets per k , seed 7. Error bars are ± 1 sd across subsets.

4.4 Pre-specified question ledger

RQ	Result	Status
RQ1	Strong cross-product consensus in both fragment sets; $D(k)$ declines rapidly and stably with k	Supported
RQ2	$k_{\text{eff}} \approx 1.022$ (eight half-mission fragments); ≈ 1.004 (four full-mission products)	Quantified (descriptive)
RQ3	Agreement is dominated by shared output structure; the present design cannot partition shared sky signal from shared processing or shared systematics	Consensus established; source of shared structure not decomposed

Table 2. Outcomes for the pre-specified research questions of §2.1. RQ2's status is descriptive by design (§5); RQ3's answer is the paper's central interpretive boundary.

5. The Effective-Channel Boundary: Consensus Is Not Independence

The central interpretive discipline of this paper — introduced as a limitations bullet in the v0.2 text and promoted here to a section — is that the convergence documented above is a **record-consensus demonstration, not an independence audit**. Two facts force the distinction. First, the reference in $D(k)$ is the full-set mean, which is not independent of the channels being aggregated: $D(n_maps) = 0$ by construction, so the curve measures internal consensus, not error against an external truth. Second, all fragments derive from the same Planck observations and share upstream frequency maps and processing assumptions, so raw output correlation conflates shared signal with shared systematics.

The descriptive effective-channel count makes the redundancy concrete. At the measured correlations, $k_eff = k / (1 + (k - 1)\bar{r})$ evaluates to **1.022** for the eight half-mission fragments and **1.004** for the four full-mission products: under this adjustment, the entire ensemble carries approximately **one** effective independent channel. Averaging eight of these records is, in independent-evidence terms, barely different from consulting one — which is precisely what strong record redundancy looks like from the inside, and precisely why consensus among them must not be read as accumulating independent confirmation. The k_eff form itself is a first-order, equicorrelation-style descriptive adjustment using raw output correlation — a warning against overreading redundancy as independence, not an estimate of the physical information degrees of freedom in the Planck products; the audited protocol that a genuine independence claim requires — an independently specified channel model (e.g., FFP10-style noise simulations), a model-predicted variance-reduction curve, and a k_eff audit on residuals against an independent reference — is specified in the companion methods note (Jones 2026a). A companion empirical paper executing that full audit against this analysis infrastructure is the designed extension.

6. End-to-End Re-Execution from Primary Archives at Conversion

At T16 conversion (2026-07), the robustness configuration (Run B) was re-executed end-to-end in a fresh environment, using nothing not publicly available: the PR3 common intensity mask and all four full-mission component maps (≈ 7 GB) were re-downloaded from the IRSA Planck archive via the URLs in the package's own manifest; the 10^6 -pixel masked sample was regenerated with the package's own seeded sampling function (seed 7); Stokes-I extraction and mean-centering used the package's own extraction function; and the pairwise and aggregation analyses were recomputed with the locked metric definitions. Every reported digit reproduces: $\bar{r} = 0.9951544$; off-diagonal correlation range 0.9924926–0.9979500; RMS range $6.8637 \times 10^{-6} - 1.3182 \times 10^{-5}$ — exact to all seven published decimal places. The re-execution additionally recovered the full pairwise matrix (Figure 1) and the $D(k)$ curve (Figure 2), and its artifacts (sample indices, per-map sampled vectors, matrices, curve values) accompany this deposit.

Scope of the re-execution: Run B was reproduced end-to-end; Runs A and A' (the eight half-mission maps, ≈ 13 GB) were not re-downloaded at conversion, and their reproduction path is identical in kind — same manifest, same functions, same seed discipline. The division of labor stands as designed: the half-mission set remains the primary analytic fragment set, and the core-4 run is the conversion-proven public-recovery demonstration. Four months after the original run, in a fresh environment and through a separately executed recovery workflow, the published result was recovered from public materials alone. For a diagnostic whose registration mechanism is the weakest in the family (§2.2), this is the operative guarantee, demonstrated rather than promised.

7. Deviations

The package's Update Integrity Standard-style deviations log (Jones 2026e) contains **zero entries**: no source substitutions, no fallback invocations, no protocol departures were recorded for the executed runs. One maintenance event is documented in the Registered Analysis Report and reported here for completeness: a numeric-stabilization patch to the descriptive exponential-with-floor fit routine (clamping the log-amplitude and model exponent to float64-safe bounds, and skipping non-finite candidate fits during the floor search), applied to prevent overflow when fitting against k_{eff} on the highly correlated core-4 inputs. The anchored v0.1 package remains unchanged and carries the original routine; the maintenance-patched script is deposited separately alongside the conversion artifacts as `analysis_cmb_consensus_patched_2026_07.py`. The patch affects only the descriptive curve summary; the primary consensus metrics — $\bar{r}$, the pairwise tables, and $D(k)$ — are identical between the unpatched and patched runs (Table 1, last two rows). No conclusion in this paper depends on that fit routine.

8. Interpretation (Bounded)

What the result licenses. Under the PR3 common intensity mask, in temperature, the published component-separated products behave as strongly redundant records of the same sky signal: pairwise agreement is high in both fragment sets ($\bar{r} = 0.9757$ across half-mission fragments; 0.9952 across full-mission products), aggregation suppresses fragment-specific residuals along a rapidly declining, seed- and sample-stable $D(k)$ curve, and the structure of the result is recoverable to every published digit from public materials alone (§6). This is a robustness statement about the published products — the released PR3 temperature reconstruction is stable across component-separation pipelines within the masked region — and it is a worked, reproducible instance of a records-based consensus diagnostic of the kind the standards layer describes (Jones 2026c, 2026d).

What the result does not establish. It does not establish independence: at the measured correlations the ensemble carries ≈ 1 effective channel (§5), so the convergence documents consensus among redundant records, not accumulation of independent evidence — the audited-independence question is open by design and assigned to the companion protocol (Jones 2026a). It does not test Λ CDM, evaluate anomalies, or constrain parameters; it is not a re-analysis of the Planck likelihood. It is temperature-only (polarization products carry different noise and systematics), mask-conditional (quantities are defined under the PR3 common intensity mask), and pixel-space (no harmonic-domain claims). And it makes no ontological claim about the early universe: within UCT the result is an operational example of redundant-record stabilization; outside UCT it is a reproducibility check, and both readings are licensed by the same numbers.

9. Limitations and Falsifier Conditions

9.1 Scope conditions

Not independent experiments. All fragments derive from the same Planck observations and share upstream frequency maps and processing assumptions; §5 quantifies the consequence and bounds the claim.

Mask dependence. Results are computed under the PR3 common intensity mask; alternative masks or inpainting choices may change quantitative values.

Temperature only. Stokes-I; polarization is untested here.

Descriptive fitting only. The exponential-with-floor summaries of $D(k)$ are shape aids, not physical models; the k_{eff} axis is a first-order descriptive adjustment (§5).

Registration strength. Per §2.2: internal protocol lock plus versioned package; no third-party filing; empty deviations log. The compensating guarantee — end-to-end public reproducibility — is demonstrated in §6, but a future run of this diagnostic class should file the preregistration before execution as a matter of hygiene.

9.2 Falsifiers

The interpretation in §8 should be revised or withdrawn if: (i) the audited-independence protocol (independent channel model, model-predicted variance-reduction curve, residual-based k_{eff} against an independent reference; Jones 2026a) shows the observed convergence to be attributable to shared systematics in a way that voids the consensus reading of the published products; (ii) a pre-specified alternative-mask or inpainting variant materially breaks the pairwise agreement or the $D(k)$ decline, with materiality thresholds pre-specified in that variant's own locked protocol; (iii) a polarization run under the analogous locked protocol fails to show consensus where instrument understanding predicts it, or shows it where systematics are known to differ — either divergence localizing a pipeline failure mode; or (iv) re-execution from the public archives fails to reproduce the published values — a condition this deposit has already passed once (§6). Of these, (i) — the executed independence audit against this analysis infrastructure — is the designed companion and is ordered first.

10. Provenance and Reproducibility

Artifact	Identity / location
Reproducibility package	CMB_Record_Consensus_Test_Package_v0.1 — SHA-256 8daa6d30f3607bf2874137a85496cea05b363e4fa2c9aa59962cc48222ee0794 (recorded at conversion; the anchor for this deposit)
Package contents	analysis_cmb_consensus.py (full pipeline); maps_hm_default.csv / maps_core4_full.csv (locked map manifests with IRSA URLs); scripts/download_planck_pr3_cmb_hm.sh; OSF_Prereg_Template.md; sources.md; deviations.md (empty); requirements.txt; README.md
Data inputs (public)	Planck PR3 component maps (Commander, NILC, SEVEM, SMICA; full- mission and HM1/HM2) and COM_Mask_CMB-common-Mask- Int_2048_R3.00.fits, via the IRSA Planck archive
Original run record	Internal Registered Analysis Report v0.1 (2026-02-18): protocol §A, execution record §B, run inventory, range diagnostics, patch note, seed- sensitivity table
Run outputs (per configuration)	outputs/, outputs_smoke/, outputs_core4/, outputs_core4_patched/ — curve_Dk.csv, pairwise_corr.csv, pairwise_rms.csv, table_fit_summary.csv, metadata.json, figures
Conversion re-execution artifacts	Seed-7 masked-pixel index set; per-map sampled vectors; recomputed pairwise matrices and $D(k)$ curve (runB_repro.json); packaged as CMB_Record_Consensus_Verification_Artifacts_2026_07.zip, accompanying this deposit alongside the original package (SHA-256 524e81c1ef83bab9cc23674b3c48da9a323fd68275c3a84c26a45d07cc054069). Deposited alongside as loose supplements:

	analysis_cmb_consensus_patched_2026_07.py (the §7 maintenance patch) and download_from_manifest.sh (generic manifest downloader for either fragment set); the two anchored archives remain unchanged
Reproduction	bash download_from_manifest.sh maps_core4_full.csv (supplemental, this deposit); python analysis_cmb_consensus.py --maps maps_core4_full.csv --mask data/COM_Mask_CMB-common-Mask-Int_2048_R3.00.fits --n_sample 1000000 --seed 7

Table 3. Provenance chain. Raw Planck FITS files are large and publicly archived; per the package's design, run artifacts are archived rather than raw data duplicated, and §6 demonstrates that this suffices for exact recovery.

References

- Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., and Bartelmann, M. (2005). HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *The Astrophysical Journal* 622(2): 759–771. <https://doi.org/10.1086/427976>.
- Jones, J. C. (2026a). Auditing Independence in Multi-Channel Measurement: An Agreement-Curve Diagnostic for Genuine versus Correlated Redundancy (Methods-S₁, v1.0). HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/7U8SK>.
- Jones, J. C. (2026b). Empirical Demonstration of S₁, S₂, and S₃ in Deployed AI Systems (v1.0). HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/JPXCU>.
- Jones, J. C. (2026c). Objectivity from Records: An Exponential Consensus Bound (TN-S₁, v1.0). HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/6M7N3>.
- Jones, J. C. (2026d). Records Across Nature, Life, and Mind (v2.0). HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/7H6DY>.
- Jones, J. C. (2026e). Update Integrity Standard (UIS) (v1.0). HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/DWM29>.
- Planck Collaboration (2020). Planck 2018 results. IV. Diffuse component separation. *Astronomy & Astrophysics* 641: A4. <https://doi.org/10.1051/0004-6361/201833881>.

Document status. This is the v1.0 T16 deposit version of Redundant Record Consensus in Planck PR3 CMB Component Maps. The executed independence audit against this infrastructure (per Jones 2026a) is the designed companion. This paper deposits under the T16 empirical umbrella alongside Entropy as Record, Rice Hysteresis, COGITATE iEEG, and the AI-substrate demonstration (Jones 2026b).

Library note. This paper is part of the Universal Collapse Theory library, published by HoldingLight LLC. It is a physics-wing entry of the T16 empirical corpus. For a reading guide and full architecture, visit universalcollapse.com/roadmap.

AI Disclosure. AI tools were used to assist with manuscript preparation, drafting, organization, and editorial refinement. In this T16 empirical paper, AI assistance additionally extended to data evaluation: analysis-pipeline development in the original February 2026 execution, and, at the July 2026 conversion, the end-to-end re-execution of the robustness configuration from the primary public archives reported in §6 (using the package's own sampling and extraction functions), computation of the effective-channel boundary numbers of §5, and regeneration of the figures. No AI system is a subject of study in this paper.

The underlying theory, structural decisions, analysis, criteria, and conclusions are the author's own, and the author takes full responsibility for the manuscript and its contents.

Citation. Jones, J. C. (2026). Redundant Record Consensus in Planck PR3 CMB Component Maps (v1.0). HoldingLight LLC.

Contact. Inquiries about methodology, factual corrections, or replication results should be directed to contact@universalcollapse.com.