Tree rings record the broad drought pattern and almost none of the local one

A preregistered space-time test on 664 western North American tree-ring sites and 399 drought cells, with planted fields and four kinds of control

Zain Dana Harper

Version 1.0, 9 November 2026. Finding F97 in the findings ledger. Strong signal (tier B*).

Abstract

Tree-ring records are often judged one site at a time, by comparing a site's yearly growth with the climate at that spot. We split that comparison by spatial scale as well as by time scale. The data are June to August drought from a gridded index, 399 one-degree cells in western North America from 1903 to 2025, and 664 tree-ring chronologies in the same box, complete from 1903 to 1980. In each of 30 cells of a grid of six spatial bands and five time bands we measured the share of the drought signal that the trees pass. The preregistered contrast between broad, slow scales and fine, fast scales was 0.64 (95% interval 0.62 to 0.68), against 95th percentiles of 0.22 for phase-randomised trees and 0.30 for trees moved to shuffled sites, with an adjusted p of 0.035. Below 55 km the trees pass at most 0.004 of the drought signal in every time band, and above 460 km they pass 0.27 to 0.91. Three tests on the drought field came back null. The drought field's covariance did not depart from a space part times a time part, no travel speed could be claimed, and broad and fine patterns lose their memory at the same rate, with a dispersion exponent of 0.02 (-0.12 to 0.17). The trees act as a spatial filter.

1 What this means for a reader

A tree-ring site that matches its local drought record only weakly is not a weak drought recorder in every sense. In the American West the trees carry most of the broad, regional pattern of summer drought and almost none of the local detail. When a proxy is scored site by site, ask how much of its lost skill is lost in space rather than in time.

2 Introduction

Tree rings are the main yearly record of drought before instruments. Cook, Meko, Stahle and Cleaveland [1] built a grid of summer drought reconstructions for the United States from a dense network of chronologies. Their method fits each grid point from the chronologies within a search radius, and they reported that a radius of 450 km worked best. The method is local by design. The reconstructed maps still kept the spatial features of drought, which the authors noted although the method was not built to be spatial.

The drought index in this study is the self-calibrating Palmer Drought Severity Index, the scPDSI. Wells, Goddard and Hayes [2] built it so that the index's constants are fitted to each location and values compare across places. We use the gridded scPDSI distributed with CRU TS 4.10 [3] [4].

Averaging trees over larger areas raises the shared climate signal. Duchesne and colleagues [6] showed with forest inventory cores in Quebec that averaging at the stand level removes tree-level noise, that covariance among tree-ring signals falls with distance, and that reliable climate signals can be drawn at scales from 37 to 583,000 square km.

These studies treat space and time one at a time. A field that varies in both can hold structure neither view shows: a front that moves, patterns that fade faster at fine scales than at broad ones, or a filter that passes

some combinations of scale and period and not others. This paper asks four questions at once. Does the drought field's covariance split into a space part times a time part? Do its anomalies travel? Do broad and fine patterns fade at different rates? And which combinations of spatial scale and period do the tree rings pass?

3 Preregistration and design

The design was written and hash-locked on 30 October 2026 at 04:20 UTC as a preregistration, before any outcome on the real data was computed. The real data had been built and only their counts printed: 399 cells by 123 years and 664 sites by 78 years. Development before the lock used synthetic fields only, and the preregistration lists every change made during that development.

The locked file names the data, the band edges, the four statistics, the synthetic fields, each control, the forecasts and the pass rules. No unit, constant or rule changed after the lock. Each test passes only when the real statistic beats the 95th percentile of every scoring control, its adjusted p value is 0.05 or less, and, where named, its bootstrap interval sits above 0. A test that misses any rule fails, which makes each rule a kill condition.

Forecasts, written before the run: the covariance would be non-separable in both fields; travel would most likely be null or slow, under 100 km per year; the dispersion exponent would lie between 0 and 2, closer to the diffusion reading than the wave reading; and the tree filter test would pass.

4 Data and provenance

Drought: the June to August mean of the scPDSI from CRU TS 4.10 [4], 0.5 degree, 1903 to 2025, released under the Open Database License 1.0. The 0.5 degree cells between 125 and 100 W and 30 and 50 N were averaged to one-degree cells, keeping only one-degree cells with all four parts present in every year. This gave 399 cells and 123 years.

Trees: standard ring-width chronologies from the International Tree-Ring Data Bank [5], held by NOAA's National Centers for Environmental Information as public data, with a request to cite the contributors. Every site in the same box with a gap-free index from 1903 to 1980 entered, which gave 664 sites and 78 years. Each site was paired with the June to August scPDSI of its nearest 0.5 degree cell with data in every year.

Every series had its linear trend removed and was scaled to unit variance. Positions were projected about 112.5 W and 40 N, in km. The input files and code were hashed into the run record before the run.

5 Methods

Bands in space and time

Let K_b be the matrix that smooths values on the sites with a normalised Gaussian of width σ_b, where σ_1 to σ_5 are 55, 95, 160, 270 and 460 km. Set K_0 to the identity and K_6 to zero. The space band b of a field x is

$$x_b = (K_b - K_{b+1})x \tag{1}$$

for b from 0 to 5, so band 0 holds what varies below about 55 km and band 5 what varies above about 460 km, and the six bands add back to x. Each band is then split into five time bands by masks on the Fourier transform in time with edges at 1/32, 1/16, 1/8 and 1/4 cycles per year: periods over 32 years, 16 to 32, 8 to 16, 4 to 8 and 2 to 4 years. Each masked series is kept as a complex analytic signal.

The tree filter

For space band b and time band w, let a_x and a_y be the analytic signals of drought at the tree sites and of the trees. Summing over sites and years, with S_{xy} the sum of $a_x^* a_y$ and S_{xx}, S_{yy} the sums of squared magnitudes, the

passed share is the squared coherence

$$\gamma^2_{bw} = \frac{|S_{xy}|^2}{S_{xx} S_{yy}} \tag{2}$$

The preregistered contrast C is the mean of γ^2 over the four broad, slow cells (space bands 4 and 5, periods over 16 years) minus its mean over the four fine, fast cells (space bands 0 and 1, periods under 8 years):

$$C = \frac{1}{4} \sum_{L} \gamma^2_{bw} - \frac{1}{4} \sum_{F} \gamma^2_{bw} \tag{3}$$

where L and F name those two sets of cells.

Separability

Let $\rho(h, u)$ be the correlation between sites at distance h and time lag u, pooled over pairs in 100 km bins to 800 km and lags of 1, 2 and 3 years. A separable covariance has $\rho(h, u) = g_u \rho(h, 0)$. With g_u fitted by least squares, the statistic is the sum of squared departures

$$D = \sum_{h} \sum_{u} (\rho(h, u) - g_u \rho(h, 0))^2 \tag{4}$$

Travel

For site pairs within 1,200 km, the lag-one cross-correlation s_{ij} between site i in one year and site j in the next is pooled in 40 km offset cells and fitted as

$$s_{ij} = c \rho_0(|d_{ij} - v|) \tag{5}$$

where d_{ij} is the offset between the sites, ρ_0 the lag-zero correlation profile by distance, c a scale and v a velocity in km per year.

Dispersion

The decay rate of a pattern with wavenumber k is modelled as $a k^{\beta}$, with power per log wavenumber proportional to k^{-q} and white site noise. Each interior space band's autocovariance at lags 0 to 8 is predicted from the measured transfer of plane waves through that band on these exact sites and from the bias that linear detrending adds, and the four parameters are fitted to bands 1 to 4. The dispersion exponent β is the result. Beside it, each band's e-folding time is reported.

p values, intervals and multiplicity

Each control was drawn 200 times, or 300 times for the separable surrogate. A test's p value is the larger of its scoring controls' p values from draws,

$$p = \frac{k + 1}{n + 1} \tag{6}$$

with k the draws at or above the real value and n the draws. The seven real-data tests were adjusted together with the Benjamini-Hochberg correction at q 0.05. Every 95% interval comes from a moving-block bootstrap in time with blocks of 10 years.

6 Controls and planted checks

A planted check ran in the same job before the real data were scored. Synthetic fields were built on the drought cell positions with 123 years, from 800 random plane waves plus white noise: a separable field, a wave field carried at 120 km per year toward 30 degrees, a diffusing field and pure noise. Five realisations of each were scored, and the instrument passed every rule. The separable field read as separable (median p 0.37 against

a rule of at least 0.05), the planted wave's speed came back within 2.2 percent and its direction within 0.66 degrees, and the diffusion exponent read 1.88 against a rule of 1.7 to 2.3. On the separable field the exponent read -0.01 and on the wave 0.81. Those are the readings a real exponent is compared against. The same checks on the tree positions with 78 years were reported and not scored; there the diffusion exponent read 1.64, below the scoring range.

The real statistics faced these controls, each drawn many times to give a null distribution. Separability: a fitted separable field with the same lag-zero spatial covariance and pooled time autocovariance, and shuffled site coordinates. Travel: the lag-one matrix with its symmetric part kept and its antisymmetric part permuted across pairs in 100 km distance bins, and shuffled coordinates. Dispersion: shuffled coordinates and phase randomisation of each site. The tree filter: phase-randomised trees, and tree series reassigned to shuffled sites. The last control keeps every tree's own spectrum and breaks only the link to its place.

7 Results

The tree filter

The preregistered contrast between broad, slow scales and fine, fast scales was 0.64 (95% interval 0.62 to 0.68), against 95th percentiles of 0.22 for phase-randomised trees and 0.30 for trees moved to shuffled sites, with an adjusted p of 0.035. Figure 1 sets the contrast against both controls.

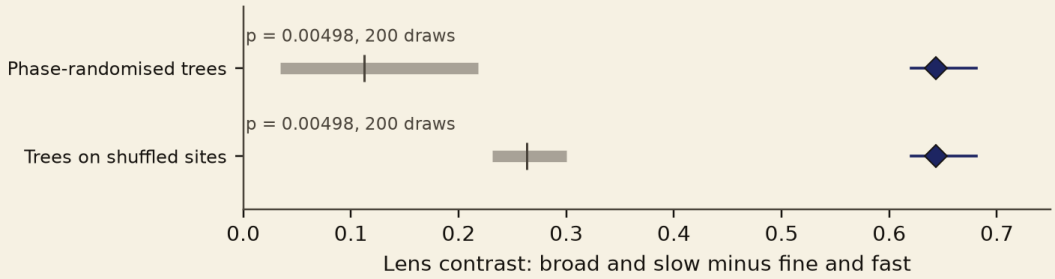


Figure 1. The lens contrast (diamond, with its bootstrap interval) against the 5th to 95th percentile range (bars) and mean (ticks) of each control, 200 draws each. Drawn from the run's result file.

Figure 2 gives the passed share in all 30 cells. The share rises with spatial scale in every time band. Below 55 km the trees pass at most 0.004 of the drought signal in every time band, and above 460 km they pass 0.27 to 0.91. Between 95 and 160 km they pass 0.07 to 0.11, and between 270 and 460 km 0.13 to 0.53. Time scale matters most at the broad end: above 460 km the share is 0.91 at periods over 32 years and 0.27 at 2 to 4 years. The phase-randomised control's mean share is at most 0.18, in that broadest and slowest cell.

The trees' own variance sits where their filter is closed. Figure 3 shows that 0.61 of the tree variance lies in the finest space band, against 0.12 for the drought index at the same sites.

The drought field

Table 1 gives all seven real-data tests. The drought field's covariance did not depart from a space part times a time part: its departure D was 0.0011 against a separable-surrogate 95th percentile of 0.0085 (p 0.72). The trees came back the same way (p 0.35). Both fields beat shuffled coordinates, which only shows that nearby sites are alike.

The drought fit for travel pointed 34.5 degrees north of east at 64.5 km per year and beat the antisymmetric shuffle (95th percentile 32 km per year). It did not beat shuffled coordinates, whose 95th percentile was 467 km per year, and no travel speed could be claimed: the bootstrap interval of the speed along the fitted direction



Figure 2. Passed share (squared coherence) between June to August drought at the tree sites and the 664 tree-ring chronologies, 1903 to 1980, per space band and time band. Small figures give the 95% bootstrap interval. Drawn from the run's result file.

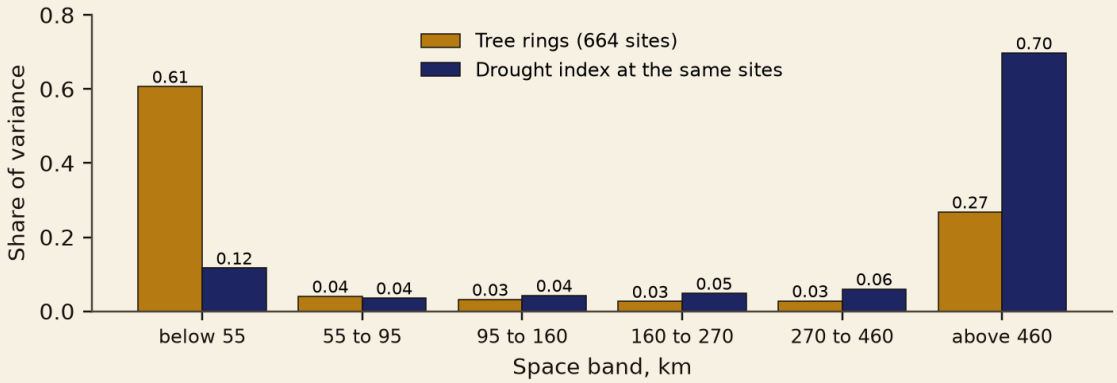


Figure 3. Share of variance in each space band, summed over time bands, for the tree-ring chronologies and for the drought index at the same sites, 1903 to 1980. Drawn from the published space-time panel file.

ran from -41 to 242 km per year and the adjusted p was 0.66.

Broad and fine patterns lose their memory at the same rate, with a dispersion exponent of 0.02 (-0.12 to 0.17). The e-folding times of the four interior bands were 0.76 to 0.83 years. Neither 1 nor 2 lies inside the interval, and the reading sits at the instrument's separable value, not at its wave or diffusion value. The trees' dispersion fit stopped at its lower bound of -1 and says nothing.
Table 2 sets the forecasts against the outcomes.

8 Discussion

The tree-ring network behaves as a spatial filter. It passes most of the drought signal at scales of several hundred km and above and almost none below 55 km. A site-by-site comparison of rings and local drought averages over that whole range and reports one number, so it reads as a weak recorder. In our findings ledger, a site-by-site comparison in time only, run after the preregistered test as exploration, gave pooled coherence of 0.24, 0.18, 0.19, 0.20 and 0.07 across the five time bands: weak and nearly flat in period. The space-time view places the weakness in space.

Table 1. The seven real-data tests. Controls are the 95th percentiles of the scoring controls. Adjusted p is after Benjamini-Hochberg over all seven.

Test	Estimate (95% interval)	Controls, 95th percentile	Adjusted p	Verdict
Separability, drought	D 0.0011 (0.0004 to 0.0125)	separable 0.0085; shuffled 0.00008	0.91	fail
Separability, trees	D 0.0017 (0.0012 to 0.0063)	separable 0.0042; shuffled 0.00013	0.80	fail
Travel, drought	64.5 km per year; along direction -41 to 242	antisymmetric 32; shuffled 467	0.66	fail
Travel, trees	8.3 km per year; along direction -185 to 401	antisymmetric 14.6; shuffled 462	0.91	fail
Dispersion, drought	beta 0.02 (-0.12 to 0.17)	shuffled 0.14; phase 0.14	0.80	fail
Dispersion, trees	beta -1, at its bound (-1 to 0.91)	shuffled 3.5; phase 3.0	1.0	fail
Tree filter contrast	0.64 (0.62 to 0.68)	phase 0.22; shuffled sites 0.30	0.035	pass

Table 2. Forecasts fixed in the preregistration against what the run gave.

Question	Forecast	Outcome
Separability	non-separable in both fields	separable in both; forecast failed
Travel	most likely null or slow, under 100 km per year	null; held
Dispersion	beta between 0 and 2, nearer diffusion	0.02, at the separable reading; partly held
Tree filter	pass	pass; held

The result agrees with two lines of earlier work. Cook and colleagues [1] found that chronologies within a radius of about 450 km predict a drought grid point best, which is the scale at which our filter opens. Duchesne and colleagues [6] found that averaging over stands removes tree-level noise and that the shared signal falls with distance. Our grid puts both on one map of scale and period, with a number and an interval in each cell.

The three nulls on the drought field matter in their own right. If the field had a travelling front or patterns that faded at scale-dependent rates, separate views in space and in time would miss them. Here they did not: at one-degree cells and yearly steps, the separate views lose nothing a coupled view could find. The coupling that matters lies in the tree filter. Each null is an honest null from a working instrument, since the planted wave and the planted diffusion were each recovered.

Part of the filter may belong to the reference grid. Harris and colleagues [3] report that CRU interpolates precipitation with a correlation decay distance of about 450 km, close to where the filter opens, so a smooth reference would itself pass little fine-scale signal to agree with. The next study needs station data or a denser grid as the reference.

9 Limitations

- The drought grid is smooth by construction, so part of the low fine-scale coherence may come from the reference and not from the trees.
- Tree sites cluster, so the finest band partly reads site noise and how dense the network is.
- The nulls on travel and dispersion hold for summer scPDSI at one degree, yearly steps, a box about 2,000 km across and 1903 to 2025. Monthly data, a larger domain or another index could show more.
- The scPDSI is a model built from temperature and precipitation, not a measurement of soil water.
- The trees' dispersion fit stopped at its bound and tells us nothing. On the tree positions the planted diffusion read 1.64, below the scoring range; that check was reported, not scored.

- The site-by-site comparison in the discussion was exploratory and not preregistered.
- Scoring was done by our own code and was not blind.

10 Verified

Verified: on 9 November 2026 the contrast and its interval, both control percentiles, every p value and adjusted p value, the passed share and interval in all 30 cells, the travel and dispersion estimates with their intervals, the e-folding times and the synthetic instrument readings were re-read from the run's result file. They match the findings ledger. The 0.61 and 0.12 variance shares were summed by the figure script from the published space-time panel file. The site-by-site coherence values in the discussion were read from the findings ledger, not from a result file. The SHA-256 of the preregistration and of the result file match the hashes in the run record. All three figures were drawn by a script that reads only the result file and the panel file. The study was not re-run for this paper. The literature was found by searches run on 9 November 2026 and read at the level of abstracts, publisher pages and search summaries; the full texts were not read. The 450 km decay distance for CRU precipitation is taken from our earlier record of Harris and colleagues and was not re-read on that date.

11 Corrections

- 9 November 2026: first published version. The draft summary called the fitted travel direction east-northeast; the result file gives 34.5 degrees north of east, which this paper uses.

12 Reproduction

- Data: download the CRU TS 4.10 scPDSI grid and the ITRDB standard chronologies for 125 to 100 W and 30 to 50 N. Average June to August drought into complete one-degree cells; keep chronologies gap-free from 1903 to 1980 and pair each with its nearest 0.5 degree cell.
- Remove each series' linear trend and scale it to unit variance.
- Build space bands with eq. 1 and time bands with the edges above. Compute eq. 2 in every cell and the contrast, eq. 3.
- Compute eqs. 4 and 5 and the dispersion fit; score each against its controls with eq. 6 and adjust the seven tests together.
- Preregistration SHA-256: 6d8b274a6784ca81aa2050dfe6fa7005402d05301405d59fb376e0b70a97d2eb. Plan SHA-256: 2641370a2fb740f53e0472b2a3039c4cb0b537533daa6c27edb0f39d49bae1ff. Run 20261030T042024Z, seed 2026102941. Result file SHA-256: 6d1917adb4c50b78125146ee5dcaa42c 6081a05edc4bcb1920e72fbc3b60cf0d.
- Cost: the analysis took about 27 minutes on one desktop CPU. The shared machine drew at most 63 Wh over the run.

13 Appendix A. The preregistered verdicts

Each verdict is a set of rules written before the run, and it passes only when every rule holds. Values are read from the run's result file.

14 Glossary

Each term is defined where it is first used in the text. The same definitions, with the field each comes from and what each excludes, are collected here.

- **self-calibrating Palmer Drought Severity Index (scPDSI)**. A monthly drought index computed from temper-

Table 3. Preregistered verdicts and the rules that decided them.

Verdict	Deciding rule	Value	Holds
Instrument, separable	median p of a planted separable field at least 0.05	0.37	yes
Instrument, wave	median speed error at most 10 percent and direction error at most 10 degrees	2.2 percent; 0.66 degrees	yes
Instrument, diffusion	median beta of a planted diffusing field from 1.7 to 2.3	1.88	yes
Separability, drought	D above both controls' 95th percentiles and adjusted p at most 0.05	0.0011 against 0.0085	no
Separability, trees	the same rules for the tree field	0.0017 against 0.0042	no
Travel, drought	speed above both controls, adjusted p at most 0.05, interval along the direction above 0	64.5 against 467; -41 to 242	no
Travel, trees	the same rules for the tree field	8.3 against 14.6	no
Dispersion, drought	beta above both controls, adjusted p at most 0.05, interval above 0	0.02 against 0.14	no
Dispersion, trees	the same rules for the tree field	-1 against 3.5	no
Tree filter	contrast above both controls, adjusted p at most 0.05, interval above 0	0.64 against 0.30; p 0.035; 0.62 to 0.68	yes

ature and precipitation through a simple soil water balance, with its constants calibrated to each location so values compare across places. Field: climatology and hydrology. Excludes a direct measurement of soil moisture or of streamflow.

- **preregistration**. A dated written plan for a study, fixed and hash-locked before any result is seen: the question, the data, the statistic, the controls, the forecasts and the kill conditions. Field: research method (metascience). Excludes a plan revised after looking at the outcome, and a protocol written only for the record after the run.
- **control**. A version of the data with the structure under test removed (for example shuffled or pooled), scored the same way as the real data, so the real score has something to beat. Field: statistics and experimental design. Excludes a comparison with a different dataset or with a published number.
- **kill condition**. A rule written in the preregistration that, when met, decides against the claim with no room for judgement after the run. Field: research method (our usage, after Popper's falsification). Excludes reasons found after the run for setting a result aside.
- **tree-ring chronology**. One yearly series for a site, built by dating many cores, removing each core's age and size trend and averaging the results into a dimensionless index. Field: dendrochronology. Excludes a single core's raw ring widths, and a climate reconstruction made from the series.
- **space band**. The part of a field that varies between two spatial scales, taken as the difference of two Gaussian smoothings of the site values with the two scales as widths; the six bands add back to the field. Field: signal processing on irregular sites, our usage. Excludes an exact wavenumber, because the bands overlap and the sites are irregular.
- **time band**. The part of a yearly series whose period lies between two limits, kept by masking its Fourier transform. Field: spectral analysis. Excludes a moving average or a fitted trend.
- **squared coherence**. The share of one series' variance in a frequency band that a linear filter of the other series can account for, between 0 and 1; here pooled over sites and years in each space and time band and called the passed share. Field: spectral analysis. Excludes a nonlinear dependence, and the size of the response, which the gain measures.
- **separable covariance**. A covariance that factors into a function of distance times a function of time lag, so the spatial pattern of correlation has the same shape at every lag. Field: space-time statistics. Excludes fields whose anomalies travel or whose broad and fine patterns fade at different rates.

- **dispersion exponent**. The power beta in a decay rate that grows as the wavenumber to the power beta; diffusion gives about 2, and fine and broad patterns that fade at the same rate give 0. Field: physics of waves and diffusion. Excludes a speed of travel, which the propagation test measures.
- **p value (from draws)**. The share of control draws scoring at or above the real score, with one added to the count and to the number of draws, so it is never zero. Field: statistics. Excludes a p value from a textbook formula that assumes a distribution.
- **Benjamini-Hochberg correction**. A rule that adjusts a set of p values so that, among the tests called significant, the expected share of false calls stays at or below a chosen rate q, here 0.05. Field: statistics, multiple testing. Excludes the stricter Bonferroni rule, which bounds the chance of any false call.
- **95% interval**. The range that holds the middle 95 percent of the bootstrap or control draws for a number; we report it beside the number. Field: statistics. Excludes a statement that the true value lies inside with 95 percent probability.
- **moving-block bootstrap**. An interval made by recomputing a statistic many times on years resampled in runs of consecutive years, here 10, so that year-to-year correlation is kept inside each run. Field: statistics. Excludes resampling single years as if they were independent.
- **planted check**. Data built so that the right answer is known, run through the same analysis before the real data, to show that the analysis can see the effect when it is there. Field: research method; called a positive control in experimental science. Excludes any check made on the real data, which cannot show what the analysis misses.
- **null distribution**. The spread of scores a control gives over many random draws; its 95th percentile is the bar a real score must clear. Field: statistics. Excludes a theoretical distribution assumed rather than drawn from the data.
- **phase randomisation**. A copy of a series with the same power spectrum and random Fourier phases, so it keeps the series' autocorrelation and loses its timing against any other series. Field: time-series statistics, surrogate data. Excludes shuffling years, which also destroys the autocorrelation.
- **findings ledger**. The append-only, dated list of every result we have recorded, numbered F1, F2 and on; corrections are added as new dated lines and old entries are never rewritten. Field: our records. Excludes a curated list of successes.
- **honest null**. A working test that found no effect beyond a stated bound, a forecast that failed, or an error of our own found and fixed, reported with the same weight as a positive result. Field: research method (our usage). Excludes a void run, and a result left unpublished because it was negative.
- **verdict**. The pass or fail of one preregistered rule set against the run's numbers, read off by code. Field: research method (our usage). Excludes an interpretation written after reading the numbers.

15 References

[1] Cook, E. R., Meko, D. M., Stahle, D. W. and Cleaveland, M. K. (1999). Drought reconstructions for the continental United States. Journal of Climate 12, 1145-1162. doi:10.1175/1520-0442(1999)012<1145:DRFTCU>2.0.CO;2

[2] Wells, N., Goddard, S. and Hayes, M. J. (2004). A self-calibrating Palmer Drought Severity Index. Journal of Climate 17, 2335-2351. doi:10.1175/1520-0442(2004)017<2335:ASPDSI>2.0.CO;2

[3] Harris, I., Osborn, T. J., Jones, P. and Lister, D. (2020). Version 4 of the CRU TS monthly high-resolution gridded multivariate climate dataset. Scientific Data 7, 109. doi:10.1038/s41597-020-0453-3

[4] Climatic Research Unit, University of East Anglia. CRU TS 4.10 gridded climate data, including the self-calibrating PDSI. Open Database License 1.0. https://crudata.uea.ac.uk/cru/data/hrg/

[5] International Tree-Ring Data Bank. NOAA National Centers for Environmental Information, World Data Service for Paleoclimatology. Standard chronologies, North America. https://www.ncei.noaa.gov/pub/data/paleo/treering/chronologies/

[6] Duchesne, L., D'Orangeville, L., Ouimet, R., Houle, D. and Kneeshaw, D. (2017). Extracting coherent tree-ring climatic signals across spatial scales from extensive forest inventory data. PLoS ONE 12, e0189444. doi:10.1371/journal.pone.0189444