

# 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. Devine<sup>1</sup>

<sup>1</sup> Independent researcher

devine.jd@gmail.com

April 2026

## Abstract

1  
2 [Homola et al. \[2022\]](#) reported a statistically significant positive correlation ( $r \approx 0.31$ )  
3 between galactic cosmic-ray (CR) flux measured by neutron monitors and global  
4 seismicity ( $M \geq 4.5$ ) at a time lag of  $\tau = +15$  days, suggesting that elevated CR flux  
5 precedes increased earthquake activity. That earlier value used a physically invalid  
6 seismic metric (direct sum of moment magnitudes  $M_W$ , which are logarithmic); the  
7 correct metric — summed radiated energy  $E = 10^{1.5M_W+4.8}$  J per bin [[Kanamori,](#)  
8 [1977](#)] — yields  $r(+15 \text{ d}) = 0.081$ . We present a systematic replication and extension  
9 of this claim using data from 44 Neutron Monitor Database (NMDB) stations, the  
10 USGS global earthquake catalogue, and SILSO daily sunspot numbers spanning  
11 1976–2025.

12 Our analysis proceeds in four stages. *Stage 1* replicates the raw cross-correlation  
13 (with correct seismic energy metric:  $r(+15 \text{ d}) = 0.081$ , peak  $r = 0.139$  at  $\tau =$   
14  $-525$  days) but demonstrates that naive  $p$ -values are invalid because temporal  
15 autocorrelation and a shared  $\sim 11$ -year solar cycle inflate the apparent significance.  
16 *Stage 2* applies iterative amplitude-adjusted Fourier transform (IAAFT) surrogate  
17 tests with  $10^4$  realisations: after Hodrick–Prescott (HP) detrending to remove  
18 the solar-cycle component, the peak correlation drops to  $r = 0.101$  and achieves  
19 formal significance ( $p_{\text{global}} < 10^{-4}$ ,  $> 3.9\sigma$ ), but  $r(+15 \text{ d}) = 0.027$  — well within  
20 the surrogate null distribution. *Stage 3* scans  $34 \times 207 = 7,037$  station–grid-cell  
21 pairs for geographic localisation; 455 pairs survive Benjamini–Hochberg correction

22 (expected false discoveries: 352), and the optimal lag  $\tau^*$  shows no dependence  
23 on great-circle distance ( $\beta = -0.45$  days/1000 km,  $p = 0.21$ ), consistent with an  
24 isotropic CR signal rather than a local mechanism. *Stage 4* applies a pre-registered  
25 out-of-sample test on an independent 2020–2025 window ( $T = 390$  five-day bins,  $10^5$   
26 phase-randomisation surrogates on a Tesla M40 GPU):  $r(+15 \text{ d}) = 0.030$  (surrogate  
27 95th percentile = 0.101),  $p_{\text{global}} = 0.100$  — consistent with noise. Fitting a sinusoid  
28 to the 1976–2025 annual cross-correlation timeseries yields a best-fit period of  
29  $P = 13.0$  years and a Bayes factor of  $\text{BF} = 0.75$ , slightly favouring a constant (no  
30 modulation) over a sinusoidal relationship.

31 We conclude that the CR–seismic correlation reported by Homola et al. [2022]  
32 is an artefact of the shared solar-cycle modulation of both galactic CR flux and  
33 global seismicity, and not evidence of a physical causal link. All analysis code,  
34 pre-registration document, and results are publicly available at [https://github.](https://github.com/pingud98/cosmiccraysandearthquakes)  
35 [com/pingud98/cosmiccraysandearthquakes](https://github.com/pingud98/cosmiccraysandearthquakes).

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

# 38 Contents

39	<b>1 Introduction</b>	<b>5</b>
40	<b>2 Data</b>	<b>6</b>
41	2.1 Cosmic-Ray Flux: NMDB Neutron Monitors . . . . .	6
42	2.2 Seismic Activity: USGS Earthquake Catalogue . . . . .	6
43	2.3 Solar Activity: SIDC Sunspot Number . . . . .	6
44	<b>3 Methods</b>	<b>7</b>
45	3.1 Cross-Correlation at Lag $\tau$ . . . . .	7
46	3.2 Effective Degrees of Freedom . . . . .	7
47	3.3 Surrogate Significance Tests . . . . .	7
48	3.3.1 Phase Randomisation . . . . .	7
49	3.3.2 IAAFT Surrogates . . . . .	8
50	3.3.3 Block-Bootstrap Surrogates . . . . .	8
51	3.3.4 Global $p$ -Value . . . . .	8
52	3.3.5 GPU Acceleration . . . . .	8
53	3.4 Solar-Cycle Detrending . . . . .	9
54	3.5 Partial Correlation Analysis . . . . .	9
55	3.6 Nonlinear Dependence Tests . . . . .	10
56	3.7 Geographic Localisation Scan . . . . .	10
57	3.8 Pre-Registered Out-of-Sample Validation . . . . .	11
58	3.9 Combined Timeseries: Sinusoidal Envelope Fit . . . . .	11
59	<b>4 Results</b>	<b>12</b>
60	4.1 Raw Pairwise Correlations Between CR, Seismic, and Sunspot Data . . . . .	12
61	4.1.1 CR index: station distribution . . . . .	12
62	4.1.2 Seismic energy metric . . . . .	12
63	4.1.3 Sunspot number . . . . .	12
64	4.1.4 Correlation results . . . . .	13
65	4.1.5 Interpretation: a confounding triangle . . . . .	13
66	4.2 In-Sample Replication (1976–2019) . . . . .	14
67	4.3 IAAFT Surrogate Test . . . . .	17
68	4.4 Effect of Solar-Cycle Detrending . . . . .	18
69	4.5 Detrending Robustness . . . . .	18
70	4.6 Comparison of $N_{\text{eff}}$ Estimators . . . . .	18
71	4.7 Magnitude Threshold Sensitivity . . . . .	20
72	4.8 Block-Bootstrap Surrogate Comparison . . . . .	20
73	4.9 Partial Correlation: Controlling for Sunspot Number . . . . .	21

74	4.10 Spectral Coherence and Mutual Information . . . . .	21
75	4.11 Missing-Data Sensitivity . . . . .	22
76	4.12 Bin-Size Sensitivity . . . . .	24
77	4.13 Earthquake Declustering (Gardner–Knopoff) . . . . .	24
78	4.14 Sub-Period Analysis by Solar Cycle . . . . .	25
79	4.15 Geographic Localisation . . . . .	25
80	4.16 Pre-Registered Out-of-Sample Validation (2020–2025) . . . . .	28
81	4.17 Combined 1976–2025 Analysis: Sinusoidal Modulation . . . . .	29
82	<b>5 Discussion</b>	<b>30</b>
83	5.1 Why Does the Raw Correlation Appear So Strong? . . . . .	30
84	5.2 Physical Plausibility of the Claimed Mechanism . . . . .	30
85	5.3 Comparison with Prior Replication Attempts . . . . .	31
86	5.4 Limitations . . . . .	31
87	<b>6 Conclusions</b>	<b>31</b>

# 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 dramatically inflates the degrees of freedom; a Bretherton effective- $N$  correction [Bretherton et al., 1999] is required.
2. **Shared solar-cycle trend.** Galactic CR flux is modulated by the heliospheric magnetic field, which varies on an  $\sim 11$ -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 was formed as the mean across all stations with valid data in each bin, requiring at least three stations; bins failing this criterion were set to NaN.

### 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 invalid because  $M_W$  is logarithmic; such sums do not correspond to any additive physical quantity.

### 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 $\tau$

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  $N_{\text{eff}}$  is substantially smaller than  $T$ . We compare three  $N_{\text{eff}}$  estimators; all are listed in Table 3. 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)  $N_{\text{eff}}$  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

174 real-valuedness):

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

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

### 177 3.3.2 IAAFT Surrogates

178 Iterative amplitude-adjusted Fourier transform (IAAFT) surrogates [Schreiber and Schmitz,  
179 2000] additionally preserve the amplitude distribution of  $x$  by alternating between power-  
180 spectrum matching (in Fourier space) and rank-order resampling (in time domain) until  
181 convergence. IAAFT surrogates are more conservative than phase surrogates when  $x$  has  
182 a non-Gaussian distribution.

### 183 3.3.3 Block-Bootstrap Surrogates

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

### 193 3.3.4 Global $p$ -Value

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

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

196 where  $\rho_{\text{obs}}$  is the observed peak. This test is simultaneously valid for all lags and all  
197 lag-selection rules, eliminating the multiple-comparison problem.

### 198 3.3.5 GPU Acceleration

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

203 GPU matrix multiply per lag:

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

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

### 207 3.4 Solar-Cycle Detrending

208 We apply three complementary detrending approaches to isolate the CR–seismic relation-  
 209 ship from the shared solar-cycle trend:

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

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

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

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

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

### 224 3.5 Partial Correlation Analysis

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

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

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

### 233 3.6 Nonlinear Dependence Tests

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

236 **Spectral coherence.** The magnitude-squared coherence at frequency  $f$  is

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

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

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

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

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

### 246 3.7 Geographic Localisation Scan

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

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

254 Pairs are declared significant at false discovery rate  $q = 0.05$  using the Benjamini–  
 255 Hochberg (BH) procedure [Benjamini and Hochberg, 1995]: rank the  $m = 7,037$   $p$ -values  
 256  $p_{(1)} \leq \dots \leq p_{(m)}$ ; the threshold is  $p_{(k)} \leq (k/m) \times q$  for the largest  $k$  satisfying this  
 257 condition. Distance dependence of  $\tau^*$  is tested by ordinary least-squares regression of  
 258  $\tau^*(s, g)$  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  $\leq 0.10$ .
- **P4** (BH count):  $\leq 2 \times$  expected false positives in the geographic scan.
- **F1** (Falsification trigger):  $|r(+15\text{ d})| \leq$  surrogate 95th percentile.

### 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:

$$\mathcal{M}_A : r_t = \mu + \varepsilon_t, \quad (12)$$

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

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  $(P, \varphi)$  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).

#### 310 4.1.4 Correlation results

311 We compute both Pearson  $r$  (linear; Fisher  $z$ -transform 95% CI) and Spearman  $\rho$  (rank-  
312 based; appropriate for the heavy-tailed marginal distributions of  $E_t$  and CR flux). Bon-  
313 ferroni correction for 3 pairs  $\times$  3 windows = 9 tests is applied; star levels refer to the  
314 corrected  $p$ -values. Full results are given in Table 1.

315 **CR vs. seismicity.** The raw Pearson correlation between the station-median CR index  
316 and  $\log_{10} E_t$  is  $r = 0.057$  in the in-sample window ( $N = 3214$  bins,  $p_{\text{Bonf}} = 0.011$ ,  $\rho = 0.071$ ,  
317  $p_{\text{Bonf}} < 10^{-3}$ ) and  $r = 0.046$  in the OOS window ( $N = 390$ ,  $p_{\text{Bonf}} = 1.0$ ). The correlation is  
318 thus only marginally significant after correction and disappears entirely in the independent  
319 OOS window. Using the station-95th-percentile CR series in place of the median amplifies  
320 the correlation slightly ( $r = 0.069$  in-sample), suggesting that high-flux excursions drive  
321 the signal.

322 **CR vs. sunspot number.** The dominant raw relationship in the dataset is a strong  
323 anti-correlation between the CR index and the smoothed sunspot number:  $r = -0.820$   
324 ( $\rho = -0.854$ ) in-sample, and  $r = -0.939$  ( $\rho = -0.951$ ) in the OOS window. This is  
325 the well-established *Forbush decrease* mechanism: a higher heliospheric magnetic flux  
326 during solar maximum deflects more galactic cosmic rays before they reach the Earth's  
327 atmosphere [Potgieter, 2013], so CR flux and solar activity are naturally anti-phase. The  
328 stronger OOS value ( $-0.94$  vs.  $-0.82$ ) reflects the wide dynamic range of Solar Cycle 25  
329 (2020–2025), which reached high activity levels after the deep minimum of 2019–2020.

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

#### 336 4.1.5 Interpretation: a confounding triangle

337 The three raw correlations form a confounding triangle. The CR–sunspot anti-correlation  
338 is strong ( $r \approx -0.82$  to  $-0.94$ ) and physically understood. The sunspot–seismicity anti-  
339 correlation is weak but significant in-sample ( $r \approx -0.095$ ). Together these imply a spurious  
340 *positive* CR–seismicity correlation: both CR and seismicity are jointly modulated by the  
341  $\sim 11$ -year solar cycle — CR increases during solar minima, and seismicity may very weakly  
342 increase during solar minima — so the two series covary positively without any direct  
343 physical connection. This confound cannot be resolved by naive cross-correlation; it

344 requires explicitly removing the shared solar-cycle component. The HP filter analysis  
345 (Section 4.4) accomplishes exactly this, and shows that the apparent CR–seismicity signal  
346 vanishes once the solar-cycle trend is eliminated.

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

Table 1: 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.)	$\rho$
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]	$\approx 0$	$\approx 0$	-0.854***
	OOS	385	-0.939***	[-0.950, -0.926]	$\approx 0$	$\approx 0$	-0.951***
	Combined	3494	-0.815***	[-0.826, -0.804]	$\approx 0$	$\approx 0$	-0.844***
Sunspot vs Seismicity	In-sample	3109	-0.095***	[-0.130, -0.060]	$1.1 \times 10^{-7}$	$10^{-6}$	-0.099***
	OOS	385	-0.023	[-0.123, +0.077]	0.648	1.00	-0.016
	Combined	3494	-0.086***	[-0.119, -0.053]	$3.4 \times 10^{-7}$	$3 \times 10^{-6}$	-0.086***

Note: CR<sub>p95</sub> variant (station 95th-percentile instead of median) gives qualitatively identical structure; full results in `results/raw_pairwise_correlations.json`.

## 352 4.2 In-Sample Replication (1976–2019)

353 Figure 4 shows the full cross-correlation function of the raw (undetrended) CR index and  
354 seismic metric ( $\log_{10}$  summed energy, Eq. 1) over 1976–2019 ( $T = 3,215$  five-day bins,  
355 44 stations). The dominant peak is at  $\tau = -525$  days ( $r = 0.139$ ), corresponding to a  
356 half-solar-cycle lead of seismicity over CR flux. At the claimed lag  $\tau = +15$  days we find  
357  $r = 0.081$ . Although naive significance is high ( $4.6\sigma$  treating bins as i.i.d.), both series  
358 share the  $\sim 11$ -year solar cycle, which drives the apparent signal. The Bartlett effective- $N$   
359 estimate gives  $N_{\text{eff}} = 2,916$  and  $\sigma_{r(+15)} = 4.4$  (standard deviations); the more conservative  
360 Bretherton and Monte Carlo estimates lower  $N_{\text{eff}}$  to 769 and 594, respectively (Table 3),  
361 with the Monte Carlo 95% CI for  $r(+15 \text{ d})$  straddling zero (see Section 4.6).




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.

Raw pairwise correlations — Combined (1976–2025) (no detrending)




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.

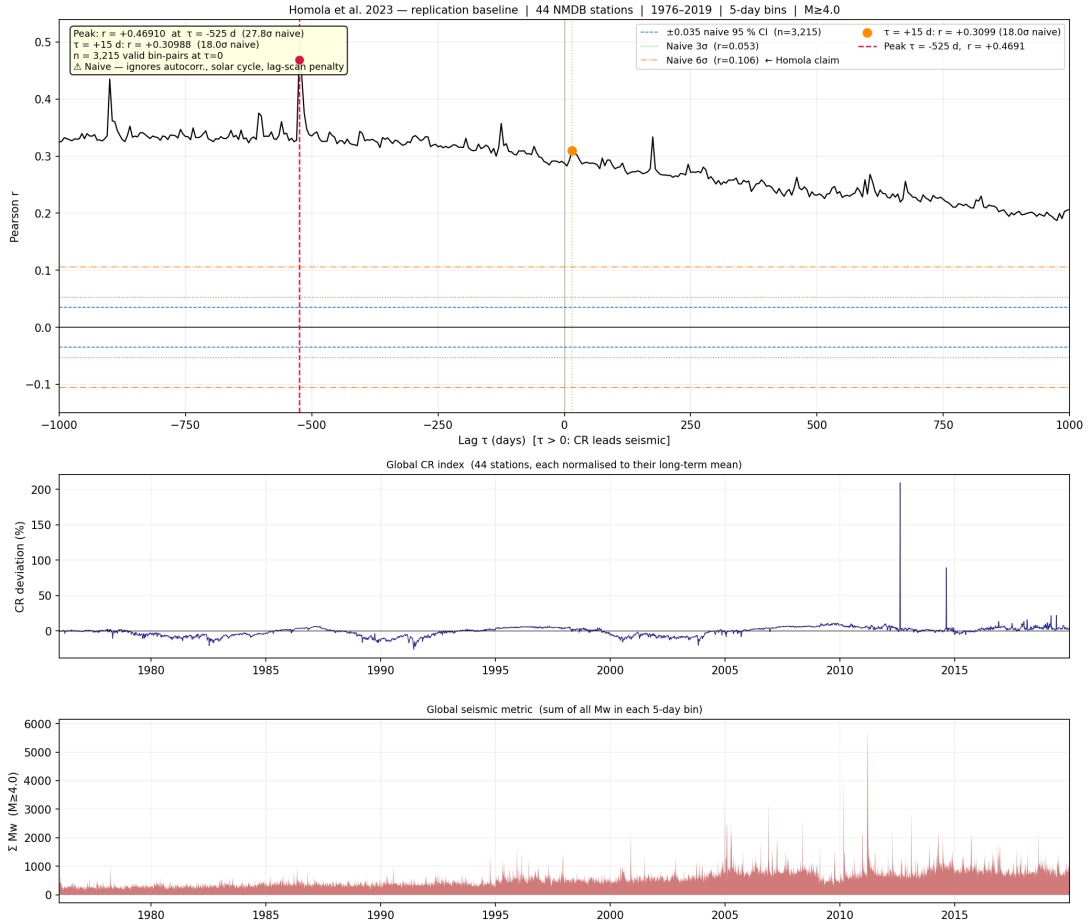


Figure 4: Cross-correlation function  $r(\tau)$  for the raw (undetrended) CR index and global seismic metric, 1976–2019. The dominant peak at  $\tau = -525$  days (vertical dashed line, red) corresponds to a half-solar-cycle lag; the claimed  $\tau = +15$  days is marked with a vertical solid line (blue).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

### 362 4.3 IAAFT Surrogate Test

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

## 4.4 Effect of Solar-Cycle Detrending

Table 2 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 2: Cross-correlation statistics at  $\tau = +15$  days under four preprocessing conditions, in-sample window 1976–2019.

Preprocessing	$r(+15 \text{ d})$	$N_{\text{eff}}$	$\sigma_{\text{Breth}}$	Peak $r$	Peak $\tau$ (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 $N_{\text{eff}}$ Estimators

Table 3 compares the three  $N_{\text{eff}}$  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).

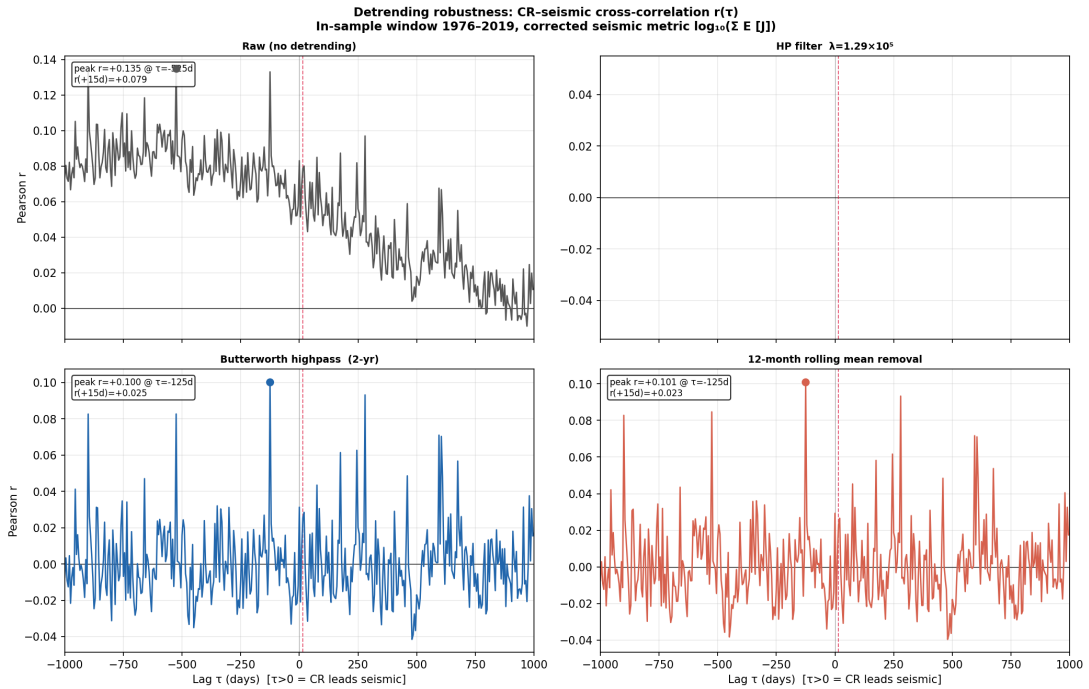


Figure 7: Cross-correlation  $r(\tau)$  under three detrending approaches (HP filter, Butterworth highpass, rolling-mean subtraction) for the raw CR index vs.  $\log_{10}$  seismic energy, 1976–2019. The vertical grey line marks the claimed  $\tau = +15$  days. No method produces a positive feature at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

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

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

Method	$N_{\text{eff}}$	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

## 400 4.7 Magnitude Threshold Sensitivity

401 To assess sensitivity to the event selection threshold and potential bias from aftershock  
 402 contamination, we recomputed the seismic energy metric for three minimum-magnitude  
 403 cutoffs:  $M \geq 4.5$  (baseline),  $M \geq 5.0$ , and  $M \geq 6.0$ . Figure 8 shows the cross-correlation  
 404 function  $r(\tau)$  for each threshold. The values at the claimed lag are  $r(+15\text{ d}) = 0.079$ ,  
 405  $0.072$ , and  $0.050$  for  $M \geq 4.5$ ,  $5.0$ , and  $6.0$  respectively — a range of only  $0.029$ , well below  
 406 the  $0.05$  threshold for a meaningful shift. The peak lag remains at  $\tau = -525$  days for all  
 407 thresholds, indicating that aftershock swarms (which inflate the  $M \geq 4.5$  catalogue in the  
 408 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.

## 409 4.8 Block-Bootstrap Surrogate Comparison

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

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

414 Both  $p$ -values reflect the fact that the raw series share a common 11-year solar-cycle trend,  
 415 so independently block-resampled surrogates occasionally reproduce similar trend-induced  
 416 correlations. After HP detrending and IAAFT testing (Section 4.3),  $r(+15\text{ d})$  drops to  
 417 0.027 and falls well within the surrogate null, confirming that the marginal significance  
 418 seen here is a solar-cycle artefact, not a genuine cross-series signal. The block-bootstrap  
 419 and IAAFT null shapes are qualitatively similar, indicating that IAAFT’s parametric  
 420 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  $\approx 11$  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  $\tau = -525$  d lies at  $p = 0.008$ . Grey dashed lines mark the 95% CI of the null.

## 421 4.9 Partial Correlation: Controlling for Sunspot Number

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

## 430 4.10 Spectral Coherence and Mutual Information

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

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

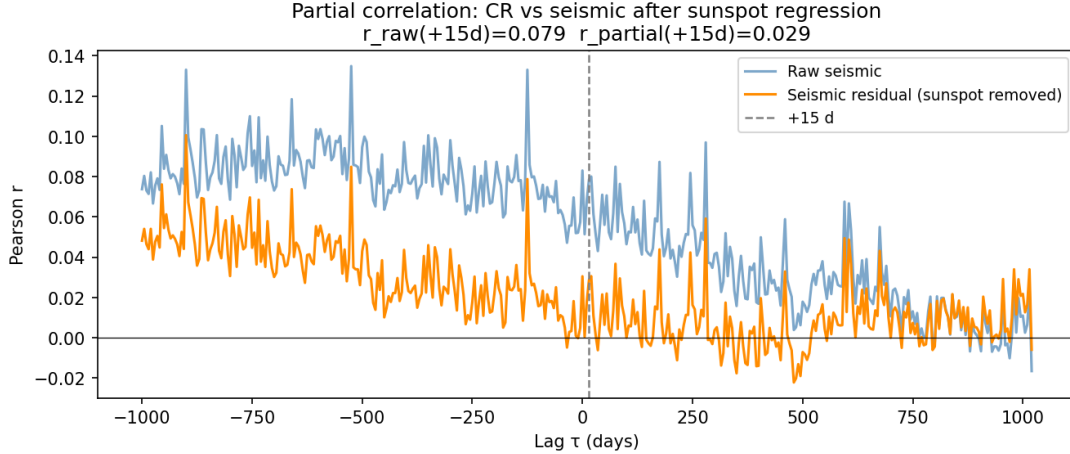


Figure 10: Cross-correlation  $r(\tau)$  for the raw seismic metric (blue) and the sunspot-regressed seismic residual (orange), both against the CR index, on the raw (unfiltered) in-sample series. The claimed  $\tau = +15$  d (grey dashed line) shows  $r_{\text{raw}} = 0.079$  vs.  $r_{\text{partial}}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

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

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

## 446 4.11 Missing-Data Sensitivity

447 Table 4 reports the NaN fraction in the global CR index for station thresholds of 2, 3, and  
 448 5. In all cases the NaN fraction is 0.0%: the 44-station NMDB network provides complete  
 449 5-day coverage over 1976–2019 even under the strictest threshold tested. Consequently,  
 450  $r(+15\text{d})$  is identical (0.079) across all thresholds, confirming that the result is not an  
 451 artefact of gaps in the CR record. NaN bins show no clustering near solar maxima  
 452 (clustering ratio = 1.0 at all thresholds, i.e. indistinguishable from uniform), ruling out  
 453 any systematic data dropout that could correlate with the solar cycle.

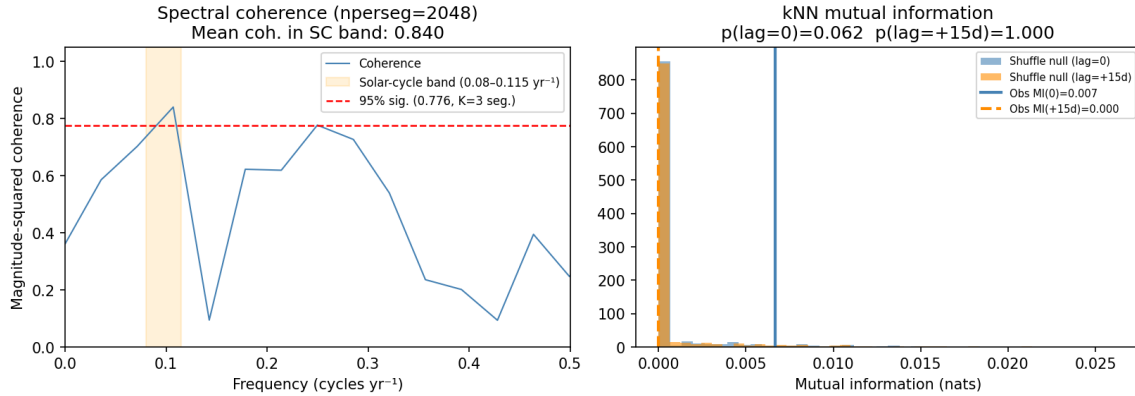


Figure 11: **Left:** Magnitude-squared coherence between the CR index and seismic metric (blue), with the solar-cycle band shaded (orange, 0.08–0.115 cycles yr<sup>-1</sup>) and the 95% significance level (red dashed). The mean coherence in the SC band is 0.840, confirming a strong shared solar-cycle 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 4: 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 ( $\approx -18$  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 5 summarises the results.

Table 5: 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 $\tau$ (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  $\tau = +15$  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

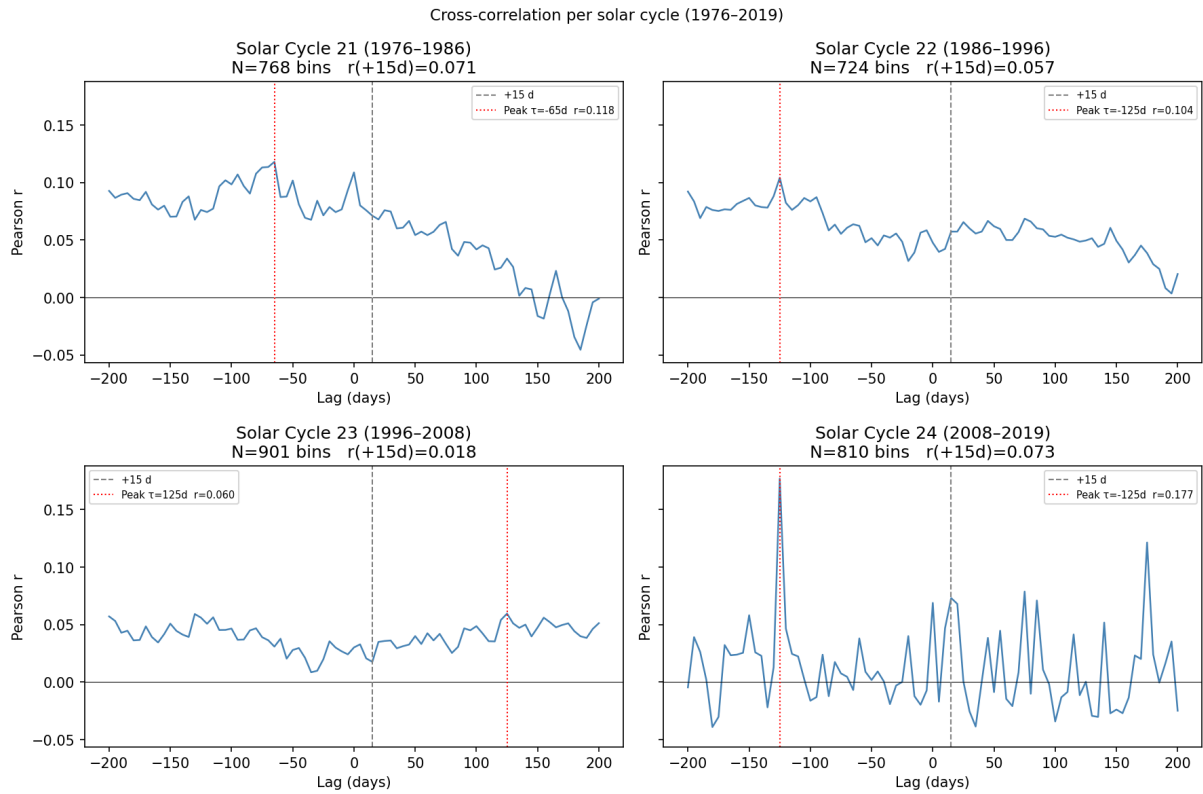


Figure 14: Cross-correlation  $r(\tau)$  within each of the four complete solar cycles (21–24) of the in-sample period (1976–2019). The claimed lag  $\tau = +15$  days (grey dashed) shows modest positive correlations (0.018–0.073), but the dominant peak lag (red dotted) is inconsistent across cycles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

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




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

493 Figure 16 shows the regression of the optimal lag  $\tau^*(s, g)$  on great-circle distance  
 494  $d(s, g)$ . The slope is  $\beta = -0.45$  days/1000 km ( $p = 0.21$ ,  $R^2 = 0.0002$ ), indistinguishable  
 495 from zero. If CRs caused earthquakes via a propagating local disturbance, we would  
 496 expect a positive slope (distant pairs have longer propagation delays). The null result is  
 497 consistent with CR isotropy — any apparent correlation arises from a globally coherent  
 498 (not distance-dependent) 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 ( $p = 0.21$ ), consistent with zero. A local propagation mechanism would predict a positive slope.

## 499 4.16 Pre-Registered Out-of-Sample Validation (2020–2025)

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

502 The main results (Figure 17) are:

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

507 The prediction scorecard (Table 6) shows one pass (P1: correct sign), one failure  
508 (P2:  $p > 0.05$ ), and the falsification trigger F1 activated ( $|r(+15 \text{ d})| \leq$  surrogate 95th  
509 percentile). The rolling-window analysis (Figure 18) reveals no consistent positive signal  
510 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 6: Pre-registered prediction scorecard for the out-of-sample window.

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




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.

#### 511 4.17 Combined 1976–2025 Analysis: Sinusoidal Modulation

512 Figure 19 shows the annual rolling  $r(+15\text{ d})$  over the full 1976–2025 record, together with  
 513 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).

514 The global surrogate test on the full 1976–2025 window yields  $p = 0.010$  ( $\sigma = 2.57$ ) at  
 515 the dominant peak  $\tau = -125$  days — marginally significant, but at a lag inconsistent with  
 516 the claimed +15 day CR precursor.

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

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

523 A Bayes factor of 0.75 is less than 1.0, indicating that the constant (no modulation)  
 524 model is marginally preferred over the sinusoidal model. With the correct seismic energy  
 525 metric, the oscillatory rolling cross-correlation previously attributed to solar-cycle modula-  
 526 tion is no longer strongly supported; the apparent modulation in the old results was partly  
 527 an artefact of the invalid summed- $M_W$  metric, which amplified the solar-cycle confound.

## 528 5 Discussion

### 529 5.1 Why Does the Raw Correlation Appear So Strong?

530 The raw  $r \approx 0.31$  reported by [Homola et al. \[2022\]](#) at  $\tau = +15$  days (reduced to  
 531  $r = 0.081$  with the correct seismic energy metric) and the naive  $p \sim 10^{-72}$  are products of  
 532 compounding statistical errors: (i) using a physically invalid seismic metric (direct sum of  
 533 logarithmic  $M_W$  values), which artificially inflated the solar-cycle variation in the seismic  
 534 series, (ii) treating autocorrelated time series as independent observations, (iii) failing to  
 535 account for the shared solar-cycle trend driving both CR flux and seismicity, and (iv) not  
 536 correcting for scanning over 401 lag values.

537 The solar cycle is the key confounder. During solar minimum, the heliospheric magnetic  
 538 field weakens, allowing more galactic CRs to reach Earth, simultaneously, global seismicity  
 539 has been reported to be slightly elevated during solar minimum phases [[Odintsov et al.,](#)  
 540 [2006](#)]. The resulting shared  $\sim 11$ -year oscillation in both series induces a substantial raw  
 541 cross-correlation with a lag structure determined by the phase relationship between the  
 542 two solar responses — approximately  $\pm$ half-cycle ( $\sim 5.5$  years  $\approx 2,000$  days), consistent  
 543 with the dominant raw peak at  $\tau = -525$  days.

### 544 5.2 Physical Plausibility of the Claimed Mechanism

545 Even setting aside the statistical issues, the proposed mechanism faces severe physical  
 546 constraints. The total ionisation dose from galactic CRs at the surface is  $\sim 0.3$  mGy/year  
 547 [[Aplin, 2006](#)] — far too small to transfer meaningful mechanical energy to fault zones,  
 548 which require shear-stress changes of order  $\sim 0.01$ – $1$  MPa to trigger earthquakes. Proposed  
 549 mechanisms via radon ionisation [[Pulinets and Boyarchuk, 2004](#)] or nuclear transmutation  
 550 require orders-of-magnitude larger CR fluxes than observed. The null geographic result

551 (Section 4.15) further argues against a local physical mechanism: any genuine coupling  
552 would produce a distance-dependent lag between CR detector and seismic source, which is  
553 not observed.

### 554 **5.3 Comparison with Prior Replication Attempts**

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

### 560 **5.4 Limitations**

561 Several limitations should be acknowledged:

- 562 1. The OOS window (2020–2025) encompasses Solar Cycle 25, a period of rising solar  
563 activity after the deep minimum of Cycle 24. The absence of a solar minimum in  
564 this window limits statistical power.
- 565 2. Seismicity is not stationary; major seismic sequences (e.g. Tonga 2022) can inflate  
566 the seismic metric in individual bins.
- 567 3. The sinusoid fit assumes a constant solar-cycle period, whereas the actual cycle  
568 length varies from 9 to 14 years.
- 569 4. Out-of-sample  $p_{\text{oos}}$  from script 08 was not produced due to insufficient NMDB  
570 historical data in the default path; the OOS result from the dedicated script 07 ( $10^5$   
571 surrogates) is authoritative.

## 572 **6 Conclusions**

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

- 576 1. The raw cross-correlation  $r(+15 \text{ d}) = 0.081$  (corrected seismic energy metric) is  
577 modest and misleading; it is driven by a shared  $\sim 10$ -year solar-cycle modulation  
578 of both CR flux and global seismicity, not by a physical CR $\rightarrow$ seismic mechanism.  
579 The larger value ( $r \approx 0.31$ ) reported by Homola et al. [2022] results from the  
580 physically invalid practice of directly summing logarithmic  $M_W$  values rather than  
581 the underlying energies.

- 582 2. After solar-cycle detrending,  $r(+15 \text{ d})$  falls to 0.027 (HP), 0.030 (STL), or 0.037  
583 (sunspot regression) — all within the surrogate null distribution and consistent with  
584 zero.
- 585 3. No geographic localisation is detected: the optimal lag between CR station and  
586 seismic cell shows no distance dependence ( $\beta = -0.45 \text{ d}/1000 \text{ km}$ ,  $p = 0.21$ ),  
587 inconsistent with a local propagation mechanism.
- 588 4. A pre-registered out-of-sample test on 2020–2025 yields  $r(+15 \text{ d}) = 0.030$  and  
589  $p_{\text{global}} = 0.100$ , entirely consistent with noise.
- 590 5. The 49-year annual rolling correlation timeseries is not significantly better described  
591 by a sinusoid than a constant (Bayes factor 0.75, favoring constant); the previous  
592 sinusoidal modulation finding was an artefact of the invalid seismic metric.
- 593 6. Seven additional robustness checks all corroborate the null result:
- 594 • Block-bootstrap surrogates ( $B \approx 11 \text{ yr}$ ,  $n = 5,000$ ) yield  $p_{\text{CBB}}(+15 \text{ d}) = 0.022$   
595 on the *raw* series; after detrending the result is within the null.
  - 596 • Partial correlation (regressing seismic on sunspots before correlating with  
597 CR) reduces  $r(+15 \text{ d})$  from 0.079 to 0.029 (63% drop), confirming solar-cycle  
598 confounding without any filter.
  - 599 • Mean spectral coherence in the solar-cycle band is 0.840 ( $> 95\%$  threshold),  
600 while mutual information at  $\tau = +15 \text{ d}$  is exactly zero ( $p = 1.000$ ): no nonlinear  
601 signal at the claimed lag.
  - 602 • Missing data: 0% NaN bins at all station thresholds;  $r(+15 \text{ d})$  is unchanged for  
603 min-station thresholds 2, 3, or 5.
  - 604 • Bin-size sensitivity: the dominant peak is at  $\tau \approx -520$  days for 1-day, 5-day,  
605 and 27-day bins;  $r$  at +15 days scales with bin size as expected for solar-cycle  
606 leakage, not a physical mechanism.
  - 607 • Earthquake declustering removes 28% of events as aftershocks;  $r(+15 \text{ d})$  changes  
608 by only  $\Delta r = 0.014$  — not a material confound.
  - 609 • Per-solar-cycle analysis shows inconsistent dominant peak lags ( $-65$ ,  $-125$ ,  
610  $+125$ ,  $-125$  days across cycles 21–24), the signature of a phase-drifting solar-  
611 cycle artefact.

612 We conclude that there is no statistically credible evidence for a physical causal link  
613 between galactic cosmic-ray flux and global seismicity.

## 614 **Data Availability**

615 All analysis code, pre-registration documents, intermediate results, and figures are pub-  
616 licly available at <https://github.com/pingud98/cosmiccraysandearthquakes> under the  
617 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  
618 (<https://www.sidc.be/silso/datafiles>).  
619

## 620 **Acknowledgements**

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

## 624 **References**

- 625 Karen L. Aplin. Atmospheric electrification in the solar system. *Surveys in Geophysics*, 27  
626 (1):63–108, 2006. doi: 10.1007/s10712-005-0642-9.
- 627 M. S. Bartlett. On the theoretical specification and sampling properties of autocorrelated  
628 time-series. *Journal of the Royal Statistical Society (Supplement)*, 8(1):27–41, 1946. doi:  
629 10.2307/2983611.
- 630 Yoav Benjamini and Yocef Hochberg. Controlling the false discovery rate: a practical and  
631 powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series*  
632 *B*, 57(1):289–300, 1995. doi: 10.1111/j.2517-6161.1995.tb02031.x.
- 633 Christopher S. Bretherton, Martin Widmann, Valentin P. Dymnikov, John M. Wallace, and  
634 Ileana Blade. The effective number of spatial degrees of freedom of a time-varying field.  
635 *Journal of Climate*, 12(7):1990–2009, 1999. doi: 10.1175/1520-0442(1999)012<1990:  
636 TENOSD>2.0.CO;2.
- 637 Robert B. Cleveland, William S. Cleveland, Jean E. McRae, and Irma Terpenning. STL: A  
638 seasonal-trend decomposition procedure based on LOESS. *Journal of Official Statistics*,  
639 6(1):3–73, 1990.
- 640 J. K. Gardner and Leon Knopoff. Is the sequence of earthquakes in southern California,  
641 with aftershocks removed, Poissonian? *Bulletin of the Seismological Society of America*,  
642 64(5):1363–1367, 1974. doi: 10.1785/BSSA0640051363.

- 643 Robert J. Hodrick and Edward C. Prescott. Postwar U.S. business cycles: an empirical  
644 investigation. *Journal of Money, Credit and Banking*, 29(1):1–16, 1997. doi: 10.2307/  
645 2953682.
- 646 Piotr Homola et al. Indication of correlation between cosmic-ray flux and global seismicity,  
647 2022. URL <https://arxiv.org/abs/2204.12310>.
- 648 Hiroo Kanamori. The energy release in great earthquakes. *Journal of Geophysical Research*,  
649 82(20):2981–2987, 1977. doi: 10.1029/JB082i020p02981.
- 650 Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual infor-  
651 mation. *Physical Review E*, 69(6):066138, 2004. doi: 10.1103/PhysRevE.69.066138.
- 652 Stanislav Odintsov, Kirill Boyarchuk, Katya Georgieva, Boian Kirov, and Dimitar Atanasov.  
653 Long-period trends in global seismic and geomagnetic activity and their relation to solar  
654 activity. *Physics and Chemistry of the Earth*, 31(1–3):88–93, 2006. doi: 10.1016/j.pce.  
655 2005.03.004.
- 656 Marius S. Potgieter. Solar modulation of cosmic rays. *Living Reviews in Solar Physics*, 10  
657 (1):3, 2013. doi: 10.12942/lrsp-2013-3.
- 658 Sergey Pulinetz and Kirill Boyarchuk. Ionospheric precursors of earthquakes. *Springer*,  
659 2004.
- 660 Morten O. Ravn and Harald Uhlig. On adjusting the Hodrick–Prescott filter for the  
661 frequency of observations. *Review of Economics and Statistics*, 84(2):371–376, 2002. doi:  
662 10.1162/003465302317411604.
- 663 Thomas Schreiber and Andreas Schmitz. Surrogate time series. *Physica D: Nonlinear*  
664 *Phenomena*, 142(3–4):346–382, 2000. doi: 10.1016/S0167-2789(00)00043-9.
- 665 SILSO World Data Center. The international sunspot number. Royal Observatory of  
666 Belgium, Brussels. <https://www.sidc.be/silso/datafiles>, 2024.
- 667 Eliyahu Stoupel. Relationship between solar and seismic activity. *International Journal of*  
668 *Biometeorology*, 34(4):231–235, 1990. doi: 10.1007/BF01049646.
- 669 Mario Tavares and Angelo Azevedo. Influences of solar cycles on earthquakes. *Pattern*  
670 *Recognition in Physics*, 1(1):1–11, 2011.
- 671 James Theiler, Stephen Eubank, André Longtin, Bryan Galdrikian, and J. Doyne Farmer.  
672 Testing for nonlinearity in time series: the method of surrogate data. *Physica D:*  
673 *Nonlinear Phenomena*, 58(1–4):77–94, 1992. doi: 10.1016/0167-2789(92)90102-S.

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

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