

Uneven evidence for a shared planet, sampling and publication are coupled worldwide

Lingli Fan ^a, Yiming Lei ^a, Yufei Chen ^{a,b}, Qian Huang ^{d*}, Hui Jia ^e, Shan Jiang ^f, Jingchun Liu ^a,
Haoliang Lu ^a, Chongling Yan ^a, Hualong Hong ^{a,b,c*}

Affiliations

a. Key Laboratory of Ministry of Education for Coastal and Wetland Ecosystems, Xiamen University, Xiamen 361102, China

b. Coastal and Ocean Management Institute (COMI), Xiamen University, Xiamen 361102, China

c. National Observation and Research Station for the Taiwan Strait Marine Ecosystem (Xiamen University), Zhangzhou 363305, China

d. Institute of Geosciences, University of Mainz, Johann-Joachim-Becher-Weg 21, Mainz 55128, Germany

e. School of the Environment and Safety Engineering, Jiangsu University, Zhenjiang 212013, China

f. State Key Laboratory of Estuarine and Coastal Research, East China Normal University, Shanghai 200241, China

Author email addresses and ORCID

Lingli Fan (a) — 33120251152187@stu.xmu.edu.cn

Yiming Lei (a) — 33120241152277@stu.xmu.edu.cn

Yufei Chen (a,b) — 33120241152351@stu.xmu.edu.cn

Qian Huang (d) * — qhuang01@uni-mainz.de

Hui Jia (e) — 1000005403@ujs.edu.cn

Shan Jiang (f) — sjjiang@sklec.ecnu.edu.cn

Jingchun Liu (a) — liujingchun@xmu.edu.cn

Haoliang Lu (a) — luhl@xmu.edu.cn, 0000-0003-0720-3137

Chongling Yan (a) — ycl@xmu.edu.cn

Hualong Hong (a,b,c) * — honghl@xmu.edu.cn, 0000-0003-4284-8096

29 **Abstract**

30 Natural archives record environmental change over timescales that instrumental series cannot reach,
31 and global environmental assessments rest on the evidence they supply. To our knowledge, however,
32 where that evidence is produced, and whether it reflects the processes it describes, has not been
33 established. Here we reconstruct the sampling and publication geography of 33,145 studies of past
34 environments, tracing each to the Sustainable Development Goal (SDG) it addresses. Sampling and
35 publication are not decoupled but nearly isomorphic. Twenty-six countries account for 89% of
36 publication credit and 81% of sampling, and all eleven goals with evidence flows draw 82–90% of
37 their evidence from them. Those countries hold 57.5% of world population and produce that
38 evidence roughly 28 times more intensively per head. National income acts on research output as a
39 threshold rather than a gradient. Below about 41,000 PPP dollars per head additional income buys no
40 additional output, while above it output rises steeply with income. Because credit and sampling are
41 read from affiliations and abstracts, the analysis captures where research is led rather than the
42 provenance of every sample. The evidence base for global goals therefore reflects the research
43 geography of a minority rather than the distribution of the processes being measured. Our findings
44 provide the first baseline against which commitments to widen participation can be measured, and
45 identify analytical capacity, rather than access to samples, as the binding constraint on whose
46 evidence the planet is studied.

47 **Keywords:** Bibliometrics; Palaeoclimatology; Research inequality; Geographic bias; Open science

48 **1 Introduction**

49 Natural archives, lacustrine and marine sediments, ice cores, tree rings, speleothems, corals, peat and
50 loess, record the environmental evolution of the Earth system over timescales that instrumental series
51 cannot reach (Bradley, 2015). They supply the baseline against which contemporary change is

52 judged (Kaufman et al., 2020) and are the empirical foundation for testing and improving climate
53 models (Masson-Delmotte et al., 2013; Braconnot et al., 2012), for evaluating the anomalous
54 magnitude of industrial-era warming (Abram et al., 2016; Tierney et al., 2020), and for assessing the
55 pace at which the cryosphere responds to forcing (Zemp et al., 2019). Over the past half century the
56 field has moved from single-site description to quantitative reconstruction and from regional
57 compilations to global syntheses.

58 That transition was made possible by methodological advances rather than by a change in the
59 archives themselves. Improvements in radiocarbon and uranium-series dating made absolute
60 chronologies routine (Bronk Ramsey, 2009; Edwards et al., 1987), age-depth modelling became
61 statistically explicit (Blaauw and Christen, 2011), and proxy calibration systems were standardised.
62 In parallel, community data infrastructure lowered the cost of reuse, the International Tree-Ring Data
63 Bank turned scattered ring-width measurements into a globally distributed network (Grissino-Mayer
64 and Fritts, 1997; Esper et al., 2002), and comparable global compilations followed for other carriers
65 (Kaufman et al., 2020). The same period saw the analytical repertoire of speleothem and loess
66 research broaden (Fairchild et al., 2006; Liu and Ding, 1998). Bibliometric analysis has become the
67 standard instrument for describing how such a field is organised, because it can characterise
68 publication dynamics, geographic concentration and thematic structure from a complete corpus
69 rather than from a sample of expert opinion (Aria and Cuccurullo, 2017; Donthu et al., 2021; Zupic
70 and Čater, 2015). Applied to the geosciences, however, it has largely described growth curves,
71 national rankings and keyword co-occurrence, typically through co-citation and keyword mapping
72 (Chen, 2006); it rarely asks who physically acquires the material, and almost never connects archive
73 type to the type of question asked.

74 A separate body of work has shown that this distinction matters. Studies of sampling effort and of
75 research leadership in biodiversity and Earth sciences find that the countries where material is

76 collected are frequently not the countries where it is analysed, a pattern described as parachute
77 science and quantified for fossil localities, marine field programmes and tropical ecology (Raja et al.,
78 2022; Stefanoudis et al., 2021; Asase et al., 2021). In these analyses, publication counts alone
79 systematically overstate the scientific capacity of sampling-rich countries and understate the
80 concentration of interpretive capacity (Raja et al., 2022). The economic determinants of research
81 output have a longer pedigree, national scientific production scales with economic size, and the
82 scaling exponent carries information about whether returns to income are increasing or diminishing
83 (May, 1997; King, 2004), while whether that relation is uniform across the income range or changes
84 slope at a threshold is an empirical question with direct implications for capacity-building policy.
85 Neither question has been tested across all natural-archive classes.

86 Here we address those questions with the union of two bibliographic sources covering the field, the
87 Web of Science Core Collection and OpenAlex. We ask three questions. The first concerns whether
88 there is a spatial asymmetry between where samples are taken and where knowledge is produced,
89 and how it relates to national income. The second concerns the hierarchical structure that links
90 archive carriers, proxy indicators, scientific questions and societal goals. The third, and most
91 consequential, concerns whether the concentration of capability distorts the evidence base that
92 decisions about global environmental goals actually rest on. Answering them requires a measure of
93 sampling that is independent of publication counts, a country classification built on both dimensions,
94 and a record-level network that preserves the multi-label character of the underlying data.

95 **2 Results**

96 **2.1 Production and sampling are isomorphic, not decoupled**

97 Production and sampling are tightly coupled across countries and both are extremely concentrated
 98 (Figure 1a-c). Cross-classifying 195 countries gives 26 in the high-high class and 39 double-zero
 99 countries, and two of the nine cross-classes are empty, so no country publishes heavily without
 100 sampling or samples heavily without publishing. The two dimensions correlate at Spearman $\rho =$
 101 0.884 and Pearson $r = 0.938$, and the high-high class, 13.33% of countries, accounts for 89.11% of
 102 fractional publications and 81.29% of sampling-location articles. The Gini coefficient is 0.910 for
 103 publications and 0.847 for sampling. The concentration has a continental expression, with Europe
 104 leading on publications (45.1% of fractional credit) and Asia on sampling (34.3%), and the ten
 105 largest producers run from the United States to the Russian Federation.

106 The signed wedge between publication and sampling shares is small everywhere and robust to how
 107 multi-label records are treated (Figure 1d, e; Table S3). It is indistinguishable from zero for all nine
 108 archive sub-classes tested, and assigning every record to its first sub-class instead leaves the two
 109 sediment Ginis unchanged at 0.8748 and 0.8559. Only polar ice cores reverse the sign, with a median
 110 difference of -0.0003. Across the eleven classes the size of the wedge tracks how widely an archive
 111 is sampled rather than how many records it carries or how expensive the material is to obtain.

112 The asymmetry is not uniform across archive types (Figure 1d, e; Tables S3, S4). Marine and lake
 113 sediments are the closest available natural experiment, being alike in material and indicator
 114 repertoire and differing mainly in whether sampling requires a research vessel or a drilling
 115 programme. Their sampling concentration differs by +0.0189 Gini points, marine 0.8748 against lake
 116 0.8559 with a 95% bootstrap interval of 0.0133 to 0.0237, so facility dependence does contribute to
 117 inequality. Polar ice cores are more concentrated still, 0.9742 against 0.9052 for the distributed tree-
 118 ring network. Their binding constraint is not cost but sovereignty, of the 2,824 polar ice-core records,
 119 1,388 were sampled in Antarctica and 642 in Greenland, and 2,265 carry at least one location outside
 120 the 195-country roster, because Antarctica has no sovereign state and Greenland is not one. The

121 archive that is hardest to reach is also the one whose sampling lies furthest outside the system of
122 sovereign states.

123 The layers differ sharply in how narrowly they are held (Figures S2, S3). Archive classes are
124 strongly unequal in size, with sediment carrying 70.85% of all labelled records against 0.17% for
125 other natural carriers, and no record matches two or more top-level classes. Indicators concentrate
126 more narrowly than questions, biological material leading at 39.72% against 33.86% for chemical
127 composition, the two classes that together account for most of the field's proxy repertoire (Croudace
128 and Rothwell, 2015; Castañeda and Schouten, 2011), whereas cognitive goals are stable over time
129 and led by reconstruction and inversion at 33.49%.

130 **2.2 Income, scaling and the elasticity step**

131 Income separates the corpus more sharply than geography (Figure 2a). High-income countries, 64 of
132 195, produce 70.52% of fractional publications against 0.09% for the 25 low-income countries, and
133 the gradient is monotonic across the four World Bank groups. Income must be measured in constant
134 purchasing power, because the same regression returns an exponent of 1.009 on current US dollars
135 and 1.351 on constant 2021 PPP international dollars, a difference that changes the shape of the
136 relationship rather than its scale.

137 The residuals of the country-level fit are asymmetric, and the asymmetry sits at the low end (Figure
138 2b-d). Grouping by fractional output, the median residual is indistinguishable from zero in every
139 group above five publications, and the only group whose bootstrap interval excludes zero is the one
140 just entering the literature, whose median residual is -0.336 log₁₀ units. The relation is also not a
141 single power law, sampling is a zero-inflated response, with 85 of 195 countries recording none and
142 42 publishing nothing, and a Tweedie generalised additive model needs 3.52 effective degrees of
143 freedom where a straight line would use one.

144 A single exponent is nevertheless a poor description, because elasticity is a property of countries
145 rather than of the field (Tables S2, S5). Country-specific elasticities span more than an order of
146 magnitude and 42 of the 56 countries with enough observations lie above the pooled value. They are
147 not randomly distributed, since a model in which the elasticity is flat below an income threshold and
148 steps upward above it fits far better than a single trend (AIC 289.6 against 307.7). The step, about
149 3.96 elasticities at 41,458 PPP international dollars per head, is the whole of the signal (Figure 2e),
150 because once it is in the model the residual slope is indistinguishable from zero while the step
151 interval excludes it.

152 **2.3 How country groupings connect to archives, indicators, questions and goals**

153 The network resolves the whole corpus and is overwhelmingly concentrated (Figure 3). Because
154 each record contributes a total weight of one, the flow axis reads directly as records, and high-high
155 countries carry 89.18% of all flow with all twenty largest pathways originating from them. Four
156 goals absorb 95.9% of that flow, SDG 13 (Climate Action) taking 60.34%, SDG 15 (Life on Land)
157 13.14%, SDG 14 (Life below Water) 12.17% and SDG 6 (Clean Water and Sanitation) 10.25%.

158 The layers are unequally concentrated. Indicators are led by biological material at 39.72% and
159 concentrate more narrowly than questions, whose leading category carries 33.49%, and the same
160 ordering holds by the Herfindahl index and by normalised entropy. The community converges more
161 on what to measure than on what to ask.

162 The concentration propagates to the goals themselves (Figure 3b). All eleven goals with non-zero
163 flow draw between 81.9% and 90.1% of their evidence from the twenty-six high-high countries,
164 while lower-middle-income countries contribute 0.00% to every goal and low-income countries
165 between 0.00% and 0.12%. Those twenty-six countries hold 57.5% of world population, and per

166 head of population the evidence reaching these goals is produced roughly twenty-eight times more
167 intensively than in the lower-middle, low-income and low-low classes combined.

168 **3 Discussion**

169 Analytical capacity, not access to material, organises this field. Twenty-six countries supply 89.11%
170 of its publications and 81.29% of its sampling, and the countries outside that group are outside both,
171 so what separates the field is not who collects samples but who can analyse them. That distinction
172 decides where reform should aim, and it frames the three questions this discussion takes up, namely
173 what produces the inequality, what it costs the evidence base on which global policy depends, and
174 what would close it. Income separates the corpus more sharply than geography, and the elasticity of
175 output to income is a step rather than a gradient, indistinguishable from zero below about 41,458 PPP
176 international dollars per head and jumping roughly fourfold above it. Because archive classes are
177 separate communities with separate geographies, and because the knowledge network funnels 95.9%
178 of its flow into four Sustainable Development Goals, concentrating capability narrows the questions
179 the field is able to answer.

180 **3.1 Isomorphism and the capacity gap**

181 This finding qualifies rather than contradicts the parachute-science literature. That literature
182 documents a mismatch between where material is collected and where it is analysed, comparing field
183 sites with author affiliations (Raja et al., 2022; Stefanoudis et al., 2021; Asase et al., 2021). Our
184 sampling measure is country-level rather than site-level, and at that resolution the mismatch
185 disappears, because the countries that publish are the countries whose territory is sampled. The
186 asymmetry lies instead between the small group that does both and the large group that does neither,
187 so framing the problem as decoupling would point policy towards sample-sharing agreements

188 whereas framing it as isomorphism points towards building analytical capacity where sampling
189 already happens. The stratified analysis sharpens this and identifies where the coupling breaks down.

190 Income separates the field more sharply than geography, and only constant purchasing power reveals
191 the shape of the relation rather than its scale. Elasticity is likewise a property of countries rather than
192 of the field, and its distribution is diagnostic, in that 42 of 56 well-observed countries lie above the
193 pooled value, and the step at 41,458 PPP international dollars per head is the whole of the signal. The
194 step survives normalising output by population, which removes a component moving against income
195 and accounts for about 8% of the raw-total step (Tables S2, S5). What matters for the argument is not
196 the magnitude of the exponent but that countries below the threshold convert resources into
197 publications at the same, near-zero rate, so no amount of income below it predicts who publishes.

198 The shape of the inequality identifies the mechanism more precisely than the correlation does.
199 Comparing each class's share of publications with its share of sampling gives a signed wedge, with
200 the high-high countries holding 89.11% of publications but 81.29% of sampling, a surplus of 7.8
201 points, while the 39 medium-publication, medium-sampling countries run a deficit of 7.2 points. The
202 top therefore publishes more than it samples and the bottom samples far more than it publishes, and
203 in a field where every publication rests on a physical archive the difference is made up by analysing
204 archives that are not on the author's own territory. Four mechanisms can produce this pattern, and the
205 frame of the question is reproduced by training, so archives in capacity-poor countries are used
206 disproportionately to answer questions set elsewhere.

207 One alternative explanation has to be ruled out, namely that the coupling is a bookkeeping identity, if
208 every paper were written by an author from the country where its sample was taken, the two counts
209 would move together by construction. Three observations argue against this. A purely mechanical
210 relation would leave every residual at zero, whereas the one-to-five publication group sits below the
211 fitted line with an interval that excludes zero; it would give a correlation of one, whereas the

observed Pearson r of 0.938 leaves a measurable wedge; and it would not make the sign of the wedge vary systematically with class, yet the wedge is positive at the top and negative in the middle and lower tiers. What the design cannot exclude is that a single common cause, the size of a national research system, drives both dimensions, and income is only a partial proxy for it.

3.2 The evidence base for global goals is not globally produced

Concentration would be a matter of scholarly geography were it not for what the network delivers. Because every record resolves to a Sustainable Development Goal, the network measures not only who publishes but who supplies the evidence attached to each goal, and that distribution is narrower than the publication distribution it derives from. All eleven goals with non-zero flow draw between 81.9% and 90.1% of their evidence from the twenty-six high-high countries, while lower-middle-income countries contribute 0.00% to every goal and low-income countries between 0.00% and 0.12%. The gap is one of intensity as well as totals, since those countries hold 57.5% of world population and per head produce the evidence reaching these goals roughly twenty-eight times more intensively than the poorest classes combined.

This matters because the assessments that inform the Sustainable Development Goals are built from the same literature this network describes. Where that literature comes from a narrow set of countries, its spatial and ecological coverage inherits the research geography of those countries rather than following the distribution of the processes being measured, a limitation the Intergovernmental Panel on Climate Change (IPCC) has itself recognised in its calls for geographic balance (Masson-Delmotte et al., 2013). The point concerns the evidence base and not the quality of individual studies. It also decides the remedy, because if the deficit were one of sampling access the answer would be agreements that share material, whereas a deficit of analytical capacity has to be met by building that capacity where the archives and the affected populations are.

235 That diagnosis locates the gap in analytical capacity, and the international apparatus is not aimed at it.
236 The governance gap over this object is structural, because instruments written for genetic resources
237 assume that value and benefit flow from access to a physical sample, whereas natural archives yield
238 their value chiefly through analysis. The Convention on Biological Diversity (CBD) and its Nagoya
239 Protocol (Secretariat of the Convention on Biological Diversity, 1992; Secretariat of the Convention
240 on Biological Diversity, 2011), the Kunming-Montreal Global Biodiversity Framework (Secretariat
241 of the Convention on Biological Diversity, 2022) and the recent decisions on digital sequence
242 information (Secretariat of the Convention on Biological Diversity, 2024; Nature, 2024; Sett et al.,
243 2024) all regulate the movement of material or of its digital representation, and the new agreement
244 on marine biodiversity beyond national jurisdiction (United Nations, 2023) follows the same logic.
245 That is the opposite of the CBD model, and it is why extending an access-and-benefit-sharing
246 architecture to natural archives cannot by itself close the gap.

247 The generic instruments are also asymmetric between reading and writing. The UNESCO
248 Recommendation on Open Science names equity and capacity building among its commitments
249 (UNESCO, 2021a), against a baseline in which four out of five countries invest less than one per
250 cent of GDP in research (UNESCO, 2021b), while Research4Life and the diamond open-access
251 sector address access to the literature (Springer Nature, 2026; Yang and Huang, 2025), the
252 Systematic Observations Financing Facility (SOFF) pays for observing capacity in the least
253 developed countries (World Meteorological Organization, 2022) and the UN Decade of Ocean
254 Science runs a dedicated capacity workstream (UNESCO-IOC, 2021). Editorial calls for geoscience
255 to globalise, and evidence that authors from low- and lower-middle-income countries remain under-
256 represented even in the journals most committed to equitable development, describe the same
257 asymmetry (Dangles et al., 2022; Nature Geoscience, 2015). Nothing among them targets the

conversion of sampling capacity into publication output, and the countries with one to five publications sit exactly on that boundary.

3.3 What would close the gap

The field nevertheless holds levers that are specific to it, and they are more concrete than the generic instruments. Training is the first, because the elasticity step is a capacity threshold rather than a funding one, and the countries just below it already have the archives and the field programmes. Its concrete form is the twinned doctoral and postdoctoral programme, in which a student is registered in a laboratory in a high-high country and a laboratory in a low-publication country and divides their time between them, and the constraint is not the willingness of either partner but the absence of a funding line that treats analysis, rather than sampling, as the object of capacity building. Regional synthesis networks are the second, since they let countries with archives but thin analytical capacity contribute to questions that would otherwise be answered elsewhere; the community databases already built for this purpose, Neotoma and PANGAEA among them, show that the infrastructure exists (Damerow, 2025; Williams et al., 2018; Diepenbroek et al., 2002). Governed access is the third, and it exists but is not equitable, as ocean drilling protects new cores with a one-year post-cruise moratorium while land-based archives carry no equivalent (International Ocean Discovery Program, 2024). The deep-time initiatives that have tried to build the missing capacity at scale remain the exception (Stephenson et al., 2020; Deep-time Digital Earth, 2022), and the most recent global assessment reaches the same conclusion (UNESCO, 2026).

Shared infrastructure is the fourth lever, and the largest attempt to build it is Chinese-led. Across this corpus the share of publications produced by the twenty-six high-high countries falls only from 93.2% in the first period to 88.9% in the last, so the question is where the chain is weakest rather than which link to abandon. Sample access, repository and curation systems, and analytical

laboratories are complementary rather than substitutable, and strengthening one without the others leaves the chain no stronger; what the evidence here identifies is that the analytical link is the binding one, not that the others do not matter.

The community's own governance is where an equity norm would have to live, and it is half built. Past Global Changes requires that the data behind any publication it acknowledges be publicly available and encourages inclusive authorship of large syntheses (Past Global Changes, 2018), and its code of conduct carries a sanctions route, but that code governs meetings rather than fieldwork or sampling, and the one concrete co-authorship norm in the proxy community comes from a database rather than a programme. The most quotable development in this layer is recent and negative, in June 2025 Past Global Changes discontinued the mobility fellowships aimed specifically at early-career scientists in low- and middle-income countries, to concentrate resources on working groups. The most transferable instrument, however, is a journal policy. Climate of the Past requires authors to address power inequalities at the proposal, data-access, fieldwork, publication and legacy stages, including in-country leadership of publications, and grounds the requirement in the geoscience union's 2023 statement recognising helicopter research as scientific neocolonialism (European Geosciences Union, 2023). Elsevier's research data guidelines already let a journal choose between encouraging, requiring and enforcing data deposit, so Global and Planetary Change is in a position to adopt the same architecture, a sample-and-data return requirement, a declaration of in-country collaboration and a published authorship norm would cost the journal nothing and would place the step this paper isolates inside a reviewer-facing obligation rather than a recommendation. The last lever belongs to the sampling country and is rarely used, access to an archive on national territory is granted by permit, and a permit condition can require a named local collaborator, deposit of a reference collection in a national repository, return of derived data or registration of the samples, as biodiversity research has begun to do (Haelewaters et al., 2021; Trisos et al., 2021; Ramirez-

305 Castaneda et al., 2022), and where sampling overlaps Indigenous territories the CARE principles set
306 the standard for the data side (Carroll et al., 2020). Because sampling in this field is geographically
307 spread while publication is not, a permit condition is the one instrument that acts precisely where the
308 asymmetry lies.

309 If this interpretation is correct, the gap should close in specific and measurable ways. The elasticity
310 step should flatten as analytical capacity spreads; the residual of the one-to-five publication group
311 should move towards zero; the number of countries with at least one fractional publication, 153 of
312 195 here, should rise; the share of flow reaching prediction and simulation, at present 0.33%, should
313 increase; and the wedge between publication and sampling shares at the top of the distribution should
314 narrow. Each is computable from the next release of the corpus, and we report them here as a
315 baseline rather than a forecast.

316 Three options follow for the intergovernmental bodies already active in this space, and all three work
317 within mandates they already hold. The first is to make analytical capacity an explicit criterion in the
318 capacity-building workstreams they fund, so that a proposal which builds a laboratory or a twinned
319 programme competes on the same footing as one that buys a ship or a drill. The second is to attach
320 an authorship condition to sample access, requiring that work derived from material collected under
321 a national permit includes authors from the country of origin, which converts an access regime into a
322 capacity instrument without changing its legal architecture. The third is to extend the curation duties
323 that data repositories already carry to the laboratory partnerships that generate measurements, so that
324 the instruments and the training that produced a dataset are recorded alongside it and can be credited
325 and repeated.

326 **3.4 Limitations and outlook**

327 Four kinds of limitation bound these results, and they differ in how far they reach. The corpus is a
328 union of two bibliographic sources, the Web of Science Core Collection and OpenAlex, and it
329 remains biased towards English-language and indexed literature; a further 16,044 records retrieved
330 by the same queries were excluded as out of scope by a model judgement, so the analysed corpus is a
331 screened rather than an exhaustive set (Table S7). Field labels are model-derived, and 2.54% of
332 records could not be assigned to an archive class while 19.23% of sampling-location fragments were
333 not resolvable to a country, so archive shares describe the classified subset. Country attribution rests
334 on authorship locations under fractional counting and on stated sampling locations, and neither
335 captures who owns a sample or its data.

336 The definition of sampling is itself a choice with consequences. It is taken throughout as the
337 geography of the archives a paper studies rather than the fieldwork its authors performed, which is
338 why 1,293 reviews and syntheses enter the corpus. Excluding them moves the sediment and ice-core
339 contrasts by 0.0001 and 0.0005 Gini and sampling concentration by 0.0006, so the choice does not
340 drive the conclusions (Table S8); the stratified analysis, which requires an archive sub-class and a
341 resolvable sampling country, retains 32,302 of the 33,145 records, and the unclassified remainder is
342 not randomly distributed over time.

343 The interpretation the paper rests on is an inference rather than a measurement. That the pattern is a
344 capacity gap rather than a sampling gap follows from two measured quantities, publication and
345 sampling, and the network, funding and training variables that would measure analytical capacity
346 directly are not observed here. The Sustainable Development Goal layer encodes goals named in
347 abstracts rather than verified contributions, and the sampling year is the publication year, a proxy for
348 sampling activity rather than a record of it.

349 Three lines of work follow. First, the isomorphism result should be tested at site level, where the
350 parachute-science literature operates, to establish whether the country-level coupling survives finer

resolution. Second, the relation between income and output should be examined with instrumentation and funding variables rather than income alone, since the contrast between marine and lake sediments shows that the threshold is not simply the cost of field equipment, and those variables need to separate the capacity to acquire material from the capacity to analyse it. Third, the network should be re-estimated as the corpus grows, so that the topology reported here can serve as a baseline against which capacity-building interventions are measured.

3.5 Conclusions

From 33,145 records assembled from the Web of Science Core Collection and OpenAlex, 31,964 of them published between 1980 and 2025, this study reconstructs the global structure of paleoenvironmental natural-archive research. Output grew from 26 articles in 1980 to a plateau averaging about 1857.8 per year, following the saturating form long expected of quantitative studies of science (de Solla Price, 1963), passing through four periods of 2,489, 5,685, 8,471 and 15,319 articles respectively; a further 1,181 records carry publication years of 2026 or 2027 and fall outside the analysis window, which is why the four periods sum to 31,964 rather than to the full corpus. Knowledge production and sample acquisition are isomorphic, since 26 high-high countries produce 89.11% of publications and 81.29% of sampling-location articles, the two extreme mixed classes are empty, and income group separates the corpus more sharply than geography alone. Archive physics constrains what can be measured but not what is asked, and societal goals are reached almost entirely through reconstruction and inversion.

Income-output elasticities are not a property of the field, since they differ by more than an order of magnitude between countries and follow a step rather than a gradient, below an income threshold of 41,458 PPP international dollars per head they are indistinguishable from zero, and above it they jump by roughly four. The consequence is not merely uneven credit but an uneven evidence base.

Every Sustainable Development Goal draws between 81.9% and 90.1% of its evidence from those same 26 countries, and per head of population the evidence reaching them is produced roughly twenty-eight times more intensively in the high-high class than in the lower-middle, low-income and low-low classes combined. The record that global environmental assessments rest on is therefore shaped by the research geography of a minority of the world's population rather than by the distribution of the processes being measured. The instruments that exist for closing the gap, open-science policy, access and benefit-sharing, sample registration, community repositories, drilling moratoria and regional synthesis networks, mostly regulate access to samples and data, whereas the binding constraint identified here is analytical capacity in the countries that host the archives. Building that capacity is what would let the evidence base for a shared global future represent the planet rather than the part of it equipped to measure it; future work should test whether such investment changes the topology of this network.

4 Methods

4.1 Data and search

Two searches of the Web of Science Core Collection (Clarivate) were run, the first on 31 December 2025 and the second on 6 October 2026, and the second was merged with OpenAlex. The first search returned 11,174 records and served to establish the profile structure of the field, that is, the recurring types of study this literature contains, defined by sampling region, sample type, indicator type, scientific question and the Sustainable Development Goals addressed. These profiles were derived from the titles and abstracts of the first-search corpus by clustering, and DeepSeek Pro was then used to extract the key information they carry and to consolidate the category options that this paper uses throughout. The second search returned a further 42,063 records, and the same concept blocks run against OpenAlex returned 27,014. After de-duplication on DOI and normalised title 49,189

remained, of which 1,181 fell outside 1980-2025. A decision model then judged each record's relevance to palaeoenvironmental archive research, excluding 16,044 that lay outside the field and leaving the analysed corpus of 33,145. Retrieval and screening follow a PRISMA-style flow (Table S7; Figure S6).

4.2 Field extraction

Every record of the merged corpus was then labelled with Jev (TypeSafe System One), a decision model that returns a calibrated distribution over a fixed option set. The option set is the one consolidated from the first-search profile clustering, and Jev assigned each record to those fixed options across ten fields, covering article type, archive carrier at major and sub-class level, proxy indicator, scientific question, Sustainable Development Goal, sampling location and an in-scope judgement. Because the model must choose within the option set, every field is populated for every retained record; completion therefore carries no information about accuracy. Accuracy was assessed separately on a stratified manual sample and a held-out test set. The full extraction prompt and validation are in the Supplementary Information. Two AI tools were used in this workflow. DeepSeek Pro extracted the key information carried by the first-search article profiles and consolidated the category options, and Jev (TypeSafe System One) then labelled every record of the merged corpus. The authors reviewed and validated both sets of labels against a stratified manual sample and a held-out test set.

4.3 Classification and analysis

Countries are cross-classified on fractional publications and sampling-location articles. Publication credit uses fractional counting, the convention of the U.S. National Center for Science and Engineering Statistics (NCSES), dividing each paper among the countries of its authors. Sampling credit is assigned to the countries named as sampling locations in the abstract. Each dimension is

split at its own P25 and P75 over countries with a non-zero value, giving 26 high-high and 39 double-zero countries. Concentration is measured by the Gini coefficient with 2,000 bootstrap replicates. Scaling is estimated by ordinary least squares on country-year observations, by a mixed model with country-level random slopes, and by a Tweedie generalised additive model for the relation between publication and sampling. Change points were located by binary segmentation with the number of segments selected by BIC, and a five-layer network traces each record from country class through archive, indicator and question to goal. The detailed specification of every test, including variable definitions, model selection and robustness checks, is given in the Supplementary Information.

Code and data availability

The bibliographic corpus itself cannot be redistributed because of the Web of Science licence. The derived record-level annotation table, the complete search string, the extraction prompt and the manual validation sample are archived on Zenodo at <https://doi.org/10.5281/zenodo.23294431> under a CC BY-NC-SA 4.0 licence. The analysis code is available at <https://github.com/honghlaries/natural-archive-bibliometrics> under an MIT licence.

Author contributions

Lingli Fan: Conceptualization, Investigation, Software, Formal analysis, Writing - original draft.
Yiming Lei: Validation, Writing - Review & Editing.
Yufei Chen: Validation, Writing - Review & Editing.
Qian Huang: Conceptualization, Resources, Validation, Writing - review & editing.
Hui Jia: Resources, Writing - review & editing.

442 Shan Jiang: Writing - review & editing.

443 Jingchun Liu: Writing - review & editing.

444 Haoliang Lu: Funding acquisition, Writing - review & editing.

445 Chongling Yan: Writing - review & editing.

446 Hualong Hong: Conceptualization, Supervision, Data Curation, Funding acquisition, Writing -
447 review & editing.

448 **Competing interests**

449 The authors declare that they have no known competing financial interests or personal relationships
450 that could have appeared to influence the work reported in this paper.

451 **Declaration of generative AI use**

452 During the preparation of this work the authors used two AI tools, and both were applied to
453 bibliographic metadata rather than to the scientific content. DeepSeek Pro was used on the corpus
454 returned by the first Web of Science search to extract key information from titles and abstracts and to
455 consolidate the category options that define the article profile types described in the Methods,
456 namely sampling region, sample type, indicator type, scientific question and the Sustainable
457 Development Goals addressed. Jev (TypeSafe System One) was then used to label every record of
458 the merged corpus, 58,201 records in total, against those fixed options, returning archive carrier at
459 major and sub-class level, proxy indicator, scientific question, SDG framing, sampling location and
460 an in-scope judgement; the 33,145 records retained after screening carry the labels analysed here.
461 The authors reviewed and validated the extracted labels against a stratified manual sample and a
462 held-out test set, and take full responsibility for the content of the publication. Generative AI tools

were not used to create or alter any figure, image or other artwork, and were not used to draft or edit the text of the manuscript.

Financial support

This work was funded by the National Key R&D Program of China (2022YFF0803100) and the National Natural Science Foundation of China (U25A20801).

Acknowledgements

The authors thank the Web of Science Core Collection (Clarivate) and OpenAlex for access to the bibliographic records analysed in this study, and the developers of the R and Python packages used in the analysis.

References

- Abram, N. J., McGregor, H. V., Tierney, J. E., Evans, M. N., McKay, N. P., and Kaufman, D. S.: Early onset of industrial-era warming across the oceans and continents, *Nature*, 536, 411–418, <https://doi.org/10.1038/nature19082>, 2016.
- Aria, M., and Cuccurullo, C.: bibliometrix: An R-tool for comprehensive science mapping analysis, *J. Informetr.*, 11, 959–975, <https://doi.org/10.1016/j.joi.2017.08.007>, 2017.
- Asase, A., Mzumara-Gawa, T. I., Owino, J. O., Peterson, A. T., and Saupe, E.: Replacing “parachute science” with “global science” in ecology and conservation biology, *Conservation Science and Practice* 4, <https://doi.org/10.1111/csp2.517>, 2021.
- Blaauw, M., and Christen, J. A.: Flexible paleoclimate age-depth models using an autoregressive gamma process, *Bayesian Anal.*, 6, 457–474, <https://doi.org/10.1214/11-ba618>, 2011.
- Braconnot, P., Harrison, S. P., Kageyama, M., Bartlein, P. J., Masson-Delmotte, V., Abe-Ouchi, A., Otto-Bliesner, B., and Zhao, Y.: Evaluation of climate models using palaeoclimatic data, *Nat. Clim. Change*, 2, 417–424, <https://doi.org/10.1038/nclimate1456>, 2012.

Bradley, R. S.: Paleoclimatology, 2015.

Bronk Ramsey, C.: Bayesian analysis of radiocarbon dates, *Radiocarbon*, 51, 337–360, <https://doi.org/10.1017/s0033822200033865>, 2009.

Carroll, S. R., Garba, I., Figueroa-Rodriguez, O. L., Holbrook, J., Lovett, R., Materechera, S., Parsons, M., Raseroka, K., Rodriguez-Lonebear, D., Rowe, R., Sara, R., Walker, J. D., Anderson, J., and Hudson, M.: The CARE principles for Indigenous data governance, *Data Sci. J.*, 19, 43, <https://doi.org/10.5334/dsj-2020-043>, 2020.

Castañeda, I. S., and Schouten, S.: A review of molecular organic proxies for examining modern and ancient lacustrine environments, *Quat. Sci. Rev.*, 30, 2851–2891, <https://doi.org/10.1016/j.quascirev.2011.07.009>, 2011.

Chen, C.: CiteSpace II: Detecting and visualizing emerging trends and transient patterns in scientific literature, *J. Assoc. Inf. Sci. Technol.*, 57, 359–377, <https://doi.org/10.1002/asi.20317>, 2006.

Croudace, I. W., and Rothwell, R. G.: Micro-XRF studies of sediment cores, <https://doi.org/10.1007/978-94-017-9849-5>, 2015.

Damerow, J., Queste, B. Y., Vogt, M., and Novellino, A.: Opening doors to physical sample tracking and attribution in Earth and environmental sciences, *Sci. Data*, 12, 1047, <https://doi.org/10.1038/s41597-025-05295-z>, 2025.

Dangles, O., Struelens, Q., Ba, M. P., Bonzi-Coulibaly, Y. L., Charvis, P., Emmanuel, E., González Almario, C., Hanich, L., Koita, O., León-Velarde, F., Mburu, Y. K., Ntoumi, F., Restrepo, S., and Vidal, L.: Insufficient yet improving involvement of the global south in top sustainability science publications, *PLOS ONE* 17, <https://doi.org/10.1371/journal.pone.0273083>, 2022.

de Solla Price, D. J.: Little science, big science, Columbia University Press, <https://doi.org/10.7312/pric91844>, 1963.

Deep-time Digital Earth: DDE Platform Global Launch at the DDE Open Science Forum held at the UNESCO HQ in Paris, Deep-time Digital Earth, <https://ddeworld.org/news/detail/189>, 2022.

Diepenbroek, M., Grobe, H., Reinke, M., Schindler, U., Schlitzer, R., Sieger, R., and Wefer, G.: PANGAEA—an information system for environmental sciences, *Comput. Geosci.*, 28, 1201–1210, [https://doi.org/10.1016/S0098-3004\(02\)00039-0](https://doi.org/10.1016/S0098-3004(02)00039-0), 2002.

Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., and Lim, W. M.: How to conduct a bibliometric analysis: An overview and guidelines, *J. Bus. Res.*, 133, 285–296, <https://doi.org/10.1016/j.jbusres.2021.04.070>, 2021.

Edwards, R. L., Chen, J. H., and Wasserburg, G. J.: ^{238}U - ^{234}U - ^{230}Th - ^{232}Th systematics and the precise measurement of time over the past 500,000 years, *Earth Planet. Sci. Lett.*, 81, 175–192, [https://doi.org/10.1016/0012-821X\(87\)90154-3](https://doi.org/10.1016/0012-821X(87)90154-3), 1987.

515 Esper, J., Cook, E. R., and Schweingruber, F. H.: Low-frequency signals in long tree-ring chronologies for reconstructing
 516 past temperature variability, *Science*, 295, 2250–2253, <https://doi.org/10.1126/science.1066208>, 2002.

517 European Geosciences Union: EGU statement on scientific neocolonialism, EGU, 2023.

518 Fairchild, I. J., Smith, C. L., Baker, A., Fuller, L., Spötl, C., Matthey, D., and McDermott, F.: Modification and
 519 preservation of environmental signals in speleothems, *Earth-Sci. Rev.*, 75, 105–153,
 520 <https://doi.org/10.1016/j.earscirev.2005.08.003>, 2006.

521 Grissino-Mayer, H. D., and Fritts, H. C.: The international tree-ring data bank: an enhanced global database serving the
 522 global scientific community, *Holocene*, 7, 235–238, <https://doi.org/10.1177/095968369700700212>, 1997.

523 Haelewaters, D., Hofmann, T. A., and Romero-Olivares, A. L.: Ten simple rules for Global North researchers to stop
 524 perpetuating helicopter research in the Global South, *PLOS Computational Biology* 17,
 525 <https://doi.org/10.1371/journal.pcbi.1009277>, 2021.

526 International Ocean Discovery Program: Sample, Data, and Obligations Policy, IODP, 2024.

527 Kaufman, D., McKay, N., Routson, C., Erb, M., Dätwyler, C., Sommer, P. S., Heiri, O., and Davis, B.: A global database
 528 of Holocene paleotemperature records, *Sci. Data*, 7, 115, <https://doi.org/10.1038/s41597-020-0445-3>, 2020.

529 King, D. A.: The scientific impact of nations, *Nature*, 430, 311–316, <https://doi.org/10.1038/430311a>, 2004.

530 Liu, T., and Ding, Z.: Chinese loess and the paleomonsoon, *Annu. Rev. Earth Planet. Sci.*, 26, 111–145,
 531 <https://doi.org/10.1146/annurev.earth.26.1.111>, 1998.

532 Masson-Delmotte, V., Schulz, M., Abe-Ouchi, A., Beer, J., Ganopolski, A., and González Rouco, J. F.: Information from
 533 paleoclimate archives, In: Stocker, <https://doi.org/10.1017/cbo9781107415324.013>, 2013.

534 May, R. M.: The scientific wealth of nations, *Science*, 275, 793–796, <https://doi.org/10.1126/science.275.5301.793>, 1997.

535 Nature: Don’t rush rules for sharing digital genetic-sequence information, *Nature*, 634, 513–514,
 536 <https://doi.org/10.1038/d41586-024-03331-3>, 2024.

537 Nature Geoscience: Globalize geoscience, *Nat. Geosci.*, 8, 491, <https://doi.org/10.1038/ngeo2485>, 2015.

538 Past Global Changes: PAGES Data Guidelines, PAGES, 2018.

539 Raja, N. B., Dunne, E. M., Matiwan, A., Khan, T. M., Nätscher, P. S., Ghilardi, A. M., and Chattopadhyay, D.: Colonial
 540 history and global economics distort our understanding of deep-time biodiversity, *Nat. Ecol. Evol.*, 6, 145–154,
 541 <https://doi.org/10.1038/s41559-021-01608-8>, 2022.

542 Ramirez-Castaneda, V., West, A., and Arias, A.: A set of principles and practical suggestions for equitable fieldwork in
 543 biology, *PNAS*, 119, e2122667119, <https://doi.org/10.1073/pnas.2122667119>, 2022.

544 Secretariat of the Convention on Biological Diversity: Convention on Biological Diversity, SCBD, 1992.

545 Secretariat of the Convention on Biological Diversity: Nagoya Protocol on Access to Genetic Resources and the Fair and

546 Equitable Sharing of Benefits Arising from their Utilization to the Convention on Biological Diversity, SCBD, 2011.

547 Secretariat of the Convention on Biological Diversity: Kunming-Montreal Global Biodiversity Framework, SCBD, 2022.

548 Secretariat of the Convention on Biological Diversity: Digital sequence information on genetic resources, decision

549 adopted at the sixteenth meeting of the Conference of the Parties to the Convention on Biological Diversity, SCBD, 2024.

550 Sett, et al.: Harmonize rules for digital sequence information benefit-sharing across UN frameworks, *Nat. Commun.*, 15,

551 8745, <https://doi.org/10.1038/s41467-024-52994-z>, 2024.

552 Springer Nature: APC waivers and discounts, Springer Nature, [https://www.springernature.com/gp/open-](https://www.springernature.com/gp/open-science/policies/journal-policies/apc-waiver-countries)

553 [science/policies/journal-policies/apc-waiver-countries](https://www.springernature.com/gp/open-science/policies/journal-policies/apc-waiver-countries), 2026.

554 Stefanoudis, P. V., Licuanan, W. Y., Morrison, T. H., Talma, S., Veitayaki, J., and Woodall, L. C.: Turning the tide of

555 parachute science, *Current Biology* 31, <https://doi.org/10.1016/j.cub.2021.01.029>, 2021.

556 Stephenson, M. H., Cheng, Q., Wang, C., Fan, J., and Oberhänsli, R.: Progress towards the establishment of the IUGS

557 Deep-time Digital Earth (DDE) programme, *Episodes*, 43, 1057–1062, <https://doi.org/10.18814/epiiugs/2020/020057>,

558 2020.

559 Tierney, J. E., Poulsen, C. J., Montañez, I. P., Bhattacharya, T., Feng, R., and Ford, H. L.: Past climates inform our future,

560 *Science* 370, <https://doi.org/10.1126/science.aay3701>, 2020.

561 Trisos, C. H., Auerbach, J., and Katti, M.: Decoloniality and anti-oppressive practices for a more ethical ecology, *Nat.*

562 *Ecol. Evol.*, 5, 1205–1212, <https://doi.org/10.1038/s41559-021-01460-w>, 2021.

563 UNESCO: UNESCO Recommendation on Open Science, UNESCO, <https://doi.org/10.54677/mnmh8546>, 2021a.

564 UNESCO: Persistent divides on science threaten to leave developing nations behind, new UNESCO report shows,

565 UNESCO, [https://www.unesco.org/en/articles/persistent-divides-science-threaten-leave-developing-nations-behind-new-](https://www.unesco.org/en/articles/persistent-divides-science-threaten-leave-developing-nations-behind-new-unesco-report-shows)

566 [unesco-report-shows](https://www.unesco.org/en/articles/persistent-divides-science-threaten-leave-developing-nations-behind-new-unesco-report-shows), 2026.

567 UNESCO: UNESCO Science Report, The Race Against Time for Smarter Development. UNESCO,

568 <https://doi.org/10.18356/9789210058575>, 2021b.


569 UNESCO-IOC: United Nations Decade of Ocean Science for Sustainable Development (2021-2030) Implementation

570 Plan, IOC Ocean Decade Series 20. UNESCO, 2021.

571 United Nations: Agreement under the United Nations Convention on the Law of the Sea on the Conservation and

572 Sustainable Use of Marine Biological Diversity of Areas beyond National Jurisdiction, United Nations, 2023.

573 Williams, et al.: The Neotoma Paleoecology Database, a multiproxy, international, community-curated data resource,
 574 Quat. Res., 89, 156-177, <https://doi.org/10.1017/qua.2017.105>, 2018.
 575 World Meteorological Organization: Systematic Observations Financing Facility (SOFF), WMO, 2022.
 576 Yang, W., and Huang, J.: A global-local strategy for global open access, Innovation, 6, 100799,
 577 <https://doi.org/10.1016/j.xinn.2025.100799>, 2025.
 578 Zemp, M., Huss, M., Thibert, E., Eckert, N., McNabb, R., and Huber, J.: Global glacier mass changes and their
 579 contributions to sea-level rise from 1961 to 2016, Nature, 568, 382–386, <https://doi.org/10.1038/s41586-019-1071-0>,
 580 2019.
 581 Zupic, I., and Čater, T.: Bibliometric methods in management and organization, Organ. Res. Methods, 18, 429–472,
 582 <https://doi.org/10.1177/1094428114562629>, 2015.



583

584

585

586

587

588

589

Figure 1. The global geography of production and sampling, and how the gap varies by archive sub-class.

(a) Nine-class classification of 195 countries on fractional publications and sampling-location articles; grey denotes entities outside the sovereign roster. (b) Annual publications and sampling-location articles per country within each class, plotted over the full 1980-2025 window on a common axis with years of no output filled as zero, so that the year in which each class first appears is directly

comparable across panels; the denominator is every country in the class including those with no output in that year, panels are the nine classes, titled with the class size, and empty classes (HL, LH) carry no observations. (c) Lorenz curves of fractional publications and sampling-location articles across countries, drawn on a 1:1 scale. (d) Gini coefficient of sampling-location articles (circles) and of fractional publications (triangles) across the 195-country roster, for each archive sub-class with at least 100 records; the horizontal segment gives the gap. (e) Median difference between a country's share of sampling-location articles and its share of fractional publications for the same sub-class, plotted against the number of countries in which that sub-class is sampled; point area is proportional to the number of records, and values below the dashed line mean that sampling countries hold a smaller share of the literature than of the sampling. In (e), the polar ice-core point sits low because Antarctica and Greenland are not sovereign countries and therefore fall outside the 195-country roster. Axis units, articles per country per year in (b); cumulative share (%) in (c); Gini coefficient (index, dimensionless) in (d); percentage points and number of countries (count) in (e).






Figure 2. Income structure and the scaling step.

(a) Publications and sampling-location articles by World Bank income group (log10 scale). (b) Country-level fractional publications against sampling-location articles on log1p axes, because both counts contain zeros. The grey dashed line is a single ordinary least squares fit; the red curve is a Tweedie generalised additive model (compound Poisson-gamma, which accommodates the zero mass) with the smoothness chosen by generalised cross-validation and the shaded band giving the 95% interval of the fitted mean. The response is zero-inflated, 85 of 195 countries record no sampling, which is why the curve is fitted on the original count scale rather than on the transformed one. (c) Median residual from the fit in (b), grouped by fractional output; bars are 95% bootstrap

614 intervals, the dashed line is zero and n gives the number of countries in each group. (d) Country-year
615 observations with a World Bank PPP series (grey), with the ordinary least squares fit of publications
616 per million inhabitants on GDP per capita in constant 2021 PPP international dollars. Coloured paths
617 trace one representative per continent, selected as the country with the largest change in income and
618 in publications per million among those with at least 30 annual observations, or, where no country on
619 the continent reaches that threshold, that continent's best-covered country. China, Germany, Canada,
620 Australia, Brazil and South Africa. Paths are coloured by sampling year, from blue (earliest) to red
621 (latest). (e) Country-specific elasticities against mean income; the two background bands and
622 horizontal fit lines are colour-coded by regime, blue below the breakpoint of 41,458 PPP
623 international dollars and red above it, where the elasticity jumps by 3.96 (HC3 $p = 0.0009$). Bars are
624 95% confidence intervals and the dashed grey line is the pooled exponent from (d). Axis units,
625 records in (a); articles in (b); $\log_{10}$ residual, dimensionless, in (c); articles per 10^6 people in (d);
626 dimensionless in (e).

627



631 (a) Flow runs from country tier through archive class and indicator class to cognitive goal and
632 Sustainable Development Goal. Every record carrying all five dimensions contributes a total weight
633 of 1: it is split across the countries of its authors, so each layer conserves the record count and the
634 widths are directly comparable. Node labels give flow in records and each node's share within its
635 layer. (b) Share of each goal's flow that originates in high-high countries, over all eleven goals with
636 non-zero flow. Bars are drawn to a common length so that the split, not the magnitude, carries the
637 comparison; the flow reaching each goal is given numerically at the right. Every goal draws between
638 81.9% and 90.1% of its evidence from the twenty-six high-high countries, and the least-developed
639 classes contribute between 0.00% and 0.12% in every case.

640

1 **Supplementary Information**

2

3

4 Lingli Fan ^a, Yiming Lei ^a, Yufei Chen ^{a,b}, Qian Huang ^{d*}, Hui Jia ^e, Shan Jiang ^f, Jingchun Liu ^a,
5 Haoliang Lu ^a, Chongling Yan ^a, Hualong Hong ^{a,b,c*}

6 **Affiliations**

7 a. Key Laboratory of Ministry of Education for Coastal and Wetland Ecosystems, Xiamen University, Xiamen 361102,
8 China

9 b. Coastal and Ocean Management Institute (COMI), Xiamen University, Xiamen 361102, China

10 c. National Observation and Research Station for the Taiwan Strait Marine Ecosystem (Xiamen University), Zhangzhou
11 363305, China

12 d. Institute of Geosciences, University of Mainz, Johann-Joachim-Becher-Weg 21, Mainz 55128, Germany

13 e. School of the Environment and Safety Engineering, Jiangsu University, Zhenjiang 212013, China

14 f. State Key Laboratory of Estuarine and Coastal Research, East China Normal University, Shanghai 200241, China

15

16 **Author email addresses and ORCID**

17 Lingli Fan (a) — 33120251152187@stu.xmu.edu.cn

18 Yiming Lei (a) — 33120241152277@stu.xmu.edu.cn

19 Yufei Chen (a,b) — 33120241152351@stu.xmu.edu.cn

20 Qian Huang (d) * — qhuang01@uni-mainz.de

21 Hui Jia (e) — 1000005403@ujs.edu.cn

22 Shan Jiang (f) — sjiang@sklec.ecnu.edu.cn

23 Jingchun Liu (a) — liujingchun@xmu.edu.cn

24 Haoliang Lu (a) — luhl@xmu.edu.cn, 0000-0003-0720-3137

25 Chongling Yan (a) — ycl@xmu.edu.cn

26 Hualong Hong (a,b,c) * — honghl@xmu.edu.cn, 0000-0003-4284-8096

27

28

29 *Supplementary text*

30 **S1. Search string**

31 TS=(chronology OR chronological OR "sediment core\$" OR "lake core\$" OR "ice core\$" OR "tree
32 ring\$" OR dendrochronolog* OR "coral skeleton*" OR "speleothem*" OR stalagmite* OR stalactite*
33 OR "cave deposit*" OR "peat core\$" OR varve* OR "laminated sediment*" OR loess OR "packrat
34 midden" OR coprolite* OR "mollusk shell\$" OR otolith* OR "ground ice" OR "permafrost core\$")
35 AND TS=(proxy OR proxies OR surrogate OR biomarker\$ OR isotope\$ OR "environmental DNA"
36 OR "ancient DNA" OR "sedimentary ancient DNA" OR sedaDNA OR eDNA OR aDNA OR
37 "geochemi*" OR "element* ratio" OR pollen OR phytolith* OR diatom* OR foraminifera* OR
38 "biogenic silica" OR BSi OR ostracod*)
39 AND DT = ("Article" OR "Early Access" OR Review)

40 **S2. Extraction summary**

41 Article profile types were derived from the titles and abstracts of the first Web of Science search (31
42 December 2025) by clustering on sampling region, sample type, indicator type, scientific question and
43 Sustainable Development Goals addressed, and DeepSeek Pro was used to extract the key information
44 carried by those profiles and to consolidate the category options. Jev (TypeSafe System One) then
45 labelled every record of the merged corpus against that fixed option set, returning a calibrated
46 distribution for each of ten fields, covering article type, archive carrier (major class and sub-class),
47 proxy indicator, scientific question, SDG framing, sampling location and an in-scope judgement.
48 Because Jev must choose from the option set, every field is populated for every record and completion
49 is 100.00% for all ten fields by construction (Table S3); completion therefore carries no information
50 about accuracy. Accuracy was assessed separately against a stratified manual sample and a held-out
51 testset. Labels produced October 2026.

52 **S3. Growth of the field and its developmental periods**

53 Annual output rose from 26 articles in 1980 to 1,978 articles in 2025. It reached a local maximum of
54 1933 articles in 2021, dipped for two years, and has averaged 1,857.8 articles per year over 2021-2025,
55 returning to a new high of 1978 in 2025. BIC-selected binary segmentation identifies three change
56 points (2000, 2009, 2017) and therefore four periods, with 2,489 records in 1980-1999, 5,685 in 2000-
57 2008, 8,471 in 2009-2016 and 15,319 in 2017-2025. Sediment archives dominate every period but their
58 share declines monotonically, from 70.87% to 70.6%, while the share of labelled records the dictionary
59 cannot classify grows from 1.29% to 2.75%. (Figure S7; Tables S2, S3, S4).

60 Composition changes within periods as well as between them. Sediment archives fall from 70.87% of
61 records in 1980-1999 to 70.6% in 2017-2025; biological-growth archives rise from 9.44% to 14.82%;
62 ice-core archives decline from 12.58% to 8.03%. Research articles dominate every period and their
63 share slips slightly, from 91.76% in 1980-1999 to 89.12% in 2017-2025, while review articles fall from
64 4.38% to 3.48% and methodological papers rise steadily from 1.00% in 1980-1999 to 1.44% in 2000-
65 2008 and 1.77% in 2017-2025. The declining review share indicates that synthesis has not kept pace
66 with primary output. (Figure S7)

67 The plateau since 2017 is therefore not saturation of the field but saturation of its largest component.
68 Because the aggregate curve is largely a sediment curve, its flattening describes the largest archive
69 class rather than the field as a whole, and the growth that continues is in the classes whose share rises
70 within the same periods.

71 The boundaries are legible in the same terms (Figure S4). The first coincides with the peak share of
72 methodological papers (1.44% in 2000-2008 against 1.0% in 1980-1999), consistent with
73 improvements in dating precision and the standardisation of proxy calibration preceding the large-scale
74 adoption of new carriers (Bronk Ramsey, 2009; Edwards et al., 1987; Blaauw and Christen, 2011); the
75 second, near 2009, with the expansion of community data networks and the maturation of tree-ring
76 programmes (Grissino-Mayer and Fritts, 1997; Esper et al., 2002); and the third with the point at which
77 total volume stops growing while composition keeps turning over.

78 **S4. Archive types are separate sub-fields, not interchangeable carriers**

79 Archive types are not interchangeable, because the material sets how much can be asked of it. Sediment,
80 the least specialised carrier, spreads across chemical composition (34.1%) and biological material
81 (34.9%). Ice cores concentrate in chemical composition (67.5%), fossil and shell assemblages in
82 biological material (84.7%), and biological-growth archives in chemistry (19.0%) and physical
83 properties such as ring width and density (1.2%). A carrier that can host many indicator classes
84 supports many questions; one that preserves a single signal supports few.

85 That difference in breadth, rather than any judgement about scientific value, is why sediment carries
86 70.85% of network flow. The field's aggregate properties - its growth, its geography, its question
87 profile - are therefore to a large extent properties of its most permissive archive, and reporting them as
88 properties of paleoenvironmental science is a category error that the five-layer network makes visible,
89 because the same statistics computed on ice cores alone would describe a different field.

90 The archive classes are also separate political economies, and the distance between them is measurable.
91 Across sub-classes the concentration of sampling ranges from 0.8559 to 0.9742, the number of
92 countries in which a class is sampled from 33 to 115, and the sampling-publication wedge is present in
93 some classes and absent or reversed in others. Tree rings, lake sediments and marine sediments are
94 sampled by dispersed communities in more than eighty countries; polar ice cores are sampled in 39 and
95 mostly outside sovereign territory. A single instrument aimed at natural archives would land on those
96 two groups very differently, and that contrast is what the polar case exposes.

97 **S5. What is measured and what is asked: indicator and question composition**

98 Indicator composition tracks archive physics, though less simply than the literature sometimes implies
99 (Figure S2); how the archive classes divide across cognitive goals is shown in Figure S1. Ice-core
100 records are dominated by chemical composition, which carries 67.5% of their flow. Biological-growth
101 archives divide between chemical composition (19.0%) and physical properties such as ring width and
102 density (1.2%), so they are not primarily physical. Fossil and shell assemblages are led by biological
103 material (84.7%), with chemistry second at 12.6%. Sediment archives, the least specialised carrier,
104 spread across chemical (34.1%) and biological indicators (34.9%), which is what continuous non-
105 destructive core scanning makes possible (Croudace and Rothwell, 2015). This is what the material
106 permits, since a tree ring measures growth, an ice core measures atmospheric chemistry, and a fossil
107 assemblage measures assemblage composition, because the signal each proxy carries is set by the
108 carrier that preserves it (Fairchild et al., 2006; Castañeda and Schouten, 2011).

109 Scientific questions, by contrast, do not track archive physics at all. Reconstruction and inversion takes
110 33.49% of network flow and mechanism and drivers 24.26% (Table S1); the network excludes records
111 with an unresolved layer, and the same ordering holds at record level. Prediction and simulation, where
112 societal value would arguably be highest, takes 0.33%. The two leading goals together account for
113 48.9% to 72.7% of flow depending on the archive class (Figure S2), a spread of twenty-four points that
114 is far narrower than the variation in indicator composition across the same classes. Three sources are
115 plausible and this corpus cannot separate them, as questions are set by training and by the reward
116 structure of the field rather than by the material at hand; reconstruction is the step that makes any
117 archive publishable, so it is a precondition of publication rather than a choice among alternatives; and
118 prediction requires model-data fusion infrastructure concentrated in the countries that already dominate
119 output, which would make the low share of prediction a capacity signal rather than a preference. What
120 the data do establish is the asymmetry itself: what can be measured is set by the archive, and what is
121 asked is set by the community.

122 **S6. Methods: data, extraction and statistical specification**

123 **Data source and search strategy.**

124 Records were assembled from two bibliographic sources. The primary source is the Web of Science
125 Core Collection (Clarivate), queried on 31 December 2025 and again on 6 October 2026 with an
126 expanded concept block and exported as field-tagged plain text; the document-type filter was kept to
127 Article, Early Access and Review, so proceedings, corrections and editorial material are excluded. To
128 test whether the resulting picture depends on that single index, the same concept blocks were then run
129 against OpenAlex, which indexes a substantially larger and more multilingual body of literature. Three
130 concept blocks were combined with AND in both sources, natural archives and chronometric carriers;
131 proxy indicators; and global change. After de-duplication on DOI, normalised title and accession
132 number, the union contains 33,145 records, of which 9,126 come from the Web of Science Core
133 Collection, 11,663 from the expanded Web of Science query and 12,356 from OpenAlex. The finer
134 type labels reported below (Method, Case Study, Data Descriptor and the like) are the article types
135 recorded in the records themselves, not a further filter. 1,181 records carry a publication year outside
136 the 1980-2025 window and are retained in the corpus total but excluded from every annual and period
137 series, so the analysed time series contains 31,964 records. Research Article accounts for 30,109
138 records (90.84%), Review for 1293, Method for 527 and Case Study for 250. The remaining 966
139 records are spread over six smaller types, the largest being Comment/Reply (20) and Editorial/Letter
140 (12), so the four named types cover 97.08% of the corpus.

141 Retrieval followed a PRISMA-style flow, reported in full in Table S7 and Figure S6. The original Web
142 of Science Core Collection search returned 11,174 records; an expanded Web of Science query
143 returned a further 42,063; and the same concept blocks run against OpenAlex returned 27,014. After
144 de-duplication on DOI, normalised title and accession number, 49,189 unique records remained, of
145 which 11,055 came from the original Web of Science search, 18,449 from the expanded Web of
146 Science query and 19,685 from OpenAlex. Every record was then screened on title and abstract against
147 a single question, does the study analyse a natural environmental archive as its primary evidence?, and
148 16,044 records were excluded because the words that had matched the search blocks were used in
149 another sense, or because the material studied was not a natural archive. Screening retained 33,145
150 records, of which 1,181 carry a publication year outside the 1980-2025 window and are held in the
151 corpus total but excluded from every annual and period series.

152 **Field extraction with a large language model.**

153 Ten fields were extracted from titles and abstracts with Jev (TypeSafe System One), a decision model
154 that returns a calibrated probability distribution over a fixed option set, article type, archive carrier
155 (major class and sub-class), proxy indicator, scientific question, SDG framing, sampling location and
156 an in-scope judgement. Because the model must choose from the option set, every field is populated for
157 every record; completion is therefore 100.00% to 100.00% across the ten fields by construction (Table
158 S5) and carries no information about accuracy. Accuracy was instead assessed against a stratified
159 manual sample and against a held-out decision testset (Supplementary Methods S2). A separate
160 screening pass applied the same model to a single question, is this article about a natural environmental
161 archive?, and its verdicts define the corpus reported here (.

162 **Classification of countries, archives and indicators.**


163 Countries are cross-classified on fractional publications and sampling-location articles. Publication
164 credit is assigned by fractional counting, the convention used by the U.S. National Center for Science
165 and Engineering Statistics (Table S6), in which each paper is divided among the countries of its authors
166 in proportion to how many authors each country contributes, so that a paper with two authors in one
167 country and one each in two others awards half the credit to the first and a quarter to each of the others.
168 Whole counting, which gives every participating country a full point, would inflate countries that
169 collaborate widely, while crediting only the first author would systematically undercount countries that
170 take part as co-authors. Each dimension is split at its own P25 and P75, computed over the countries
171 with a non-zero value on that dimension (153 countries for publications, 110 for sampling); values at or
172 below the P25 form the low band, values between the P25 and the P75 the middle band, and values
173 above the P75 the high band. A country with no publication therefore falls in the low band without
174 entering the quantile calculation, which prevents the breaks from being dragged towards zero by the 53
175 countries that have neither a publication nor a sampling-location article in the corpus. Sampling-
176 location articles are counted at the record level after normalising the free-text sampling description to
177 standard country names; this metric reproduces the archived country table exactly for all 195 countries
178 (sum of absolute differences = 0). Archive classes follow an explicit dictionary of 18 sub-classes
179 mapped to five top-level classes; records matching no sub-class are reported as unclassified rather than
180 forced into a class. The same dictionary carries an intermediate level of 19 archive sub-classes;.
181 Indicators are mapped to the what branch of the PaST (the Paleoenvironmental Standard Terms
182 thesaurus, a controlled vocabulary for palaeoenvironmental proxies and their measured properties).
183 Scientific questions are classified by cognitive goal.

184 **Temporal segmentation and statistical analysis.**


185 Annual publication volume for 1980-2025 was smoothed with a three-year centred moving average.
186 Change points were located by binary segmentation, recursively choosing the split that minimises the
187 within-segment sum of squares with a minimum segment length of five years, and the number of
188 change points was selected over $k = 0-6$ by the Bayesian Information Criterion with a per-change-point
189 penalty of $2 \ln n$. The BIC surface has no interior minimum at $k = 2$; $k = 3$ and $k = 4$ differ by less than
190 five BIC units and we adopt $k = 3$. The unrounded change points it returns are 1998, 2009 and 2019;
191 for the period summaries we report them rounded to the nearest year boundary as 2000, 2009 and 2017,
192 giving four developmental periods. The full BIC table and a ten-configuration sensitivity analysis are
193 reported in Table S6 and Table S7.

194 A logistic curve was fitted to the annual series, giving $K = 2548$, $r = 0.114$, $t_0 = 2015$ and $R^2 = 0.9816$;
195 observed output peaked at 1978 articles in 2025 and has averaged 1857.8 articles per year since 2021.
196 The fitted capacity is 2,548 against the 1857.8 articles per year averaged since 2021, so the field sits at
197 about 73% of its fitted ceiling; the recent plateau describes a slowing of growth rather than saturation.
198 Three regressions describe the relation between income and output. Country-level per-capita output
199 was regressed on GDP per capita in double-logarithmic space over all country-year observations with a
200 World Bank series. Country-specific elasticities were then estimated as the slopes of the same double-
201 logarithmic relation fitted within each country with at least ten annual observations, and their
202 dependence on income was tested with a linear mixed model with country-level random slopes against
203 a random-intercepts-only alternative. A step model, in which the elasticity is flat below a breakpoint
204 and shifts by a constant above it, was fitted by grid search over the breakpoint on the 20th-80th
205 percentile range of mean income, with HC3 heteroscedasticity-consistent standard errors and a non-
206 parametric bootstrap over countries for the breakpoint interval. The same double-logarithmic fit
207 between the two country-level dimensions was used to group countries by publication output and to test
208 each group's median residual against zero, by one-sided Wilcoxon signed-rank and sign tests with
209 bootstrap intervals. Concentration is summarised by the Gini coefficient, the Herfindahl-Hirschman
210 index and top-N shares. The five-layer network apportions multi-label records with a $1/n$ weight so that
211 every record contributes exactly one unit of flow. (Table S8)

212



216 **Figure S1. Share of each archive class falling into each cognitive goal. Cells give percentages**
217 **within an archive class and the colour scale is log1p. Axis units: per cent.**




219

220 **Figure S2. Indicator composition mapped to the what branch of the PaST thesaurus. Axis units:**
 221 **per cent of records.**


222

223
224




225
226 **Figure S3. Data-quality profile: records with missing or unresolvable fields, and the unclassified**
227 **archive residual. Axis units: number of records.**
228

229
230



231
232 **Figure S4. Cognitive goal composition across the four developmental periods. Lines give each**
233 **goal's share of the records in that period (Table S3). Axis units: per cent of records within period.**
234


235
236




237
238
239
240
241
242
243

Figure S5. The two main tests evaluated on three corpus scopes, from the expanded corpus with partial OpenAlex coverage to the analysed corpus used throughout this paper. Points are Gini differences between the two archive classes named in each panel and bars are 95% bootstrap intervals; blue marks intervals that exclude zero. Record counts for each scope appear beneath the scope label (Table S7). Axis units: Gini index difference (dimensionless).

244
245



246
247 **Figure S6. PRISMA-style flow of records from retrieval through de-duplication and screening to**
248 **the analysed corpus. Tile shading tracks the number of records at each stage (Table S7). Axis**
249 **units: records (count).**
250



253 **Figure S7. Growth of the field and its developmental periods. (a) Annual output by archive class**
254 **with the fitted logistic curve and the unclassified residual; (b) BIC as a function of the number of**
255 **change points.**

257 **Supplementary tables**

258

259 **Table S1. Cognitive goal composition by developmental period.**

period	axis_a	records_n	share_pct
1980-1999	A1 Reconstruction & inversion	896	36.0
1980-1999	A2 Description & characterisation	139	5.58
1980-1999	A3 Mechanism & drivers	469	18.84
1980-1999	A4 Relationship & response	318	12.78
1980-1999	A5 Spatiotemporal pattern	413	16.59
1980-1999	A6 Assessment & decision	34	1.37
1980-1999	A7 Method development	112	4.5
1980-1999	A8 Prediction & simulation	5	0.2
1980-1999	A9 Synthesis & review	103	4.14
2000-2008	A1 Reconstruction & inversion	2069	36.39
2000-2008	A2 Description & characterisation	206	3.62
2000-2008	A3 Mechanism & drivers	1181	20.77
2000-2008	A4 Relationship & response	799	14.05
2000-2008	A5 Spatiotemporal pattern	690	12.14
2000-2008	A6 Assessment & decision	109	1.92
2000-2008	A7 Method development	343	6.03
2000-2008	A8 Prediction & simulation	24	0.42
2000-2008	A9 Synthesis & review	264	4.64
2009-2016	A1 Reconstruction & inversion	2965	35.0
2009-2016	A2 Description & characterisation	327	3.86
2009-2016	A3 Mechanism & drivers	1946	22.97
2009-2016	A4 Relationship & response	1283	15.15
2009-2016	A5 Spatiotemporal	786	9.28

2009-2016	pattern A6 Assessment & decision	231	2.73
2009-2016	A7 Method development	564	6.66
2009-2016	A8 Prediction & simulation	29	0.34
2009-2016	A9 Synthesis & review	340	4.01
2017-2025	A1 Reconstruction & inversion	5169	31.33
2017-2025	A2 Description & characterisation	786	4.76
2017-2025	A3 Mechanism & drivers	4442	26.92
2017-2025	A4 Relationship & response	2596	15.74
2017-2025	A5 Spatiotemporal pattern	1364	8.27
2017-2025	A6 Assessment & decision	352	2.13
2017-2025	A7 Method development	1158	7.02
2017-2025	A8 Prediction & simulation	53	0.32
2017-2025	A9 Synthesis & review	578	3.5

260

261

262 **Table S2. Scaling regressions under alternative specifications.**

specification	n_obs	n_countries	slope	intercept	coefficient_A	r2	spearman
full panel (all observations)	1427	97	1.0089	-4.9352	1.161e-05	0.5276	0.7422
latest year per country (cross-section)	97	97	1.0266	-4.8878	1.295e-05	0.662	0.8142
1990-2025 panel	1403	97	1.0106	-4.9294	1.177e-05	0.5426	0.7476
below threshold (<= US\$35,119)	974	93	0.8144	-4.2148	6.099e-05	0.332	0.5509
above threshold (> US\$35,119)	453	27	2.3204	-11.0527	8.858e-12	0.3953	0.6051

263

264 **Table S3. Stratified analysis robustness checks.**

item	value
records in corpus	33145.0
records with archive_mid	33145.0
records with archive_mid and sampling country	33145.0
Unclassified records (mid level)	843.0
Unclassified: year distribution chi-square p	4.87e-89
multi-label mid records (>=2)	0.0
Gini marine, primary-label only	0.8748
Gini lake, primary-label only	0.8559
difference, primary-label only	0.019
types with n < 100 (excluded from main)	3.0
polar ice core records	2824.0
polar ice core records sampled in Antarctica	1388.0
polar ice core records sampled in Greenland	642.0
polar ice core records sampled in both Antarctica and Greenland	0.0
polar ice core records sampled outside the 195-country roster	2265.0

265

266

267 **Table S4. The two archive-type contrasts across corpus scopes. The comparisons set marine**
 268 **against lake sediment and polar ice core against tree ring; difference is the Gini gap between the**
 269 **two classes with its 95% bootstrap interval. The analysed corpus is the screened union of Web of**
 270 **Science and OpenAlex used throughout this paper.**

comparison		corpus	records	difference	ci_low	ci_high	ci_excludes_zero
Marine	vs	expanded	20255	0.0262	0.0185	0.0324	True
lake		corpus					
sediment		(partial					
		OpenAlex)					
Polar	ice	expanded	20255	0.0614	0.0542	0.0656	True
core vs tree		corpus					
ring		(partial					
		OpenAlex)					
Marine	vs	expanded	37880	0.024	0.0177	0.0284	True
lake		corpus					
sediment		(complete					
		OpenAlex)					
Polar	ice	expanded	37880	0.0592	0.0529	0.0626	True
core vs tree		corpus					
ring		(complete					
		OpenAlex)					
Marine	vs	analysed	33145	0.019	0.0133	0.0237	True
lake		corpus					
sediment		(WoS +					
		OpenAlex,					
		screened)					
Polar	ice	analysed	33145	0.069	0.0627	0.0726	True
core vs tree		corpus					
ring		(WoS +					
		OpenAlex,					
		screened)					

271

272

273 **Table S5. Robustness of the elasticity step to output normalisation.**

quantity	value
correlation log population vs log GDP per capita	-0.3657
population elasticity wrt GDP per capita	-0.6858
median country elasticity, per-million basis	2.5823
median country elasticity, raw-total basis	2.7245
breakpoint, per-million basis	41458.0
breakpoint, raw-total basis	40956.0
step at 41458, per-million basis	3.9647
step at 41458, raw-total basis	4.5349
step at 41458, population basis	0.3841
step p, per-million basis	1.43e-05
step p, raw-total basis	1.15e-06
population share of the raw-total step (%)	8.5
countries with ≥ 10 annual observations	56.0
panel observations	1389.0

274

275 **Table S6. Mapping of archive sub-classes to major classes.**

archive_mid_en	archive_class_en	note
Amber	Other natural	Amber: the workbook lists it under Other natural archives
Biogenic carbonate shell	Biological growth	
Cave & speleothem	Sediment	
Coral	Biological growth	
Ferromanganese crust	Other natural	
Glacial sediment	Sediment	
Lake sediment	Sediment	
Loess & paleosol	Sediment	
Marine microfossil	Fossil & shell	
Marine sediment	Sediment	
Paleosol	Sediment	Paleosol: same material as Loess & paleosol; the workbook assigns it to Sediment
Peat	Sediment	
Polar ice core	Ice core	
Rock glacier	Other natural	Rock glacier: no such sub-class in the workbook; assigned to Other natural archives as an ice-rock mixture
Terrestrial macrofossil	Fossil & shell	
Terrestrial microfossil	Fossil & shell	
Tree ring	Biological growth	
Volcanic sediment	Sediment	

276

277

278 **Table S7. PRISMA flow of records from retrieval to analysed corpus.**

stage	records_n
Records identified from Web of Science Core Collection, original search	11174
Records identified from Web of Science Core Collection, expanded search	42063
Records identified from OpenAlex	27014
Records after de-duplication across the three sources	49189
of which from the original Web of Science search	11055
of which from the expanded Web of Science search	18449
of which from OpenAlex	19685
Records excluded at title-and-abstract screening (not natural-archive research)	16044
Records retained for the bibliometric analysis	33145
records dated outside the 1980-2025 window	1181
records with a resolvable sampling location	33145

279

280 **Table S8. Robustness of the main tests to the inclusion of reviews and syntheses.**

specifi cation	rec ord s	reviews_ included	marin e_gini	lake _gin i	marine_m inus_lake	ice_co re_gin i	tree_ri ng_gini	ice_minus _tree_ring	sampli ng_gin i	countries_wi th_sampling
all record s (adopt ed)	331 45	1293	0.874 8	0.85 59	0.019	0.9742	0.9052	0.069	0.8556	121
review s exclud ed	318 52	0	0.873 8	0.85 47	0.0192	0.9741	0.9018	0.0722	0.8524	121