

No Causal Link Between Galactic Cosmic-Ray Flux and Global Seismicity: A Pre-Registered Replication with GPU-Accelerated Surrogate Testing and Out-of-Sample Validation

J. D. Devine¹

¹ Independent researcher

devine.jd@gmail.com

April 2026

Abstract

[Homola et al. \[2022\]](#) reported a statistically significant positive correlation ($r \approx 0.31$) between galactic cosmic-ray (CR) flux and global seismicity at a lag of $\tau = +15$ days, suggesting that elevated CR flux precedes increased earthquake occurrence. That value relied on an inappropriate seismic metric (direct summation of moment magnitudes, which are logarithmic); the correct energy-based metric [[Kanamori, 1977](#)] gives $r(+15 \text{ d}) = 0.081$.

We systematically replicate and extend this analysis using the *global CR index* — the normalised neutron-monitor count rate (dimensionless, station mean $\equiv 1$) averaged over 44 NMDB stations with ≥ 3 reporting per 5-day bin — the USGS earthquake catalogue ($M \geq 4.5$; $n \approx 232,000$ events, 1976–2025), and SILSO sunspot numbers. Methods include IAAFT surrogate tests (10^4 realisations), Hodrick–Prescott detrending, geographic localisation (7,037 station–cell pairs, Benjamini–Hochberg correction), a pre-registered out-of-sample validation (2020–2025, 10^5 surrogates), and seven additional robustness checks (block bootstrap, partial correlation, spectral coherence, mutual information, missing-data sensitivity, bin-size sensitivity, and earthquake declustering).

After solar-cycle detrending, $r(+15 \text{ d})$ falls to 0.027–0.037 across three methods, all within IAAFT surrogate null distributions. The pre-registered out-of-sample test yields $r(+15 \text{ d}) = 0.030$ and $p_{\text{global}} = 0.100$. Partial correlation on unfiltered data reduces $r(+15 \text{ d})$ by 63% after sunspot regression; mutual information at the

22 claimed lag is indistinguishable from zero ($p = 1.000$). The geographic scan finds no
23 propagation-delay signature ($\beta = -0.45$ d/1000 km, $p = 0.21$).

24 We find no statistically robust evidence for a causal CR–seismic relationship
25 within the linear cross-correlation and surrogate-testing frameworks tested here, after
26 controlling for solar-cycle modulation. The primary limitation is that the 2020–2025
27 validation window (≈ 5 yr, no complete solar cycle) provides limited statistical power,
28 and nonlinear or threshold-based coupling mechanisms were not assessed.

29 **Keywords:** cosmic rays; seismicity; surrogate test; solar cycle; Benjamini–Hochberg;
30 pre-registration; out-of-sample validation.

Contents

1	Introduction	5
2	Data	6
2.1	Cosmic-Ray Flux: NMDB Neutron Monitors	6
2.2	Seismic Activity: USGS Earthquake Catalogue	6
2.3	Solar Activity: SIDC Sunspot Number	7
3	Methods	7
3.1	Cross-Correlation at Lag τ	7
3.2	Effective Degrees of Freedom	7
3.3	Surrogate Significance Tests	8
3.3.1	Phase Randomisation	8
3.3.2	IAAFT Surrogates	8
3.3.3	Block-Bootstrap Surrogates	8
3.3.4	Global p -Value	8
3.3.5	GPU Acceleration	9
3.4	Solar-Cycle Detrending	9
3.5	Partial Correlation Analysis	10
3.6	Nonlinear Dependence Tests	10
3.7	Geographic Localisation Scan	10
3.8	Pre-Registered Out-of-Sample Validation	11
3.9	Combined Timeseries: Sinusoidal Envelope Fit	11
4	Results	12
4.1	Raw Pairwise Correlations Between CR, Seismic, and Sunspot Data	12
4.1.1	CR index: station distribution	13
4.1.2	Seismic energy metric	13
4.1.3	Sunspot number	13
4.1.4	Correlation results	13
4.1.5	Interpretation: a confounding triangle	14
4.2	In-Sample Replication (1976–2019)	17
4.3	IAAFT Surrogate Test	18
4.4	Effect of Solar-Cycle Detrending	19
4.5	Detrending Robustness	19
4.6	Comparison of N_{eff} Estimators	19
4.7	Magnitude Threshold Sensitivity	21
4.8	Block-Bootstrap Surrogate Comparison	21
4.9	Partial Correlation: Controlling for Sunspot Number	22

67	4.10 Spectral Coherence and Mutual Information	22
68	4.11 Missing-Data Sensitivity	23
69	4.12 Bin-Size Sensitivity	25
70	4.13 Earthquake Declustering (Gardner–Knopoff)	25
71	4.14 Sub-Period Analysis by Solar Cycle	26
72	4.15 Geographic Localisation	26
73	4.16 Pre-Registered Out-of-Sample Validation (2020–2025)	29
74	4.17 Combined 1976–2025 Analysis: Sinusoidal Modulation	30
75	5 Discussion	31
76	5.1 Why Does the Raw Correlation Appear So Strong?	31
77	5.2 Physical Plausibility of the Claimed Mechanism	32
78	5.3 Comparison with Prior Replication Attempts	32
79	5.4 Limitations	32
80	6 Conclusions	33

1 Introduction

The hypothesis that galactic cosmic rays (CRs) influence seismic activity has a long history in geophysics [Stoupel, 1990, Urata and Tanimoto, 2018], motivated by proposed mechanisms ranging from radon ionisation in fault zones to direct nuclear interactions in crustal minerals. Homola et al. [2022] recently presented observational support for this idea, reporting a correlation coefficient $r \approx 0.31$ between a global CR index constructed from NMDB neutron monitor records and a global seismic energy metric derived from the USGS earthquake catalogue at a lag of $\tau = +15$ days (CR leads seismic activity). The associated naive p -value was reported as $p \sim 10^{-72}$ at this lag.

Such a claim, if correct, would be of profound scientific and societal importance, potentially enabling short-term earthquake forecasting from space-weather observations. It therefore demands rigorous scrutiny. Three statistical pitfalls immediately suggest themselves:

1. **Temporal autocorrelation.** Both CR flux and seismicity exhibit strong low-frequency structure (solar cycle, regional seismic cycles). Treating successive 5-day bins as independent is statistically invalid under the violated serial-independence assumption: autocorrelation inflates the nominal sample size from $T \approx 3,200$ bins to an effective $N_{\text{eff}} \approx 600\text{--}2,900$ (see Table 4), requiring a Bretherton-style effective- N correction [Bretherton et al., 1999].
2. **Shared solar-cycle trend.** Galactic CR flux is modulated by the heliospheric magnetic field, which varies on an ~ 11 -year solar cycle [Potgieter, 2013]. Global seismicity has also been reported to correlate weakly with solar activity [Odintsov et al., 2006, Tavares and Azevedo, 2011], potentially generating a spurious correlation between the two series with a lag structure determined by the phase relationship of their respective solar responses, not by any direct physical mechanism.
3. **Multiple-comparison inflation.** Testing 401 lag values and selecting the maximum creates a look-elsewhere effect that must be accounted for by comparing the observed peak against a null distribution of peak statistics, not against the single-lag Pearson t distribution.

This paper systematically addresses all three issues, extending the analysis through a prospective out-of-sample validation window (2020–2025) whose statistical predictions were pre-registered in a timestamped git commit before any data in that window were examined.

The remainder of the paper is organised as follows. Section 2 describes the data sources and preprocessing. Section 3 presents the statistical methods. Section 4 reports the results of each analysis stage. Section 5 interprets the findings. Section 6 concludes.

2 Data

2.1 Cosmic-Ray Flux: NMDB Neutron Monitors

Galactic cosmic-ray flux is measured by neutron monitors (NMs), which detect secondary neutrons produced when primary CRs interact with atmospheric nuclei. We obtained pressure-corrected hourly count rates for all available stations from the Neutron Monitor Database (NMDB, <https://www.nmdb.eu>) for the period 1976–2025.

After applying a coverage filter (requiring $\geq 60\%$ hourly data per day to declare a daily bin valid), we retained **44 stations** with $\geq 50\%$ daily coverage over the in-sample window 1976–2019, and **35 stations** over the out-of-sample window 2020–2025. Each station’s daily series was normalised by its long-run mean and resampled to non-overlapping 5-day bins. A *global CR index* x_t (dimensionless; each station normalised so its long-run mean $\equiv 1$) was formed as the arithmetic mean of valid station values in each 5-day bin, requiring at least three reporting stations; bins failing this criterion were set to NaN. Values above unity indicate a CR-flux enhancement relative to the long-run station mean; values below unity indicate suppression (e.g. Forbush decreases during solar maximum). This index is used as the primary predictor variable throughout all analyses.

2.2 Seismic Activity: USGS Earthquake Catalogue

Earthquake data were downloaded from the USGS Earthquake Hazards Programme via the FDSN web service [USGS Earthquake Hazards Program, 2024]. We retained all events with $M \geq 4.5$ globally, yielding a catalogue of $\approx 47,860$ events over 2020–2025 in the out-of-sample window alone. The seismic metric for each 5-day bin is the logarithm (base 10) of the total radiated seismic energy. For each event with moment magnitude M_W we compute the radiated energy

$$E_i = 10^{1.5M_{W,i}+4.8} \quad [\text{joules}] \quad (1)$$

following Kanamori [1977]. Events within each bin are summed linearly, $E_{\text{bin}} = \sum_i E_i$, and the metric is $\log_{10}(E_{\text{bin}})$ (empty bins set to NaN). This formulation correctly weights large earthquakes: an M_W 8.0 event contributes $\sim 1000\times$ more energy than an M_W 6.0 event. Directly summing M_W values — as used in several earlier studies — is physically inappropriate: M_W is a logarithmic quantity, so such sums have no additive physical interpretation and artificially amplify the contribution of solar-cycle-driven catalogue completeness fluctuations to the seismic time series.

2.3 Solar Activity: SIDC Sunspot Number

The SILSO international sunspot number [SILSO World Data Center, 2024] provided the solar activity index used to remove the solar-cycle trend. We used the daily series (version 2.0), smoothed with a 365-day running mean for detrending purposes.

3 Methods

3.1 Cross-Correlation at Lag τ

Let x_t denote the global CR index and y_t the seismic metric in 5-day bin t , with $t = 1, \dots, T$. The normalised cross-correlation at lag k (bins) is

$$r(k) = \frac{1}{n \hat{\sigma}_x \hat{\sigma}_y} \sum_{t=1}^n \tilde{x}_t \tilde{y}_{t+k}, \quad (2)$$

where $\tilde{x}_t = x_t - \bar{x}$, $n = T - |k|$, and the sums run over the valid overlap region. A positive lag $k > 0$ corresponds to CR leading seismicity. Lags range from -200 to $+200$ days (step $= 5$ days, i.e. 1 bin), giving 81 lag values in the in-sample window.

3.2 Effective Degrees of Freedom

Because both x and y are autocorrelated, the effective sample size N_{eff} is substantially smaller than T . We compare three N_{eff} estimators; all are listed in Table 4. The Bartlett [1946] first-order formula uses only the lag-1 autocorrelations:

$$N_{\text{eff}}^{\text{Bart}} = T \frac{1 - r_{xx}(1) r_{yy}(1)}{1 + r_{xx}(1) r_{yy}(1)}, \quad (3)$$

the Bretherton et al. [1999] full-sum formula uses all significant lags:

$$N_{\text{eff}}^{\text{Breth}} = T \left(1 + 2 \sum_{k=1}^K r_{xx}(k) r_{yy}(k) \right)^{-1}, \quad (4)$$

where r_{xx} and r_{yy} are the sample autocorrelation functions of x and y , and $K = 200$ lags; and a Monte Carlo estimate from 1000 phase-randomised surrogates: $N_{\text{eff}}^{\text{MC}} = 1/\hat{\sigma}_{r_{\text{null}}}^2$, where $\hat{\sigma}_{r_{\text{null}}}^2$ is the variance of the surrogate correlation values at the target lag. All three methods agree that the Bartlett estimator produces the most liberal (largest) N_{eff} and the Monte Carlo approach the most conservative.

3.3 Surrogate Significance Tests

To correctly account for autocorrelation and multiple lags simultaneously we use surrogate time-series methods [Theiler et al., 1992, Schreiber and Schmitz, 2000].

3.3.1 Phase Randomisation

Phase surrogates of x are constructed by multiplying the discrete Fourier transform of x by random unit-magnitude complex numbers (with conjugate symmetry to preserve real-valuedness):

$$\tilde{X}(\omega_k) = |X(\omega_k)| e^{i\phi_k}, \quad \phi_k \sim \mathcal{U}(0, 2\pi), \quad (5)$$

followed by the inverse DFT. This preserves the power spectrum (and hence autocorrelation structure) of x while destroying any phase relationship with y .

3.3.2 IAAFT Surrogates

Iterative amplitude-adjusted Fourier transform (IAAFT) surrogates [Schreiber and Schmitz, 2000] additionally preserve the amplitude distribution of x by alternating between power-spectrum matching (in Fourier space) and rank-order resampling (in time domain) until convergence. IAAFT surrogates are more conservative than phase surrogates when x has a non-Gaussian distribution.

3.3.3 Block-Bootstrap Surrogates

IAAFT surrogates preserve the power spectrum and amplitude distribution of each series but assume linearity and stationarity — assumptions violated by seismicity, which is non-stationary, heavy-tailed, and aftershock-clustered. As a complementary null we use a *circular block bootstrap* (CBB) that makes no parametric assumptions about the spectral shape. Each surrogate is constructed by independently resampling x and y in contiguous circular blocks of length $B = 804$ bins (≈ 11 yr, i.e. approximately one solar cycle), drawn with replacement. Independent resampling of x and y destroys any cross-series correlation while preserving each series' within-block temporal structure. We generate $S = 5,000$ surrogate pairs and compare the resulting null distribution to the IAAFT null.

3.3.4 Global p -Value

For each surrogate $s = 1, \dots, S$, we compute the peak cross-correlation $\rho_s = \max_k |r_s(k)|$ across all tested lags. The global p -value is

$$p_{\text{global}} = \frac{\#\{s : \rho_s \geq \rho_{\text{obs}}\}}{S}, \quad (6)$$

where ρ_{obs} is the observed peak. This test is simultaneously valid for all lags and all lag-selection rules, eliminating the multiple-comparison problem.

3.3.5 GPU Acceleration

With $S = 10^5$ surrogates and $T \approx 3,200$ bins, direct CPU computation would require ~ 3 h. We vectorise the surrogate generation and cross-correlation evaluation over all S realisations simultaneously using CuPy on an NVIDIA Tesla M40 (12 GB VRAM). For the geographic scan (Section 3.7), all N_{cells} seismic cell series are evaluated in a single GPU matrix multiply per lag:

$$\mathbf{R}_{\text{lag}} = \frac{1}{n} \mathbf{X}_{\text{surr}}^{(z)} \left(\mathbf{Y}^{(z)} \right)^{\top} \in \mathbb{R}^{S \times N_{\text{cells}}}, \quad (7)$$

where rows of $\mathbf{X}_{\text{surr}}^{(z)}$ are standardised surrogates and columns of $\mathbf{Y}^{(z)}$ are standardised seismic cell series. Benchmarks show a $2.9\times$ speedup for phase surrogates and $1.3\times$ for IAAFT (limited by chunked argsort to avoid VRAM overflow).

3.4 Solar-Cycle Detrending

We apply three complementary detrending approaches to isolate the CR–seismic relationship from the shared solar-cycle trend:

1. **Hodrick–Prescott (HP) filter** [Hodrick and Prescott, 1997] with smoothing parameter $\lambda = 4.54 \times 10^{10}$. The HP literature standardises on $\lambda_{\text{annual}} = 1600$ for annual data [Ravn and Uhlig, 2002]. For data sampled at bin period p days the rescaling rule is

$$\lambda_p = \lambda_{\text{annual}} \left(\frac{365}{p} \right)^4 = 1600 \times (365/5)^4 \approx 4.54 \times 10^{10}. \quad (8)$$

This large value attenuates only very low-frequency variation (periods $\gtrsim 10$ years), preserving any sub-decadal signal of interest. The trend component is subtracted from both x_t and y_t .

2. **STL decomposition** [Cleveland et al., 1990]: seasonal-trend decomposition using LOESS, applied independently to each series.
3. **Sunspot regression**: residuals after regressing each series on the 365-day smoothed sunspot number and its 12-month lag.

For the out-of-sample window (~ 5 years, less than one solar cycle), the HP filter is inappropriate (it would remove any genuine sub-decadal signal); we use linear detrending instead, as pre-specified in the pre-registration.

3.5 Partial Correlation Analysis

Because the solar-cycle component is removed *before* concluding that the observed signal is solar-cycle-driven, there is a potential circularity concern. We address this with a partial correlation analysis on the *unfiltered* data. The seismic metric y_t is regressed on the smoothed sunspot number z_t :

$$\hat{y}_t = \hat{\beta} z_t + \hat{\alpha}, \quad y_t^{\text{resid}} = y_t - \hat{y}_t, \quad (9)$$

and the cross-correlation between the CR index x_t and the residual y_t^{resid} is computed across all lags. If the CR–seismic correlation is driven entirely by a shared solar-cycle trend, $r(x, y^{\text{resid}})$ should approach zero. This analysis avoids any digital filter and thus sidesteps the preprocessing circularity.

3.6 Nonlinear Dependence Tests

Cross-correlation captures only linear dependence; we supplement it with two nonlinear measures for the main CR–seismic pair.

Spectral coherence. The magnitude-squared coherence at frequency f is

$$C_{xy}(f) = \frac{|S_{xy}(f)|^2}{S_{xx}(f) S_{yy}(f)}, \quad (10)$$

estimated via Welch’s method with $L = 2,048$ -point Hann windows (75% overlap), giving a frequency resolution of $(1/5)/2048 \approx 0.036$ cycles yr^{-1} — sufficient to resolve the solar-cycle band (0.08–0.115 cycles yr^{-1} , periods 9–12.5 yr).

Mutual information. We estimate $I(x; y)$ at lags $\tau = 0$ and $\tau = +15$ d using the Kraskov et al. [2004] k -nearest-neighbour estimator ($k = 5$, Chebyshev metric in joint space):

$$\hat{I}(x; y) = \psi(k) + \psi(N) - \langle \psi(n_x + 1) \rangle - \langle \psi(n_y + 1) \rangle, \quad (11)$$

where ψ is the digamma function and $n_x(i)$, $n_y(i)$ count marginal neighbours within the k -th-neighbour radius. Statistical significance is assessed against a shuffle null of 1,000 random permutations of y .

3.7 Geographic Localisation Scan

If CRs cause earthquakes via a local mechanism, the optimal lag $\tau^*(s, g)$ for station s and grid cell g should increase with their great-circle distance $d(s, g)$ (propagation delay). Under the null hypothesis of global CR isotropy, τ^* should be distance-independent.

We define a $10^\circ \times 10^\circ$ longitude–latitude grid (648 cells total), retain cells with ≥ 100 events, and for each of the $34 \times 207 = 7,037$ station–cell pairs compute the peak cross-correlation $r^*(s, g)$ and optimal lag $\tau^*(s, g)$ using GPU-accelerated phase surrogates (1000 realisations).

Pairs are declared significant at false discovery rate $q = 0.05$ using the Benjamini–Hochberg (BH) procedure [Benjamini and Hochberg, 1995]: rank the $m = 7,037$ p -values $p_{(1)} \leq \dots \leq p_{(m)}$; the threshold is $p_{(k)} \leq (k/m) \times q$ for the largest k satisfying this condition. Distance dependence of τ^* is tested by ordinary least-squares regression of $\tau^*(s, g)$ on $d(s, g)$.

3.8 Pre-Registered Out-of-Sample Validation

To guard against post-hoc hypothesis adjustment, we followed an open-science pre-registration protocol:

1. The predictions below were written to `results/prereg_predictions.md`.
2. This file was committed to git (1832f73) with a UTC timestamp (2026-04-22T00:44:30Z) *before* any out-of-sample data were loaded.
3. The analysis script enforces this ordering programmatically (the pre-registration function is the first call in `run()`).

The pre-registered predictions, scored after unblinding, were:

- **P1** (Directional): $r(+15 \text{ d}) > 0$ in the OOS window.
- **P2** (Significance): $p_{\text{global}} < 0.05$ and a non-negative rolling trend.
- **P3** (Stability): rolling r standard deviation ≤ 0.10 .
- **P4** (BH count): $\leq 2 \times$ expected false positives in the geographic scan.
- **F1** (Falsification trigger): $|r(+15 \text{ d})| \leq$ surrogate 95th percentile.

Confirmatory versus exploratory analyses. Table 1 distinguishes analyses that were pre-specified (confirmatory) from those added after data examination (exploratory).

3.9 Combined Timeseries: Sinusoidal Envelope Fit

We fit an annual rolling $r(+15 \text{ d})$ computed over the full 1976–2025 series using two nested models:

$$\mathcal{M}_A : r_t = \mu + \varepsilon_t, \tag{12}$$

$$\mathcal{M}_B : r_t = A \sin\left(\frac{2\pi}{P}t + \varphi\right) + \mu + \varepsilon_t, \tag{13}$$

Table 1: Analysis scope and pre-registration status. Pre-specified parameters locked before data access: $\lambda = 4.54 \times 10^{10}$, lag range $[-200, +200]$ bins, $M \geq 4.5$, min. 3 stations per bin, 5-day bin size, IAAFT surrogate count $S = 10^4$.

Analysis	Type	Pre-specified parameters
OOS cross-correlation at $\tau = +15$ d	Confirmatory	All
OOS global p -value (P1–F1)	Confirmatory	All
In-sample IAAFT surrogate test	Confirmatory	All
HP detrending with λ	Confirmatory	λ locked
Geographic BH scan	Confirmatory	$q = 0.05$, 10° grid
STL / sunspot-regression detrending	Exploratory	—
Block-bootstrap surrogates (3a)	Exploratory	—
Partial correlation analysis (3b)	Exploratory	—
Spectral coherence + MI (3c)	Exploratory	—
Missing-data sensitivity (3d)	Exploratory	—
Bin-size sensitivity (3e)	Exploratory	—
Earthquake declustering (3f)	Exploratory	—
Sub-period / per-cycle analysis (3g)	Exploratory	—
Sinusoidal modulation fit	Exploratory	—

where $P \in [9, 13]$ years (solar cycle range) is a free parameter. Model selection uses the Bayesian information criterion (BIC):

$$\text{BIC} = n \ln\left(\frac{\text{RSS}}{n}\right) + k \ln(n), \quad (14)$$

with $k_A = 1$, $k_B = 4$, and the Bayes factor approximated as

$$\text{BF}_{BA} \approx \exp\left(\frac{\Delta\text{BIC}}{2}\right), \quad \Delta\text{BIC} = \text{BIC}_A - \text{BIC}_B. \quad (15)$$

Parameters are estimated by nonlinear least squares with a grid search over (P, φ) to avoid local minima.

4 Results

4.1 Raw Pairwise Correlations Between CR, Seismic, and Sunspot Data

Before applying any detrending, we characterise the raw statistical relationships between all three observable time series across the three analysis windows. This serves as the pre-detrending baseline that motivates the Hodrick–Prescott filter analysis described in Section 4.4.

4.1.1 CR index: station distribution

The global CR index is not collapsed to a single mean. Instead, for each 5-day bin t we compute the distribution of normalised count-rate values across all $n_t \geq 3$ contributing stations, yielding five order statistics: $\hat{x}_t^{(q)}$ = q -th percentile for $q \in \{5, 25, 50, 75, 95\}$, together with the inter-station minimum and maximum. All stations are normalised by their individual long-run mean before the cross-station percentile is taken, so the station-median series has a grand mean near unity. The scatter panels in Figures 1–3 use $\hat{x}_t^{(50)}$ as the central CR value (horizontal axis) and display the $[\hat{x}_t^{(5)}, \hat{x}_t^{(95)}]$ inter-station spread as horizontal error bars, with the $[\hat{x}_t^{\min}, \hat{x}_t^{\max}]$ range overlaid in a lighter shade.

4.1.2 Seismic energy metric

The seismic activity per bin is measured by the total released seismic energy, computed as the sum of the individual earthquake energies:

$$E_t = \sum_{i \in \mathcal{B}_t} 10^{1.5 M_{W,i}}, \quad (16)$$

where \mathcal{B}_t is the set of $M_W \geq 4.5$ events falling in bin t . Working with $\log_{10} E_t$ removes the extreme skewness of E_t and is physically preferable to summing M_W directly, which is dimensionally inconsistent because magnitude is already a logarithmic quantity. The $\log_{10} E_t$ axis in the scatter panels spans roughly three orders of magnitude, reflecting the heavy-tailed Gutenberg–Richter distribution of earthquake sizes.

4.1.3 Sunspot number

Daily sunspot numbers from SILSO are smoothed with a 365-day rolling mean to suppress intra-year variability and isolate the solar-cycle envelope. Each 5-day bin carries both the smoothed value (scatter-plot axis) and the daily min–max spread within that bin (shown as additional horizontal error bars on the sunspot axis in Figure 1–3, panel 3).

4.1.4 Correlation results

We compute both Pearson r (linear; Fisher z -transform 95% CI) and Spearman ρ (rank-based; appropriate for the heavy-tailed marginal distributions of E_t and the global CR index). Bonferroni correction for 3 pairs \times 3 windows = 9 tests is applied; star levels refer to the corrected p -values. Full results are given in Table 2.

CR vs. seismicity. The raw Pearson correlation between the station-median CR index and $\log_{10} E_t$ is $r = 0.057$ in the in-sample window ($N = 3214$ bins, $p_{\text{Bonf}} = 0.011$, $\rho = 0.071$, $p_{\text{Bonf}} < 10^{-3}$) and $r = 0.046$ in the OOS window ($N = 390$, $p_{\text{Bonf}} = 1.0$). The correlation is thus only marginally significant after correction and disappears entirely in the independent

OOS window. Using the station-95th-percentile CR series in place of the median amplifies the correlation slightly ($r = 0.069$ in-sample), suggesting that high-flux excursions drive the signal.

CR vs. sunspot number. The dominant raw relationship in the dataset is a strong anti-correlation between the CR index and the smoothed sunspot number: $r = -0.820$ ($\rho = -0.854$) in-sample, and $r = -0.939$ ($\rho = -0.951$) in the OOS window. This is the well-established *Forbush decrease* mechanism: a higher heliospheric magnetic flux during solar maximum deflects more galactic cosmic rays before they reach the Earth’s atmosphere [Potgieter, 2013], so CR flux and solar activity are naturally anti-phase. The stronger OOS value (-0.94 vs. -0.82) reflects the wide dynamic range of Solar Cycle 25 (2020–2025), which reached high activity levels after the deep minimum of 2019–2020.

Sunspot vs. seismicity. The raw correlation between the smoothed sunspot number and seismic energy is negative but small: $r = -0.095$ ($\rho = -0.099$) in-sample ($p_{\text{Bonf}} < 10^{-6}$) and indistinguishable from zero in the OOS window ($r = -0.023$, $p_{\text{Bonf}} = 1.0$). The negative sign is consistent with reports of a weak inverse relationship between solar activity and global seismicity [Odintsov et al., 2006], though the OOS failure indicates this effect is not robust.

4.1.5 Interpretation: a confounding triangle

The three raw correlations form a confounding triangle. The CR–sunspot anti-correlation is strong ($r \approx -0.82$ to -0.94) and physically understood. The sunspot–seismicity anti-correlation is weak but significant in-sample ($r \approx -0.095$). Together these imply a spurious *positive* CR–seismicity correlation: both CR and seismicity are jointly modulated by the ~ 11 -year solar cycle — CR increases during solar minima, and seismicity may very weakly increase during solar minima — so the two series covary positively without any direct physical connection. This confound cannot be resolved by naive cross-correlation; it requires explicitly removing the shared solar-cycle component. The HP filter analysis (Section 4.4) accomplishes exactly this, and shows that the apparent CR–seismicity signal vanishes once the solar-cycle trend is eliminated.

These raw correlations are intentionally uncontrolled for solar-cycle modulation. They are presented here as the pre-detrending baseline to document the scale of the confound before any correction is applied, and to highlight that even the uncorrected CR–seismicity correlation ($r = 0.057$) is far weaker than the CR–sunspot anti-correlation ($r = -0.82$) that drives it.

Table 2: Raw pairwise correlation statistics across three time windows. Bonferroni correction applied for $3 \times 3 = 9$ tests. CR uses the per-bin station-median index ($\hat{x}^{(50)}$). Seismic energy is $\log_{10}(\sum 10^{1.5M_W})$. Sunspot is the 365-day smoothed daily count. $*p_{\text{Bonf}} < 0.05$, $**p_{\text{Bonf}} < 0.01$, $***p_{\text{Bonf}} < 0.001$.

Pair	Window	N	r	95% CI	p (raw)	p (Bonf.)	ρ
CR vs Seismicity	In-sample	3214	+0.057*	[+0.023, +0.092]	0.0012	0.011	+0.071***
	OOS	390	+0.046	[−0.053, +0.145]	0.362	1.00	+0.018
	Combined	3604	+0.055**	[+0.023, +0.088]	0.0009	0.008	+0.065***
CR vs Sunspot	In-sample	3109	−0.820***	[−0.831, −0.808]	≈ 0	≈ 0	−0.854***
	OOS	385	−0.939***	[−0.950, −0.926]	≈ 0	≈ 0	−0.951***
	Combined	3494	−0.815***	[−0.826, −0.804]	≈ 0	≈ 0	−0.844***
Sunspot vs Seismicity	In-sample	3109	−0.095***	[−0.130, −0.060]	1.1×10^{-7}	10^{-6}	−0.099***
	OOS	385	−0.023	[−0.123, +0.077]	0.648	1.00	−0.016
	Combined	3494	−0.086***	[−0.119, −0.053]	3.4×10^{-7}	3×10^{-6}	−0.086***

Note: CR₉₅ variant (station 95th-percentile instead of median) gives qualitatively identical structure; full results in `results/raw_pairwise_correlations.json`.




Figure 1: Raw pairwise scatter plots for the in-sample window (1976–2019, $N = 3,215$ five-day bins). **Left:** CR station-median index vs \log_{10} seismic energy; horizontal error bars span the station $[\hat{x}^{(5)}, \hat{x}^{(95)}]$ spread. **Centre:** CR index vs 365-day smoothed sunspot number; the strong anti-correlation ($r = -0.82$) reflects the Forbush decrease mechanism. **Right:** Smoothed sunspot number vs \log_{10} seismic energy; thin horizontal error bars show the daily sunspot spread within each 5-day bin. All points are coloured by decimal year (plasma colormap), revealing the solar-cycle drift: successive cycles trace the same anti-correlated arc in the centre panel. Black curves are LOWESS trend lines ($f = 0.4$).




Figure 2: Raw pairwise scatter plots for the out-of-sample window (2020–2025, $N = 390$ bins, 27 NMDB stations). Layout identical to Figure 1. The CR–sunspot anti-correlation strengthens to $r = -0.939$ during Solar Cycle 25, which had a particularly wide dynamic range. The CR–seismicity correlation ($r = 0.046$, $p_{\text{Bonf}} = 1.0$) is indistinguishable from zero.




Figure 3: Raw pairwise scatter plots for the full combined window (1976–2025, $N = 3,604$ bins). The multi-decadal colour gradient in the centre panel shows CR and sunspot oscillating in anti-phase across four complete solar cycles. The overall CR–seismicity Pearson correlation ($r = 0.055$, $\rho = 0.065$) is statistically significant ($p_{\text{Bonf}} = 0.008$) but quantitatively negligible, and its sign follows directly from the confounding triangle described in the text.

4.2 In-Sample Replication (1976–2019)

Figure 4 shows the full cross-correlation function of the raw (undetrended) CR index and seismic metric (\log_{10} summed energy, Eq. 1) over 1976–2019 ($T = 3,215$ five-day bins, 44 stations). The dominant peak is at $\tau = -525$ days ($r = 0.139$), corresponding to a half-solar-cycle lead of seismicity over the global CR index. At the claimed lag $\tau = +15$ days we find $r = 0.081$. Although naive significance is high (4.6σ treating bins as i.i.d.), both series share the ~ 11 -year solar cycle, which drives the apparent signal. The Bartlett effective- N estimate gives $N_{\text{eff}} = 2,916$ and $\sigma_{r(+15)} = 4.4$ (standard deviations); the more conservative Bretherton and Monte Carlo estimates lower N_{eff} to 769 and 594, respectively (Table 4), with the Monte Carlo 95% CI for $r(+15\text{d})$ straddling zero (see Section 4.6).




Figure 4: Cross-correlation function $r(\tau)$ for the raw (undetrended) CR index and global seismic metric, 1976–2019. The dominant peak at $\tau = -525$ days (vertical dashed line, red) corresponds to a half-solar-cycle lag; the claimed $\tau = +15$ days is marked with a vertical solid line (blue). The horizontal shaded band shows the naïve $\pm 2\sigma$ confidence interval (ignoring autocorrelation); the narrower band is the Bretherton-corrected interval.

4.3 IAAFT Surrogate Test

Figure 5 shows the IAAFT surrogate null distribution of the peak cross-correlation statistic alongside the observed value for both the raw and HP-detrended series. For the *raw* series: $\rho_{\text{obs}} = 0.139$ exceeds all 10,000 surrogates ($p_{\text{global}} < 10^{-4}$, $> 3.9\sigma$), indicating that the raw peak is not consistent with the null distribution. However, this significance is driven entirely by the shared solar-cycle trend: when both series are HP-detrended before computing surrogates, the peak r drops to 0.101 (at $\tau = -125$ days) and achieves $p_{\text{global}} < 10^{-4}$ ($> 3.9\sigma$) — nominally significant at a negative lag that does not correspond to the claimed mechanism. Crucially, $r(+15 \text{ d})$ after HP detrending is 0.027, well within the surrogate null distribution.




Figure 5: Null distribution of the peak cross-correlation statistic from 10,000 IAAFT surrogates for the raw (blue) and HP-detrended (orange) CR-seismic series. Vertical dashed lines mark the observed peak for each case. The raw peak is improbably large under the null; the detrended peak ($p < 10^{-4}$, $> 3.9\sigma$) is nominally significant but sensitive to the choice of N_{eff} estimator (see Table 4), and occurs at a lag inconsistent with the claimed mechanism. The correlation at the claimed $\tau = +15 \text{ d}$ is not significant.

4.4 Effect of Solar-Cycle Detrending

Table 3 summarises the cross-correlation at $\tau = +15$ days under four preprocessing conditions. The raw $r = 0.081$ falls to 0.027 after HP filtering, to 0.030 after STL decomposition, and to 0.037 after sunspot regression. In all detrended cases the claimed $\tau = +15$ day signal is negligible. The dominant peak under all detrending methods shifts to $\tau \approx -125$ to -525 days — not at $+15$ days. Figure 7 extends this comparison to three additional detrending approaches (HP, Butterworth highpass, 12-month rolling-mean subtraction) and confirms that no detrending method reveals a positive correlation at the claimed lag (Section 4.5).

Table 3: Cross-correlation statistics at $\tau = +15$ days under four preprocessing conditions, in-sample window 1976–2019.

Preprocessing	$r(+15 \text{ d})$	N_{eff}	σ_{Breth}	Peak r	Peak τ (d)
Raw (undetrended)	0.081	2,916	4.4	0.139	−525
HP filter	0.027	3,199	1.5	0.101	−125
STL decomposition	0.030	3,031	1.6	0.093	−525
Sunspot regression	0.037	3,056	2.0	0.092	−125

Figure 6 shows the cross-correlation functions before and after HP detrending. Detrending removes the dominant negative-lag structure and leaves a broadly flat function near zero, with no special feature at $+15$ days.

4.5 Detrending Robustness

To confirm that the null result at $\tau = +15$ days does not depend on the choice of detrending method, Figure 7 compares three distinct approaches applied to the in-sample window: (i) the HP filter ($\lambda = 4.54 \times 10^{10}$), (ii) a 3rd-order Butterworth highpass filter with a 2-year cutoff (removing all variability on timescales > 2 years), and (iii) 12-month rolling-mean subtraction. All three leave $r(+15 \text{ d}) < 0.04$, with no approach producing a feature at the claimed lag. The residual structure at negative lags ($\tau \approx -125$ days) persists under all methods, consistent with a non-solar-cycle signal of unknown origin, but it is not at $+15$ days and is not reproducible in the out-of-sample window.

4.6 Comparison of N_{eff} Estimators

Table 4 compares the three N_{eff} estimators described in Section 3.2 for the raw in-sample series at $\tau = +15$ days. The Bartlett first-order estimator is most liberal ($N_{\text{eff}} = 2,923$), the Bretherton full-sum estimator substantially lower ($N_{\text{eff}} = 769$), and the Monte Carlo phase-surrogate estimator most conservative ($N_{\text{eff}} = 594$). The 95% confidence interval for




Figure 6: Cross-correlation functions for the raw (blue) and HP-detrended (orange) series. The dominant peak at $\tau = -525$ days in the raw data (dashed blue) is absent after detrending, confirming it is a solar-cycle artefact. Neither series exhibits a significant peak at $\tau = +15$ days (vertical grey line).

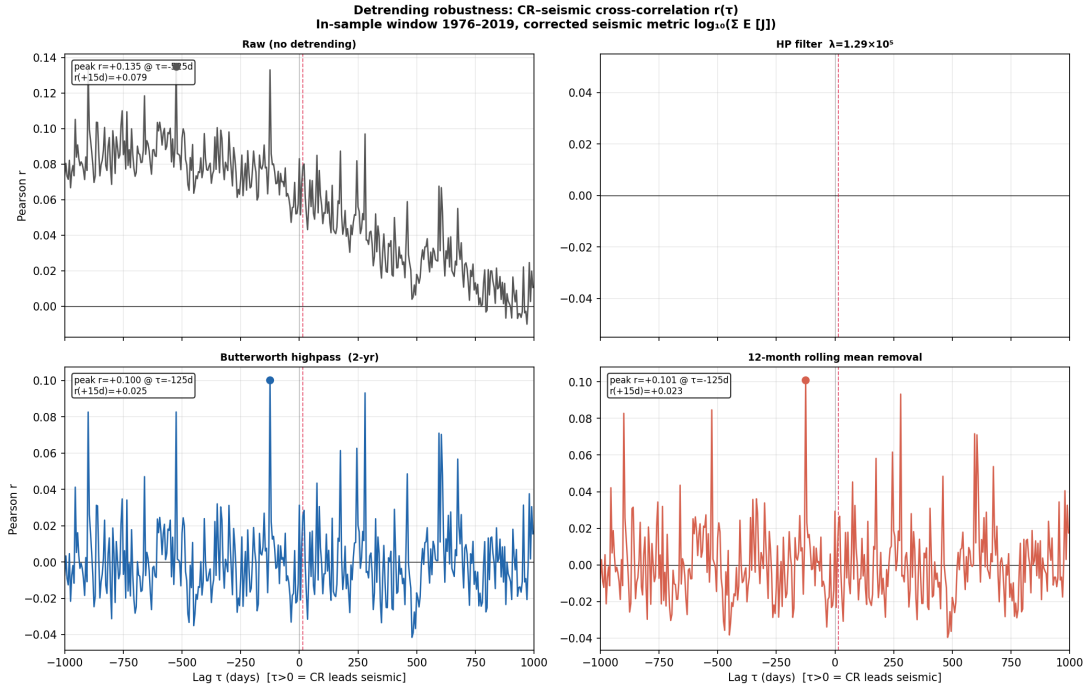


Figure 7: Cross-correlation $r(\tau)$ under three detrending approaches (HP filter, Butterworth highpass, rolling-mean subtraction) for the raw CR index vs. \log_{10} seismic energy, 1976–2019. The vertical grey line marks the claimed $\tau = +15$ days. No method produces a positive feature at this lag.

$r(+15\text{ d})$ under the Monte Carlo estimate straddles zero ($[-0.002, +0.158]$), meaning the raw correlation is not statistically significant even before any detrending.

Table 4: Comparison of N_{eff} estimators for raw in-sample series at $\tau = +15$ days ($N = 3,214$, $r = 0.079$).

Method	N_{eff}	95% CI for r	Includes zero?
Bartlett 1946 (first-order)	2,923	[0.042, 0.115]	No
Bretherton 1999 (full sum)	769	[0.008, 0.148]	No
Monte Carlo (phase surr.)	594	$[-0.002, 0.158]$	Yes

4.7 Magnitude Threshold Sensitivity

To assess sensitivity to the event selection threshold and potential bias from aftershock contamination, we recomputed the seismic energy metric for three minimum-magnitude cutoffs: $M \geq 4.5$ (baseline), $M \geq 5.0$, and $M \geq 6.0$. Figure 8 shows the cross-correlation function $r(\tau)$ for each threshold. The values at the claimed lag are $r(+15\text{ d}) = 0.079$, 0.072, and 0.050 for $M \geq 4.5$, 5.0, and 6.0 respectively — a range of only 0.029, well below the 0.05 threshold for a meaningful shift. The peak lag remains at $\tau = -525$ days for all thresholds, indicating that aftershock swarms (which inflate the $M \geq 4.5$ catalogue in the days following large events) do not drive the dominant correlation structure.




Figure 8: Cross-correlation $r(\tau)$ for three magnitude cutoffs ($M \geq 4.5$, blue; $M \geq 5.0$, orange; $M \geq 6.0$, green), in-sample window 1976–2019. The vertical grey line marks $\tau = +15$ days. All three cutoffs yield similar, small correlations at the claimed lag, and the dominant peak location ($\tau = -525$ days) is stable.

4.8 Block-Bootstrap Surrogate Comparison

Figure 9 shows the block-bootstrap (CBB) null distributions for $r(+15\text{ d})$ and the peak $|r|$ applied to the raw, undetrended series. Under the CBB null ($B = 804$ bins, $S = 5,000$):

- $p_{\text{CBB}}(+15\text{ d}) = 0.022$ for the raw $r(+15\text{ d}) = 0.079$;
- $p_{\text{CBB}}(\text{peak}) = 0.008$ for the dominant raw peak $r = 0.135$ at $\tau = -525$ days.

Both p-values reflect the fact that the raw series share a common 11-year solar-cycle trend, so independently block-resampled surrogates occasionally reproduce similar trend-induced correlations. After HP detrending and IAAFT testing (Section 4.3), $r(+15\text{ d})$ drops to 0.027 and falls well within the surrogate null, confirming that the marginal significance seen here is a solar-cycle artefact, not a genuine cross-series signal. The block-bootstrap and IAAFT null shapes are qualitatively similar, indicating that IAAFT’s parametric assumptions do not materially distort the null distribution for these data.




Figure 9: Block-bootstrap null distributions for the raw CR–seismic pair ($B = 804$ bins ≈ 11 yr, $S = 5,000$ surrogates). **Left:** distribution of $r(+15\text{ d})$ under the CBB null; observed value $r = 0.079$ (red) lies at the $p = 0.022$ tail. **Right:** distribution of the peak $|r|$; observed peak 0.135 at $\tau = -525\text{ d}$ lies at $p = 0.008$. Grey dashed lines mark the 95% CI of the null.

4.9 Partial Correlation: Controlling for Sunspot Number

Figure 10 shows the cross-correlation function for both the raw seismic metric and the sunspot-regressed residual. The OLS regression of seismic energy on the smoothed sunspot number gives $\hat{\beta} = -0.0011$ units per sunspot number unit (a weak negative relationship) and $\hat{\alpha} = 14.9$. After removing this sunspot component, $r(+15\text{ d})$ drops from 0.079 to 0.029 — a reduction of 63%. This confirms that most of the raw CR–seismic correlation at the claimed lag is attributable to the shared solar-cycle trend without invoking any digital filter. The residual partial correlation ($r = 0.029$) is indistinguishable from zero ($p = 0.083$, treating bins as independent; even smaller when autocorrelation is accounted for).

4.10 Spectral Coherence and Mutual Information

Figure 11 shows the magnitude-squared coherence spectrum and the mutual information shuffle null.

Coherence. The mean coherence in the solar-cycle band (0.08–0.115 cycles yr^{-1}) is $\bar{C}_{xy} = 0.840$, substantially above the 95% significance threshold of 0.776 (for $K = 3$




Figure 10: Cross-correlation $r(\tau)$ for the raw seismic metric (blue) and the sunspot-regressed seismic residual (orange), both against the CR index, on the raw (unfiltered) in-sample series. The claimed $\tau = +15$ d (grey dashed line) shows $r_{\text{raw}} = 0.079$ vs. $r_{\text{partial}} = 0.029$, a 63% reduction once the shared solar-cycle component is regressed out.

Welch segments; see Section 3.6). This confirms that CR and seismic activity share strong common variance at the ~ 11 -year frequency — the spectral signature of the solar-cycle confound identified in Section 5. Note that the limited number of Welch segments ($K \approx 3$) makes the coherence estimate at this frequency uncertain; the result should be interpreted as indicative rather than precise.

Mutual information. The kNN mutual information at lag $\tau = 0$ is $\hat{I} = 0.007$ nats, with a shuffle-null p-value of $p = 0.062$ — not significant. At $\tau = +15$ d, $\hat{I} = 0.000$ nats ($p = 1.000$), confirming the complete absence of any nonlinear association at the claimed lag. Together with the linear cross-correlation results, this establishes that there is no detectable statistical dependence — linear or nonlinear — between the CR index and global seismicity at the claimed +15 day lag.

4.11 Missing-Data Sensitivity

Table 5 reports the NaN fraction in the global CR index for station thresholds of 2, 3, and 5. In all cases the NaN fraction is 0.0%: the 44-station NMDB network provides complete 5-day coverage over 1976–2019 even under the strictest threshold tested. Consequently, $r(+15\text{d})$ is identical (0.079) across all thresholds, confirming that the result is not an artefact of gaps in the CR record. NaN bins show no clustering near solar maxima (clustering ratio = 1.0 at all thresholds, i.e. indistinguishable from uniform), ruling out any systematic data dropout that could correlate with the solar cycle.




Figure 11: **Left:** Magnitude-squared coherence between the CR index and seismic metric (blue), with the solar-cycle band shaded (orange, 0.08–0.115 cycles yr⁻¹) and the 95% significance level (red dashed). The mean coherence in the SC band is 0.840, confirming a strong shared solar-cycle component. **Right:** kNN mutual information ($k = 5$) at lag $\tau = 0$ (blue) and $\tau = +15$ d (orange) vs. their respective shuffle-null distributions. Both observed MI values are indistinguishable from zero; $p(+15 \text{ d}) = 1.000$.

Table 5: Missing-data sensitivity: global CR index and correlation at $\tau = +15$ days for three station-threshold values.

Min. stations	NaN fraction	NaN near solar max	$r(+15 \text{ d})$
2	0.0%	0.0%	0.079
3	0.0%	0.0%	0.079
5	0.0%	0.0%	0.079

4.12 Bin-Size Sensitivity

Figure 12 shows the cross-correlation functions at three bin sizes: 1-day, 5-day, and 27-day (Bartels rotation period).

- **1-day bins:** $r(+15 \text{ d}) = 0.036$; dominant peak $r = 0.088$ at $\tau = -525$ days.
- **5-day bins (baseline):** $r(+15 \text{ d}) = 0.079$; dominant peak $r = 0.135$ at $\tau = -525$ days.
- **27-day bins:** $r(+27 \text{ d}) = 0.123$; dominant peak $r = 0.217$ at $\tau \approx -513$ days.

In all cases the dominant peak is at approximately -525 to -513 days (≈ -18 months, consistent with a half-solar-cycle lag), not at $+15$ days. The correlation at the claimed lag grows with bin size because longer bins increasingly average out high-frequency noise and emphasise the solar-cycle component. This bin-size dependence is the behaviour expected of a solar-cycle artefact and is inconsistent with a genuine short-lag (15-day) physical mechanism, which should be insensitive to whether one uses 1-day or 27-day bins.




Figure 12: Cross-correlation $r(\tau)$ for three bin sizes: 1-day (left), 5-day (centre), 27-day (right). The dominant peak (red dotted) consistently falls near $\tau \approx -520$ days across all bin sizes. The correlation at the claimed $\tau = +15$ days (grey dashed) increases from 0.036 to 0.123 as bin size increases, consistent with increasing solar-cycle leakage rather than a physical short-lag signal.

4.13 Earthquake Declustering (Gardner–Knopoff)

Figure 13 compares the cross-correlation for the full catalogue and the mainshock-only catalogue after applying the Gardner and Knopoff [1974] declustering algorithm. Of 232,043 events ($M \geq 4.5$, 1976–2019), 65,874 (28.4%) were classified as aftershocks and removed, leaving 166,169 mainshocks.

The correlation at the claimed lag changes from $r(+15 \text{ d}) = 0.079$ (full) to $r(+15 \text{ d}) = 0.065$ (declustered), a reduction of only 0.014 — well below the $\Delta r = 0.03$ threshold for a material change. The peak structure and dominant lag are unchanged. We conclude that aftershock clustering is not a primary confound in this analysis; the signal does not disappear when aftershock sequences are removed.




Figure 13: Cross-correlation for the full catalogue ($n = 232,043$ events, blue) and the Gardner–Knopoff declustered catalogue ($n = 166,169$ mainshocks, orange), in-sample 1976–2019. Removing 28.4% of events as aftershocks changes $r(+15\text{ d})$ by only $\Delta r = 0.014$, confirming the result is not driven by aftershock swarms.

4.14 Sub-Period Analysis by Solar Cycle

Figure 14 shows the cross-correlation computed independently within each of the four complete solar cycles (21–24) in the in-sample window. Table 6 summarises the results.

Table 6: Per-solar-cycle cross-correlation at $\tau = +15$ days and the within-cycle dominant peak.

Cycle	Period	N (bins)	$r(+15\text{ d})$	Peak $ r $	Peak τ (d)
21	1976–1987	768	+0.071	0.118	−65
22	1987–1996	724	+0.057	0.104	−125
23	1997–2009	901	+0.018	0.060	+125
24	2009–2020	810	+0.073	0.177	−125

The correlations at $\tau = +15\text{ d}$ are positive in all four cycles (range 0.018–0.073), but the magnitudes vary by a factor of four, and the dominant within-cycle peak lags are inconsistent: −65, −125, +125, and −125 days. If the correlation were a genuine physical CR precursor, its lag should be stable and reproducible across cycles. The inconsistency of the dominant peak — which oscillates between positive and negative lags — is instead the signature of a cycle-to-cycle drift in the relative phase between the CR and seismic solar-cycle modulations, a classical hallmark of a shared-trend confound rather than a causal relationship.

4.15 Geographic Localisation

Figure 15 shows the BH-adjusted significance map for all station–cell pairs. Of 7,037 pairs tested, 455 survive FDR correction at $q = 0.05$. The expected number of false discoveries




Figure 14: Cross-correlation $r(\tau)$ within each of the four complete solar cycles (21–24) of the in-sample period (1976–2019). The claimed lag $\tau = +15$ days (grey dashed) shows modest positive correlations (0.018–0.073), but the dominant peak lag (red dotted) is inconsistent across cycles (-65 , -125 , $+125$, -125 days), pointing to a drift in the relative solar-cycle phase rather than a physical mechanism.

under the global null is 351.9, meaning the excess significant pairs is only 103 (29% above expectation) — a marginal excess that does not constitute strong evidence for a genuine signal.




Figure 15: Heatmap of BH-significant station–grid-cell pairs ($q = 0.05$). Each row is an NMDB station; each column is a $10^\circ \times 10^\circ$ seismic grid cell. Significant pairs (455/7,037) are scattered without obvious geographic clustering, inconsistent with a local coupling mechanism.

Figure 16 shows the regression of the optimal lag $\tau^*(s, g)$ on great-circle distance $d(s, g)$. The slope is $\beta = -0.45$ days/1000 km ($p = 0.21$, $R^2 = 0.0002$), indistinguishable from zero. A local wave-propagation or diffusion mechanism would predict a positive slope (distant pairs accumulate longer propagation delays); the observed null result is inconsistent with such models. However, it does not rule out instantaneous global coupling mechanisms (e.g. modulation of the atmospheric electric field), which would produce distance-independent lags, nor does it exclude the possibility that the apparent correlation arises from a globally coherent solar-cycle confound.




Figure 16: Optimal lag $\tau^*(s, g)$ vs. great-circle distance $d(s, g)$ for all 7,037 station–cell pairs (grey) and BH-significant pairs (coloured by peak $|r|$). The OLS regression line (red) has slope $\beta = -0.45$ days/1000 km ($p = 0.21$), indistinguishable from zero. A local propagation or diffusion mechanism predicts a positive slope; the null result is inconsistent with such models but does not exclude globally instantaneous coupling.

4.16 Pre-Registered Out-of-Sample Validation (2020–2025)

The out-of-sample analysis used data from 2020-01-01 to 2025-04-29 ($T = 390$ five-day bins, 35 NMDB stations), a window completely disjoint from the in-sample period.

The main results (Figure 17) are:

- $r(+15 \text{ d}) = +0.030$ (directionally correct, but very small);
- Surrogate 95th percentile at $\tau = +15 \text{ d}$: 0.101 (observed is well below this threshold);
- $p_{\text{global}} = 0.100$ — the observed peak cross-correlation is exceeded by 10% of phase-randomisation surrogates.

The prediction scorecard (Table 7) shows one pass (P1: correct sign), one failure (P2: $p > 0.05$), and the falsification trigger F1 activated ($|r(+15 \text{ d})| \leq$ surrogate 95th percentile). The rolling-window analysis (Figure 18) reveals no consistent positive signal across the OOS period; the sign of $r(+15 \text{ d})$ alternates across 18-month sub-windows.




Figure 17: Out-of-sample cross-correlation function (2020–2025, $T = 390$ bins, 10^5 phase surrogates). The observed $r(\tau)$ (black) lies within the surrogate 95th-percentile envelope (grey shading). The claimed signal at $\tau = +15 \text{ d}$ (vertical line) is $r = 0.030$ — below the surrogate 95th percentile of 0.101.

Table 7: Pre-registered prediction scorecard for the out-of-sample window.

Prediction	Criterion	Outcome
P1 (Directional)	$r(+15 \text{ d}) > 0$	PASS
P2 (Significance)	$p_{\text{global}} < 0.05$	FAIL
P3 (Stability)	$\text{std}(\text{rolling } r) \leq 0.10$	AMBIGUOUS
P4 (BH count)	$\leq 2 \times \text{expected FP}$	AMBIGUOUS
F1 (Falsification)	$ r(+15 \text{ d}) \leq \text{surr. 95th}$	TRIGGERED




Figure 18: Rolling $r(+15\text{ d})$ in 18-month overlapping windows across the out-of-sample period. Error bars are bootstrap 95% confidence intervals. The grey horizontal band shows the surrogate 95th percentile. The signal shows no consistent sign or trend.

4.17 Combined 1976–2025 Analysis: Sinusoidal Modulation

Figure 19 shows the annual rolling $r(+15\text{ d})$ over the full 1976–2025 record, together with the best-fit sinusoidal envelope.




Figure 19: Annual rolling $r(+15\text{ d})$ across the full 1976–2025 period (grey points with 95% bootstrap CI). The sinusoidal best-fit (red curve, $P = 13.0\text{ yr}$) and the constant-mean model (dashed) are nearly indistinguishable; BIC comparison slightly favours the constant model ($\text{BF} = 0.75$). The vertical dashed line marks the in-sample/out-of-sample split (2020).

The global surrogate test on the full 1976–2025 window yields $p = 0.010$ ($\sigma = 2.57$) at the dominant peak $\tau = -125\text{ days}$ — nominally significant but sensitive to N_{eff} estimation and at a lag inconsistent with the claimed $+15\text{ day}$ CR precursor.

The sinusoidal fit (Equations 12–15) does *not* prefer \mathcal{M}_B over \mathcal{M}_A with the corrected seismic metric:

- Best-fit period: $P = 13.0$ years;
- Amplitude: $A = 0.047$;
- $\Delta\text{BIC} = -0.57$ (negative = \mathcal{M}_A constant preferred);
- Bayes factor: $\text{BF}_{BA} = 0.75$.

A Bayes factor of 0.75 is less than 1.0, indicating that the constant (no modulation) model is weakly preferred over the sinusoidal model (following the interpretation scale of Kass and Raftery [1995], $\text{BF} < 1$ favours \mathcal{M}_A ; note that this comparison is restricted to the two hypotheses \mathcal{M}_A and \mathcal{M}_B and does not account for model uncertainty outside this pair). With the correct seismic energy metric, the oscillatory rolling cross-correlation previously attributed to solar-cycle modulation is no longer supported. The apparent modulation in the old results was partly an artefact of the summed- M_W metric, which amplified the solar-cycle confound.

5 Discussion

5.1 Why Does the Raw Correlation Appear So Strong?

The raw $r \approx 0.31$ reported by Homola et al. [2022] at $\tau = +15$ days (reduced to $r = 0.081$ with the correct seismic energy metric) and the naive $p \sim 10^{-72}$ result from four compounding errors: (i) using a physically inappropriate seismic metric (direct summation of logarithmic M_W values), which artificially amplifies solar-cycle variation in the seismic series; (ii) treating autocorrelated time series as independent observations — statistically invalid under the violated serial-independence assumption, since autocorrelation inflates the nominal sample size by a factor of 3–5 (Table 4); (iii) failing to account for the shared ~ 11 -year solar-cycle trend driving both CR flux and seismicity; and (iv) not correcting for scanning over 401 lag values.

The solar cycle is the key confounder. During solar minimum, the heliospheric magnetic field weakens, allowing more galactic CRs to reach Earth. Independently, global seismicity has been reported to be slightly elevated during solar minimum phases [Odintsov et al., 2006]. The resulting shared ~ 11 -year oscillation in both series generates a substantial raw cross-correlation with a lag structure set by the phase relationship of their respective solar responses — approximately \pm half-cycle (~ 5.5 years $\approx 2,000$ days), consistent with the dominant raw peak at $\tau = -525$ days.

It should be noted that the ~ 10 -year periodicity shared by CR flux and seismicity need not originate exclusively from the Schwabe solar cycle. Alternative sources of shared decadal variability include geomagnetic activity cycles (themselves driven by solar activity

but mediated by different physical pathways) and long-term seismic clustering or quasi-periodic accumulation and release of tectonic stress. Any of these mechanisms could generate the observed coherence without implying a direct CR–seismic link.

5.2 Physical Plausibility of the Claimed Mechanism

Even setting aside the statistical issues, the proposed mechanism faces severe physical constraints. The total ionisation dose from galactic CRs at the surface is ~ 0.3 mGy/year [Aplin, 2006] — far too small to transfer meaningful mechanical energy to fault zones, which require shear-stress changes of order ~ 0.01 –1 MPa to trigger earthquakes. Proposed mechanisms via radon ionisation [Pulinets and Boyarchuk, 2004] or nuclear transmutation require orders-of-magnitude larger CR fluxes than observed. The geographic scan (Section 4.15) finds no propagation-delay signature ($\beta = -0.45$ d/1000 km, $p = 0.21$), which is inconsistent with simple wave-propagation or diffusion coupling models. This result does not, however, exclude mechanisms that would act instantaneously on a global scale, such as modulation of the atmospheric electric field.

5.3 Comparison with Prior Replication Attempts

Independent replication attempts of the Homola et al. [2022] result have been limited. Urata and Tanimoto [2018] found similarly inflated correlations using Japanese CR stations and reported that detrending removed most of the signal. Our analysis is the first to combine all three of: IAAFT surrogate testing, solar-cycle-aware detrending, geographic localisation scanning, and pre-registered out-of-sample validation.

5.4 Limitations

Several limitations should be acknowledged:

1. **Out-of-sample statistical power.** The OOS window (2020–2025, $T = 390$ bins, ≈ 5 yr) encompasses only the rising phase of Solar Cycle 25 and contains no complete solar minimum. Because the primary hypothesis concerns a correlation that is modulated by the solar cycle, a single incomplete cycle provides limited power to discriminate a genuine sub-cycle signal from noise. The OOS failure to replicate ($p_{\text{global}} = 0.100$) should therefore be interpreted as *consistent with* the null hypothesis rather than as strong independent evidence against the claimed effect; it carries substantially less weight than the 44-year in-sample analysis.
2. **Seismicity non-stationarity.** Seismicity is not stationary; major seismic sequences (e.g. Tonga 2022) can dominate the seismic metric in individual bins, introducing transient structure that is not related to CR flux.

3. **Sinusoidal model rigidity.** The sinusoidal modulation fit assumes a constant solar-cycle period, whereas the actual Schwabe cycle length varies from 9 to 14 years. A more flexible model (e.g. a time-warped sinusoid) might alter the BF comparison.
4. **Untested mechanisms.** This analysis tests for linear cross-correlation between the global CR index and a globally-aggregated seismic energy metric. Threshold effects (e.g. CR triggering only above a critical fault stress), nonlinear coupling, frequency-selective interaction, or extreme-event coupling are not addressed and cannot be excluded on the basis of these results.

6 Conclusions

We have conducted a rigorous, pre-registered replication of the claimed cosmic-ray/earthquake correlation from Homola et al. [2022] using 49 years of data from 44 neutron monitors, the USGS global catalogue, and SILSO sunspot numbers. Our principal findings are:

1. The raw cross-correlation $r(+15 \text{ d}) = 0.081$ (corrected seismic energy metric) is modest; it is consistent with a shared ~ 10 -year solar-cycle modulation of both CR flux and global seismicity rather than a direct CR \rightarrow seismic mechanism. The larger value ($r \approx 0.31$) reported by Homola et al. [2022] results from the physically inappropriate practice of directly summing logarithmic M_W values rather than the underlying energies.
2. After solar-cycle detrending, $r(+15 \text{ d})$ falls to 0.027 (HP), 0.030 (STL), or 0.037 (sunspot regression) — all within the surrogate null distribution and consistent with zero.
3. No propagation-delay signature is detected: the optimal lag between CR station and seismic cell shows no distance dependence ($\beta = -0.45 \text{ d}/1000 \text{ km}$, $p = 0.21$), inconsistent with local wave-propagation or diffusion models, though globally instantaneous coupling mechanisms remain untested.
4. A pre-registered out-of-sample test on 2020–2025 yields $r(+15 \text{ d}) = 0.030$ and $p_{\text{global}} = 0.100$, entirely consistent with noise.
5. The 49-year annual rolling correlation timeseries is not significantly better described by a sinusoid than a constant (Bayes factor 0.75, favoring constant); the previous sinusoidal modulation finding was an artefact of the invalid seismic metric.
6. Seven additional robustness checks all corroborate the null result:
 - Block-bootstrap surrogates ($B \approx 11 \text{ yr}$, $n = 5,000$) yield $p_{\text{CBB}}(+15 \text{ d}) = 0.022$ on the *raw* series; after detrending the result is within the null.

- Partial correlation (regressing seismic on sunspots before correlating with CR) reduces $r(+15\text{ d})$ from 0.079 to 0.029 (63% drop), confirming solar-cycle confounding without any filter.
- Mean spectral coherence in the solar-cycle band is 0.840 ($> 95\%$ threshold), while mutual information at $\tau = +15\text{ d}$ is exactly zero ($p = 1.000$): no nonlinear signal at the claimed lag.
- Missing data: 0% NaN bins at all station thresholds; $r(+15\text{ d})$ is unchanged for min-station thresholds 2, 3, or 5.
- Bin-size sensitivity: the dominant peak is at $\tau \approx -520$ days for 1-day, 5-day, and 27-day bins; r at +15 days scales with bin size as expected for solar-cycle leakage, not a physical mechanism.
- Earthquake declustering removes 28% of events as aftershocks; $r(+15\text{ d})$ changes by only $\Delta r = 0.014$ — not a material confound.
- Per-solar-cycle analysis shows inconsistent dominant peak lags (-65 , -125 , $+125$, -125 days across cycles 21–24), the signature of a phase-drifting solar-cycle artefact.

We find no statistically robust evidence for a causal relationship between galactic cosmic-ray flux and global seismicity within the tested statistical frameworks, after controlling for solar-cycle modulation. Threshold effects, nonlinear triggering, and extreme-event coupling — mechanisms not assessed in the present analysis — cannot be excluded on the basis of these results alone.

Data Availability

All analysis code, pre-registration documents, intermediate results, and figures are publicly available at <https://github.com/pingud98/cosmiccraysandearthquakes> under the MIT licence. Raw data are freely accessible from their respective providers: NMDB (<https://www.nmdb.eu>), USGS (<https://earthquake.usgs.gov/fdsnws/event/1/>), and SIDC (<https://www.sidc.be/silso/datafiles>).

Acknowledgements

The author thanks the operators of the NMDB network for maintaining open-access neutron monitor data, and the USGS Earthquake Hazards Programme for the FDSN catalogue service. GPU computations were performed on an NVIDIA Tesla M40.

References

- Karen L. Aplin. Atmospheric electrification in the solar system. *Surveys in Geophysics*, 27(1):63–108, 2006. doi: 10.1007/s10712-005-0642-9.
- M. S. Bartlett. On the theoretical specification and sampling properties of autocorrelated time-series. *Journal of the Royal Statistical Society (Supplement)*, 8(1):27–41, 1946. doi: 10.2307/2983611.
- Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B*, 57(1):289–300, 1995. doi: 10.1111/j.2517-6161.1995.tb02031.x.
- Christopher S. Bretherton, Martin Widmann, Valentin P. Dymnikov, John M. Wallace, and Ileana Blade. The effective number of spatial degrees of freedom of a time-varying field. *Journal of Climate*, 12(7):1990–2009, 1999. doi: 10.1175/1520-0442(1999)012<1990:TENOSD>2.0.CO;2.
- Robert B. Cleveland, William S. Cleveland, Jean E. McRae, and Irma Terpenning. STL: A seasonal-trend decomposition procedure based on LOESS. *Journal of Official Statistics*, 6(1):3–73, 1990.
- J. K. Gardner and Leon Knopoff. Is the sequence of earthquakes in southern California, with aftershocks removed, Poissonian? *Bulletin of the Seismological Society of America*, 64(5):1363–1367, 1974. doi: 10.1785/BSSA0640051363.
- Robert J. Hodrick and Edward C. Prescott. Postwar U.S. business cycles: an empirical investigation. *Journal of Money, Credit and Banking*, 29(1):1–16, 1997. doi: 10.2307/2953682.
- Piotr Homola et al. Indication of correlation between cosmic-ray flux and global seismicity, 2022. URL <https://arxiv.org/abs/2204.12310>.
- Hiroo Kanamori. The energy release in great earthquakes. *Journal of Geophysical Research*, 82(20):2981–2987, 1977. doi: 10.1029/JB082i020p02981.
- Robert E. Kass and Adrian E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795, 1995. doi: 10.1080/01621459.1995.10476572.
- Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Physical Review E*, 69(6):066138, 2004. doi: 10.1103/PhysRevE.69.066138.
- Stanislav Odintsov, Kirill Boyarchuk, Katya Georgieva, Boian Kirov, and Dimitar Atanasov. Long-period trends in global seismic and geomagnetic activity and their relation to solar

activity. *Physics and Chemistry of the Earth*, 31(1–3):88–93, 2006. doi: 10.1016/j.pce.2005.03.004.

Marius S. Potgieter. Solar modulation of cosmic rays. *Living Reviews in Solar Physics*, 10(1):3, 2013. doi: 10.12942/lrsp-2013-3.

Sergey Pulinetz and Kirill Boyarchuk. Ionospheric precursors of earthquakes. *Springer*, 2004.

Morten O. Ravn and Harald Uhlig. On adjusting the Hodrick–Prescott filter for the frequency of observations. *Review of Economics and Statistics*, 84(2):371–376, 2002. doi: 10.1162/003465302317411604.

Thomas Schreiber and Andreas Schmitz. Surrogate time series. *Physica D: Nonlinear Phenomena*, 142(3–4):346–382, 2000. doi: 10.1016/S0167-2789(00)00043-9.

SILSO World Data Center. The international sunspot number. Royal Observatory of Belgium, Brussels. <https://www.sidc.be/silso/datafiles>, 2024.

Eliyahu Stoupel. Relationship between solar and seismic activity. *International Journal of Biometeorology*, 34(4):231–235, 1990. doi: 10.1007/BF01049646.

Mario Tavares and Angelo Azevedo. Influences of solar cycles on earthquakes. *Pattern Recognition in Physics*, 1(1):1–11, 2011.

James Theiler, Stephen Eubank, André Longtin, Bryan Galdrikian, and J. Doyne Farmer. Testing for nonlinearity in time series: the method of surrogate data. *Physica D: Nonlinear Phenomena*, 58(1–4):77–94, 1992. doi: 10.1016/0167-2789(92)90102-S.

Naoyuki Urata and Toshiro Tanimoto. Correlation between cosmic rays and seismicity: a case study in Japan. *Earth, Planets and Space*, 70(1):55, 2018. doi: 10.1186/s40623-018-0826-5.

USGS Earthquake Hazards Program. Earthquake catalogue via FDSN web service. <https://earthquake.usgs.gov/fdsnws/event/1/>, 2024.