Comment on “China’s solar expansion policy reduces bird diversity”

Chenchen Ding[1]*, Yang Liu[2]*, Zhengwang Zhang[3]*, Zhiheng Wang[1]*

[1] Institute of Ecology and State Key Laboratory for Vegetation Structure, Function and Construction (VegLab), College of Urban and Environmental Sciences, Peking University, Beijing 100871, China

[2] School of Ecology, Sun Yat-sen University, Shenzhen 518107, China.

[3] Key Laboratory for Biodiversity Science and Ecological Engineering, College of Life Sciences, Beijing Normal University, Beijing 100875, China.

* Corresponding authors. Email: chenchen.ding@pku.edu.cn (C.D.); liuy353@mail.sysu.edu.cn (Y.L.); zzw@bnu.edu.cn (Z.Z.); zhiheng.wang@pku.edu.cn (Z.W.)

One-sentence summary: The reported decline in bird diversity under China’s solar-expansion policy tracks birdwatching effort rather than birds: it disappears once effort is controlled, and the data cannot identify a national-scale ecological effect.

Abstract

Zhang *et al*. reported that solar-policy stringency reduced county-level bird diversity across China. Reanalyzing their replication package, we find the outcome variable tracks birdwatching effort; the reported effect disappears once effort is controlled, corrupted records removed, or imputed policy zeros dropped. The design cannot identify a national-scale ecological effect.

Main Text

Zhang *et al*. (*1*) link a photovoltaic policy-stringency index (PSI) to citizen-science bird records aggregated to county-months, concluding that a one-standard-deviation policy increase reduced Shannon diversity by 2.10%. An accompanying Perspective (*2*) and wide coverage presented this as evidence that solar expansion harms birds. We reproduce their preferred estimate exactly ($\beta = -0.0125$, clustered SE = 0.0037, n = 46,371). Two eLetters question the exposure measure and the evaluation boundary (*3*, *4*); our concern lies upstream of both: the outcome variable itself.

The outcome variable measures the observation process. Citizen-science records confound two processes: ecology (which species are present, in what numbers) and observation (who looks, where, when). Because detection is biased taxonomically, temporally, and spatially, raw records cannot be read as occupancy or abundance (*5*–*7*). The dependent variable (ShannonBD) is a Shannon index (H) computed from unstructured reports pooled by county-month. In the estimation sample, 38.8% of county-months come from a single birdwatcher, and the median county is observed in 11 of 120 months (Fig. 1A). Observed richness and diversity rise steeply with birdwatcher numbers before leveling off (Fig. 1B): county-months with one birder average 15 species (H = 1.90), those with over 100 birders, 142 (H = 3.60). Part of this gradient is real: birdwatchers congregate where birds abound, but two diagnostics indicate sampling accumulation dominates. First, in fixed-community simulations (supplementary materials), checklist-type recording alone reproduces the rise. Second, evenness falls steadily with effort (0.84 to 0.73), the mechanical signature of lengthening species lists rather than a real diversity gradient. The authors' own results show this signature: PSI is associated with lower richness but higher evenness (their table S19).

The counts themselves are unvalidated. In record-level platform data we hold, every count equals one in 46,997 of 100,313 checklists from 2014–2023 (46.9%) and in 43.7% of its 282,197 checklists for 2024, whose counts we verified as genuine abundances. Such records carry no abundance information: recording every species as one individual up-weights rare species, down-weights dominants, and drives Shannon H toward log richness. Corrupted counts reach the analysis unfiltered. In Xinhui County, March 2015 has $H = 8 \times 10^{-7}$ although four species were recorded (2.2–2.9 in adjacent months); arithmetically, this requires one species counted in the tens of millions. The value enters the regression as an outcome (Fig. 1C), as do 18 county-months with the same signature. County-month sums can also double-count birds: 15.7% of checklists come from group outings, filed separately by each observer. No data-quality screening is described in (*1*).

Birdwatching effort is confounded with the treatment. Summed birdwatcher counts grew 171-fold over 2014–2023 and observed county-months 25-fold, as mean PSI quadrupled (Fig. 1D and fig. S2). Effort also concentrates in accessible, economically developed counties (fig. S2A; supplementary materials). The authors control only for Duration, the log of hours per birder, a ratio that discards this growth; Duration is itself a post-treatment variable, responding to PSI ($\beta = -0.0177$, $P < 10^{-4}$), so conditioning on it is a classic "bad control" (*8*). Birdwatcher numbers serve as a negative-control outcome: PSI "reduces" birdwatchers ($\beta = -0.041$, $P = 2 \times 10^{-11}$), an association unlikely to reflect birdwatcher behavior. Three tests locate its source. It shrinks by 81% within policy-documented counties ($\beta = -0.008$, $P = 0.17$). Next calendar-year PSI "affects" current birdwatchers conditional on current PSI ($\beta = -0.051$, $P < 10^{-5}$), a causally impossible direction indicating differential county trends. And PSI does not predict realized photovoltaic area within counties ($\beta = -0.017$, $P = 0.27$; likewise where capacity is recorded, supplementary materials), which a construction-disturbance pathway would require. A design in which future policy predicts a non-ecological outcome can generate spurious effects for any variable, birds included.

The remedy is standard: model the observation process explicitly (*5*, *6*, *9*, *10*). Within the authors' specification, adding ln(birdwatchers) flips the PSI coefficient to +0.0032 ($P = 0.31$); ln(total hours), or both terms, gives the same null (Fig. 2A). The richness decline of −0.98 species per PSI unit (their table S19) shrinks 92% to −0.07 ($P = 0.37$); the effort-adjusted Shannon estimate for 2020–2023 is +0.0002 ($P = 0.96$).

The estimate is fragile and, at face value, negligible. The policy variable is also problematic: 703 of 2,344 counties (43.9% of observations) never appear in the policy-text source, and their PSI is imputed as zero although the source marks genuine zeros; restricting to documented counties yields $\beta = -0.0027$ ($P = 0.49$). Even so, PSI adds 0.0003 to R^2, 0.03% of the variance in ShannonBD against 31% absorbed by fixed effects (Fig. 2B), and a one-standard-deviation change in PSI shifts H by 0.052, 6.5% of the residual standard deviation. A coefficient this small, from a contaminated, effort-driven outcome, carries little ecological information (*11*, *12*).

We therefore estimated all 36 combinations of sample (full or policy-documented), cleaning (as published, 19 corrupted county-months removed, or ≥2 birders), metric (Shannon, Simpson, log richness), and effort adjustment (fig. S1). All seven significant negative estimates arise without effort terms; all 18 effort-adjusted paths are null (median +0.10%). Our preferred estimate is +0.14% for Shannon diversity (95% CI, −0.96 to +1.24%), which excludes the published −2.10%, an effect this design had 96% power to detect. Our reanalysis does not support the central conclusion of (***1***)**, that China's solar-expansion policy reduced bird diversity.**

Implications. This is not evidence that photovoltaic development is harmless: ground-mounted solar can displace habitat and warrants rigorous study. It is evidence that this dataset, under this design, cannot detect such an effect either way. The desert-subgroup null is being read as support for concentrating gigawatt-scale development in arid lands, yet deserts are where citizen-science coverage is sparsest; absence of evidence is not evidence of absence. Birdwatching records can support diversity inference under explicit conditions: effort-standardized richness is defensible with coverage-based rarefaction, list-length corrections, or occupancy models (*5*, *6*, *9*, *10*, *12*, *13*); abundance-weighted indices additionally require complete checklists and validated counts, neither of which holds here. Data defects and fitness for purpose must be established before citizen-science records anchor national-scale causal claims.

References

1. H. Zhang, A. Zhang, K. Wu, Y. Cai, S. Li, S. Wang, Y. Qiu, S. Huang, T. T. A. Phan, China's solar expansion policy reduces bird diversity. *Science* **393**, 831–836 (2026). doi:10.1126/science.aee0747
2. Y. Liang, Beyond carbon in renewable energy policy. *Science* **393**, 758–759 (2026).
3. H. Jin, Causal attribution and ecological interpretation of solar-policy effects. eLetter, *Science* (25 August 2026).

4. Y. Yang, J. Zhang, Y. Xian, S. Zhang, J. Zhang, Solar expansion needs system-wide biodiversity accounting. eLetter, *Science* (26 August 2026).
5. A. Johnston, W. M. Hochachka, M. E. Strimas-Mackey, V. Ruiz Gutierrez, O. J. Robinson, E. T. Miller, T. Auer, S. T. Kelling, D. Fink, Analytical guidelines to increase the value of community science data. *Divers. Distrib.* **27**, 1265–1277 (2021).
6. D. I. MacKenzie, J. D. Nichols, G. B. Lachman, S. Droege, J. A. Royle, C. A. Langtimm, Estimating site occupancy rates when detection probabilities are less than one. *Ecology* **83**, 2248–2255 (2002).
7. C. Ding, J. Ding, H. Qiao, Z. Jiang, Z. Wang, Taxonomic and spatiotemporal patterns and ecological correlates of new mammal distribution records in China. *Glob. Ecol. Biogeogr.* **34**, e70165 (2025).
8. C. Cinelli, A. Forney, J. Pearl, A crash course in good and bad controls. *Sociol. Methods Res.* **53**, 1071–1104 (2024).
9. S. Kelling, A. Johnston, A. Bonn, D. Fink, V. Ruiz-Gutierrez, R. Bonney, M. Fernandez, W. M. Hochachka, R. Julliard, R. Kraemer, R. Guralnick, Using semistructured surveys to improve citizen science data for monitoring biodiversity. *BioScience* **69**, 170–179 (2019).
10. N. J. B. Isaac, A. J. van Strien, T. A. August, M. P. de Zeeuw, D. B. Roy, Statistics for citizen science: extracting signals of change from noisy ecological data. *Methods Ecol. Evol.* **5**, 1052–1060 (2014).
11. L. Jost, Entropy and diversity. *Oikos* **113**, 363–375 (2006).
12. A. Chao, N. J. Gotelli, T. C. Hsieh, E. L. Sander, K. H. Ma, R. K. Colwell, A. M. Ellison, Rarefaction and extrapolation with Hill numbers: a framework for sampling and estimation in species diversity studies. *Ecol. Monogr.* **84**, 45–67 (2014).
13. J. K. Szabo, P. A. Vesk, P. W. J. Baxter, H. P. Possingham, Regional avian species declines estimated from volunteer-collected long-term data using List Length Analysis. *Ecol. Appl.* **20**, 2157–2169 (2010).

Acknowledgments

Funding: Beijing Natural Science Foundation grant 5264031 (C.D.).

Data and materials availability: All analyses use the original authors' public replication package (github.com/jianke22/China-s-solar-expansion-policy-reduces-bird-diversity). R scripts, derived tables, and figure source data reproducing every number in this Comment are openly available at github.com/dingchenchen6/reanalysis-zhang2026-solar-birds; the record-level China Birdwatching Record Center export carries observer identities and precise localities and is therefore not redistributed; it is obtainable from the platform under its own terms, and SHA-256 checksums of the versions used, together with the aggregated tables behind every quoted number, are archived in the repository. The archive will be deposited with a DOI upon acceptance.

Supplementary Materials

Materials and Methods; Figs. S1 and S2; Tables S1 and S2.

Figure Legends

Fig. 1. The dependent variable of Zhang *et al*. measures the observation process, not bird communities. (**A**) Distribution of the number of months (of 120) in which each county has any birdwatching record; median 11. (**B**) Observed Shannon H (green, left axis; mean ± s.d.) and species richness (blue, right axis) against the number of birdwatchers per county-month; sampling accumulation dominates this gradient (see text). 38.8% of county-months derive from a single birder. (**C**) Monthly Shannon H for Xinhui County, Guangdong: a corrupted entry collapses March 2015 to $H = 8 \times 10^{-7}$ with four species recorded (adjacent months, 2.2–2.9), a value that is arithmetically possible only if one species were counted in the tens of millions; it enters the regression unfiltered. Eighteen further county-months show ≥20 species yet $H < 0.2$. (**D**) Within the authors' estimation sample, summed birdwatcher counts (red) grew 171-fold during 2014–2023 and observed county-months (orange) 25-fold, while mean policy stringency (blue)

quadrupled; mean observed Shannon H (green) rose in step with effort. All series indexed to 2014 = 1, log scale.

Fig. 2. The reported policy effect vanishes under standard adjustments and is negligible in magnitude. (**A**) Effect of a one-standard-deviation PSI increase, expressed as % of the outcome mean with 95% CIs, under the authors' exact specification and minimal adjustments (identical data, county and year-month fixed effects, county-clustered SEs). Adding total-effort terms or restricting to counties documented in the policy source eliminates the effect, and within the 2020–2023 boom the effort-adjusted estimate is +0.0002 (P = 0.96); the richness outcome (triangles) shrinks 92% under effort control. Red dotted line: the headline −2.10%. (**B**) Variance decomposition for county-month Shannon H (the authors' ShannonBD): county and year-month fixed effects absorb 31.0%, the nine controls 2.7%, and PSI, the headline variable, 0.03%, with 66.3% unexplained.

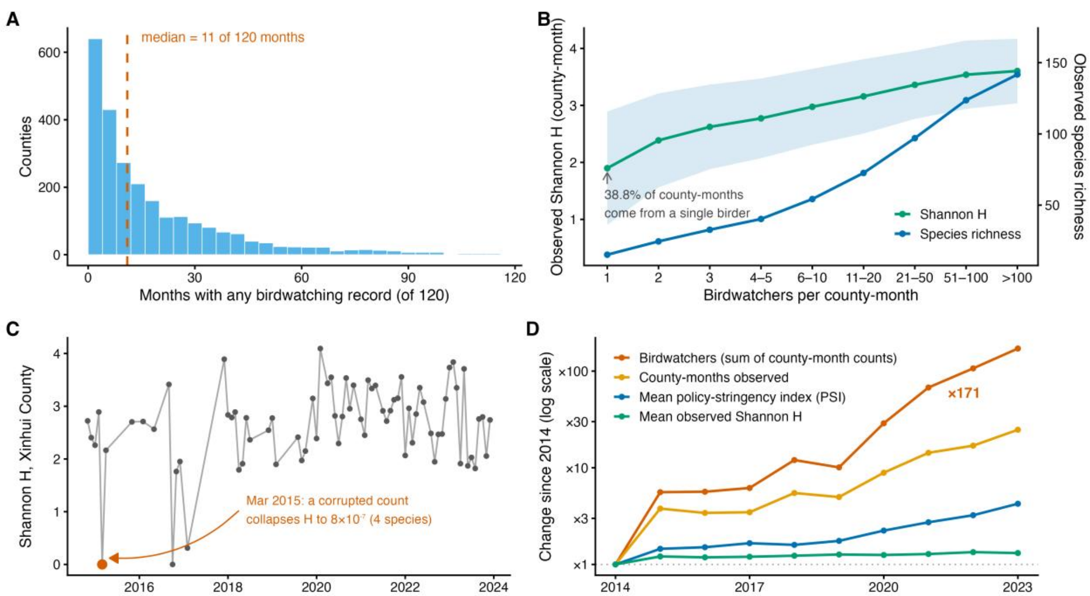

Fig. 1.

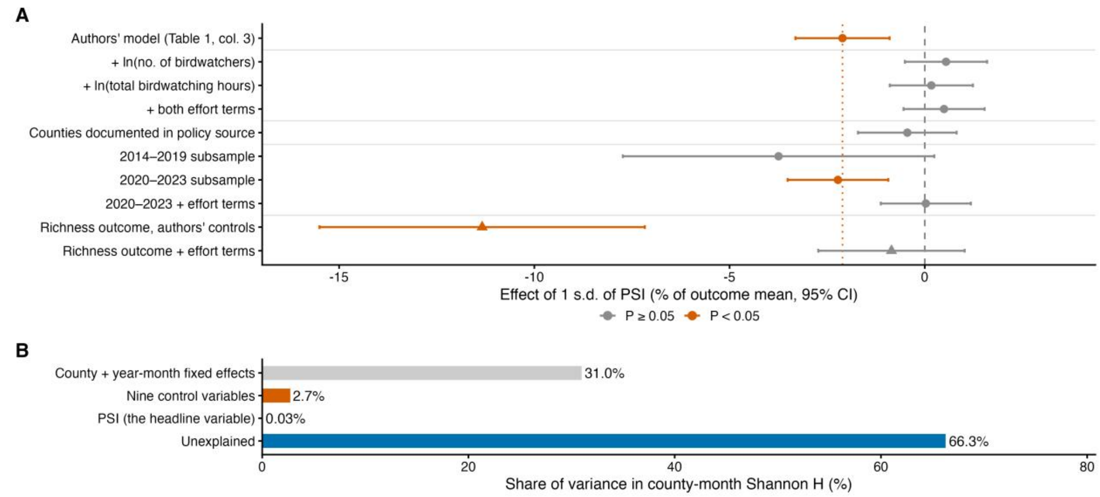

Fig. 2.

Supplementary Materials for

Comment on “China’s solar expansion policy reduces bird diversity”

Chenchen Ding. Correspondence to: chenchen.ding@pku.edu.cn

This PDF file includes: Materials and Methods; Figs. S1 and S2; Tables S1 and S2.

Materials and Methods

Data and exact replication. Variable names follow the authors’ dataset: ShannonBD and SimpsonBD are the county-month Shannon (H) and Gini–Simpson indices, PI the policy-stringency index (PSI), and Duration the log of birdwatching hours per birder. All analyses start from the replication package released by Zhang *et al*. (Stata code and county-month data; github.com/jianke22/China-s-solar-expansion-policy-reduces-bird-diversity). We rebuilt the estimation panel in R 4.5.1 (fixest, data.table, haven), following the authors’ merge order, variable construction, 1%/99% winsorization, and singleton-dropping conventions. The preferred specification regresses the county-month Shannon index on the policy-stringency index (PSI) and nine controls with county and year-month fixed effects and county-clustered standard errors. Our estimates match Table 1 of the original article to the fourth decimal place (column 3: $\beta = -0.01253$, SE = 0.00366, n = 46,371), and the headline −2.10% equals $100 \times \beta \times$ s.d.(PSI)/mean(H) = $100 \times (-0.0125) \times 4.116/2.446$.

Effort variables and adjusted specifications. The replication data include, for every county-month, the number of birdwatchers (BN) and total birdwatching hours (BT). The authors' only effort control, Duration = ln(BT/BN + 1), is a per-birder ratio and is therefore nearly invariant to growth in total effort: within the estimation sample it correlates only weakly with total effort, whereas summed BN grew 171-fold over 2014–2023. Adjusted specifications add ln(BN), ln(BT + 1), or both to the authors' model without any other change. Effects are reported as percentages of the outcome mean per 1 s.d. of PSI throughout.

Policy-source restriction and negative-control tests. The policy file covers 1,641 of the 2,344 sample counties; the authors set PSI to zero for the remaining 703 counties (43.9% of observations), although the source itself records explicit zeros where stringency is genuinely zero. The restricted sample retains policy-documented counties only. Negative-control regressions replace the outcome with ln(BN), realized photovoltaic area (converted from installed capacity as in the authors' code), or ln(nightlight + 1); a further variant adds county-specific linear trends to the authors' specification. Two details matter for interpretation. First, the county-year panel has gaps, so the pre-trend test matches each county-year explicitly to the record for YEAR + 1, rather than to the next available record, which would lie two to eight years ahead for 17.7% of county-year pairs; next calendar-year PSI is defined for 30,426 of the 46,371 rows. Second, installed capacity is missing for 54.6% of county-month rows and is set to zero in the authors' code; we retain that convention for comparability and additionally re-estimate the exposure test on the 21,054 county-months whose county-year has recorded capacity. Results are collected in Table S2.

Data-quality rules. Shannon H, Gini–Simpson D, and richness S computed from one assemblage are not independent: they must satisfy $H \leq \ln S$, $D \leq 1 - e^{-H}$ (equivalently ${}^2D \leq {}^1D$), $D \leq 1 - 1/S$, and $H > 0$ whenever $S \geq 2$. Applied to all 46,283 pre-winsorization county-months, none of these bounds is violated, so the published indices are internally consistent given their inputs. What the screen does reveal is one county-month sitting at the numerical edge of the entropy bound: Xinhui County, March 2015, with $S = 4$ and $H = 8.06 \times 10^{-7}$ against 2.2–2.9 in adjacent months. Such a value is not mathematically impossible, but it is arithmetically extreme. In the least extreme configuration, with each of the three minor species at one individual, $H = 8.06 \times 10^{-7}$ implies $3 \ln(N)/N \approx H$, hence a dominant count of $N \approx 6.5 \times 10^{7}$; larger minor counts push the required dominant count higher still. The value coincides with the 2,147,483,647-egret record verified by the China Ornithological Society, a count equal to $2^{31} - 1$, the largest signed 32-bit integer. A second screen flags 18 further county-months with the single-huge-count signature ($S \geq 20$, $H < 0.2$). "Corrupted removed" specifications exclude the union of the two screens, 19 county-months in total: (i) $S \geq 20$ with $H < 0.2$, and (ii) $S \geq 2$ with $H \leq 10^{-6}$ (script 07). Because Xinhui is also one of the counties absent from the policy source, it is excluded from the preferred specification by the policy-source restriction as well; the preferred estimates are numerically identical with and without screen (ii).

Recording-protocol simulations. To separate sampling accumulation from genuine community differences, we simulated observation of a fixed lognormal community (60 species, true Shannon $H = 3.19$) under three recording protocols. With counts recorded faithfully, observed H rises with effort and converges to the true value from below. With checklist-type recording (all counts set to one), observed H equals the log of the number of species detected; it therefore rises without bound as effort increases and exceeds the true value at high effort (4.04 at the highest effort simulated). Adding a single genuine large flock (20,000 individuals) depresses H to 0.02–0.86. A fixed community therefore suffices to generate the observed effort–diversity gradient, and the three protocols yield mutually incompatible "diversity" values from identical communities.

Species-level verification. Through a separate project we hold record-level data from the same platform: the platform's raw provincial delivery (2024 increment: 32 provinces, 8.1 million species-by-checklist rows, 282,197 checklists) and an event table built from an earlier multi-year delivery (2000–2025, 11.7 million rows). In the raw delivery the count column has zero missingness, and checklists in which every count is identical and equals the number of species recorded (the pattern expected if list length had leaked into the count column) make up at most 0.35% in any province, confirming genuine per-species abundance. Every count equals one in 43.7% of raw-delivery checklists and in 46.9% of 2014–2023 checklists in the event table, consistent with the 40.5% reported platform-wide by the Ornithological Society. Platform deliveries are cleaned with a ceiling on counts (maximum 10^5); the three astronomical records verified by the Society appear in none of them, yet one demonstrably enters the published outcome variable. County-month Shannon recomputed from the event table correlates with the published outcome at $r = 0.71$ across 17,404 overlapping county-months, indicating that the aggregated index is highly sensitive to which set of checklists is used.

Known limitations of the record-level checks. Three features of the platform's data-generating process, documented by long-standing platform users, are not addressed by either the original analysis or ours. First, observers who bird together each submit their own checklist: 7.1% of same-day, same-site groups involve two or more observers, and those groups account for 15.7% of all checklists. Second, a conspicuous flock is often recorded repeatedly: site-months with ten or more checklists account for 11.9% of checklists, and one site-month carries 119. Both forms of pseudo-replication can inflate the counts of abundant, conspicuous species when records are aggregated to county-months, and so can bias abundance-weighted indices in a systematic direction. Third, the platform has used two successive Chinese bird taxonomic checklists, involving splits and lumps that we did not reconcile. Two further screens we did not apply are the platform's own flagging of implausible records and the exclusion of multi-day, multi-county itineraries filed as a single record; the latter is a known source of error in our county-name parsing (92.3% matched). These limitations cut both ways. They affect the published dependent variable and our recomputation alike, which is why we read the correlation of $r = 0.71$ between the two as evidence that county-month aggregation is unstable, not as evidence that our values are the correct ones. The checklist-level statistics we quote are robust to pseudo-replication: retaining one checklist per same-day, same-site group changes the share of checklists in which every count equals one from 46.9% to 46.2%.

Specification curve. We estimated every combination of sample (all counties; policy-documented), cleaning (as published; the 19 corrupted county-months removed; ≥2 birders), outcome metric (Shannon; Simpson; log richness), and effort adjustment (authors' Duration only; explicit total-effort terms), yielding 36 specifications with identical fixed effects and clustering (Table S1; fig. S1). Poisson fixed-effects models for richness with a ln(BN) offset, following the catch-per-unit-effort convention, appear in Table S2.

Effort-proxy robustness. The authors' panel carries two total-effort measures, birdwatcher numbers (BN) and total hours (BT); the effort-adjusted specifications above use both. To test whether the null depends on this choice of proxy, we built a third effort proxy from the record-level platform data: the number of checklists (visits) per county-month, obtained by parsing checklist addresses into counties (address-parse coverage limits this to 17,737 county-months, 35.8% of the panel; unmatched rows are dropped). ln(visits) correlates with ln(BN) at $r = 0.78$ and with ln(BT) at $r = 0.56$, so it is related to the authors' measures without being redundant with them. Within the covered subsample, the authors' specification gives $\beta = -0.0084$ ($P = 0.17$) at baseline; adding ln(visits) gives $\beta = -0.0019$ ($P = 0.76$); restricting further to policy-documented counties gives $\beta = -0.0001$ ($P = 0.98$); and entering all three proxies together gives $\beta = +0.0060$ ($P = 0.30$). The nonsignificant result does not depend on which effort proxy is used; because the baseline in this subsample is itself nonsignificant, we read these as consistency checks across proxies rather than an independent test (script 12; output/visits_robustness.csv).

Spatial concentration of effort. Birdwatching effort concentrates where access is easy and economic activity is high. Within the assembled panel (46,528 county-months), regressing ln(BN) on ln(nightlight + 1) with year-month fixed effects and county-clustered standard errors gives $\beta = 0.163$ (SE = 0.019, $P < 10^{-15}$); population density, one of the authors' own controls, gives $\beta = 0.094$ (SE = 0.015, $P = 1 \times 10^{-10}$). Counties in the top decile of mean nightlight account for 36.0% of all recorded birdwatchers, and the bottom half for 25.8%. Recorded effort therefore tracks the urbanization–

development gradient, so county-month records reflect where observers are, not only where birds are; this cross-sectional concentration is the spatial counterpart of the temporal growth in fig. S2D (script 13; output/effort_accessibility.csv).

Power and minimum detectable effects. For the preferred corrected specification (corrupted records removed, policy-documented counties, explicit effort terms), the minimum detectable effect at 80% power and $\alpha = 0.05$ is $2.80 \times$ SE, i.e., 1.57% (Shannon), 1.20% (Simpson), and 4.04% (log richness) per 1 s.d. of PSI. The power to detect the published −2.10%, were it real, is $\Phi(0.0125/\mathrm{SE} - 1.96) = 96\%$.

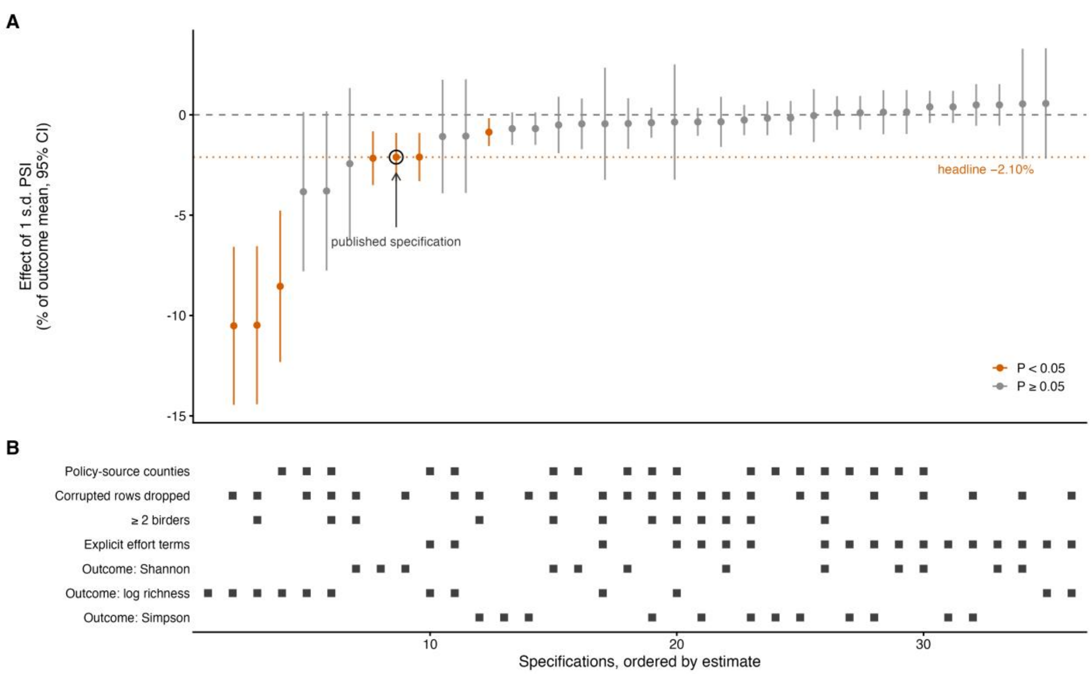

Fig. S1. Specification curve across 36 defensible analysis paths. Estimates of the PSI effect (% of outcome mean per 1 s.d., 95% CIs) for the full cross of sample, cleaning, outcome metric, and effort adjustment, ordered by estimate; the lower panel indicates each path's decisions. All seven significant negative estimates (orange) arise without effort terms; all 18 effort-adjusted paths are null. The open circle marks the published specification.

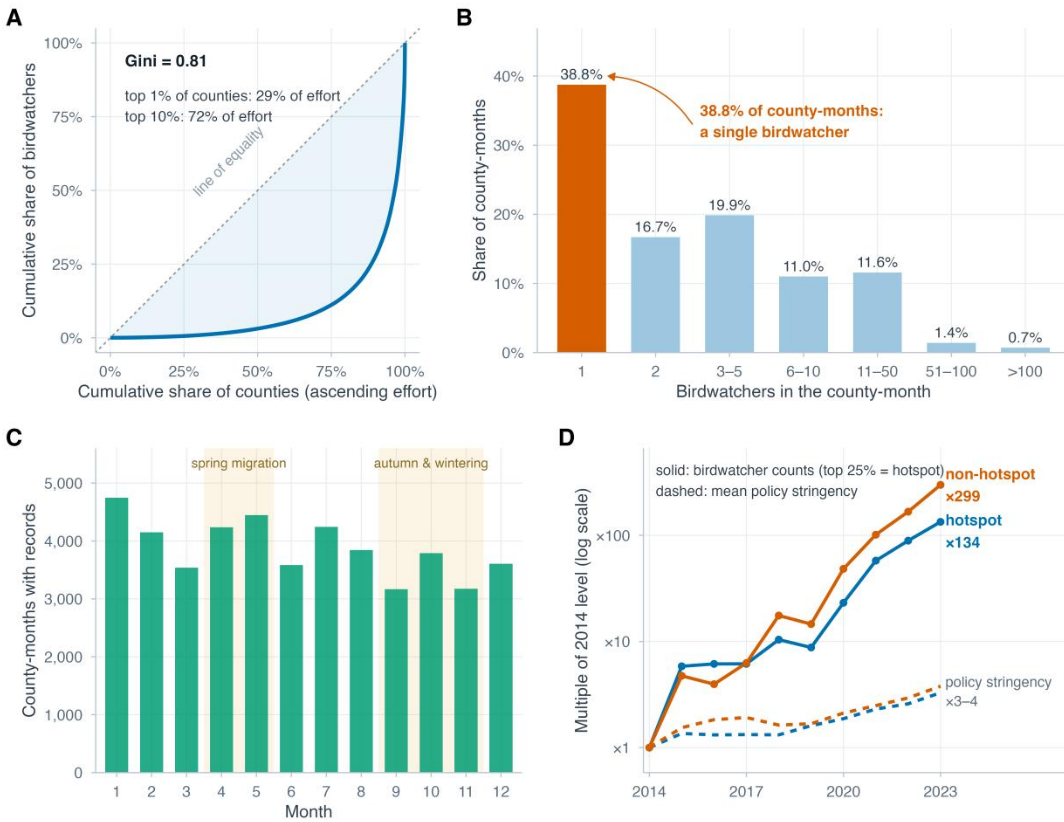

Fig. S2. Structure of the birdwatching data underlying the outcome variable. (**A**) Observation effort is highly concentrated among counties (Gini = 0.81; 1% of counties contribute 29% of effort). (**B**) 38.8% of county-months derive from a single birdwatcher. (**C**) Reporting is strongly seasonal, peaking in wintering and migration periods. (**D**) Effort growth (solid lines) outpaces policy growth (dashed) by two orders of magnitude, and non-hotspot counties (×299) outgrow hotspot counties (×134): the mismatch in timing between these two diffusion processes underlies the within-county association between policy and effort.

Table S1. Specification curve: all 36 combinations of sample, cleaning, outcome metric, and effort adjustment. All models include county and year-month fixed effects with county-clustered standard errors. Effects are % of the outcome mean per 1 s.d. of PSI; rows ordered by estimate. Rows in bold denote statistically significant negative estimates (all arise without effort terms).

#	Sample	Cleaning	Metric	Effort	Effect (%/s.d.)	95% CI	*P*	*n*
1	All counties	As published	log richness	Duration only	**-10.51**	[-14.45, -6.57]	1.8e-07	45,606
2	All counties	Corrupted removed	log richness	Duration only	**-10.48**	[-14.42, -6.54]	2e-07	45,587
3	All counties	≥2 birders	log richness	Duration only	**-8.54**	[-12.32, -4.77]	9.7e-06	27,571
4	Policy-documented	As published	log richness	Duration only	-3.83	[-7.80, +0.14]	0.059	25,245
5	Policy-documented	Corrupted removed	log richness	Duration only	-3.79	[-7.76, +0.17]	0.061	25,234
6	Policy-documented	≥2 birders	log richness	Duration only	-2.43	[-6.20, +1.33]	0.21	14,032
7	All counties	≥2 birders	Shannon	Duration only	**-2.16**	[-3.50, -0.82]	0.0016	28,052
8	All counties	As published	Shannon	Duration only	**-2.11**	[-3.31, -0.90]	0.00062	46,371
9	All counties	Corrupted removed	Shannon	Duration only	**-2.10**	[-3.31, -0.90]	0.00063	46,352
10	Policy-documented	As published	log richness	Explicit effort	-1.08	[-3.91, +1.75]	0.45	25,245
11	Policy-documented	Corrupted removed	log richness	Explicit effort	-1.06	[-3.89, +1.77]	0.46	25,234
12	All counties	≥2 birders	Simpson	Duration only	**-0.86**	[-1.55, -0.17]	0.015	28,052
13	All	As	Simpson	Duratio	-0.69	[-1.50,	0.096	46,371

#	Sample	Cleaning	Metric	Effort	Effect (%/s.d.)	95% CI	*P*	*n*
	counties	published	n	n only		+0.12]		
14	All counties	Corrupted removed	Simpson	Duration only	-0.69	[-1.50, +0.12]	0.096	46,352
15	Policy-documented	≥2 birders	Shannon	Duration only	-0.51	[-1.91, +0.90]	0.48	14,510
16	Policy-documented	As published	Shannon	Duration only	-0.45	[-1.71, +0.82]	0.49	26,007
17	All counties	≥2 birders	log richness	Explicit effort	-0.45	[-3.24, +2.35]	0.75	27,571
18	Policy-documented	Corrupted removed	Shannon	Duration only	-0.43	[-1.70, +0.83]	0.5	25,996
19	Policy-documented	≥2 birders	Simpson	Duration only	-0.39	[-1.14, +0.36]	0.3	14,510
20	Policy-documented	≥2 birders	log richness	Explicit effort	-0.36	[-3.23, +2.51]	0.81	14,032
21	All counties	≥2 birders	Simpson	Explicit effort	-0.35	[-1.05, +0.34]	0.32	28,052
22	All counties	≥2 birders	Shannon	Explicit effort	-0.35	[-1.60, +0.90]	0.58	28,052
23	Policy-documented	≥2 birders	Simpson	Explicit effort	-0.26	[-1.01, +0.50]	0.5	14,510
24	Policy-documented	As published	Simpson	Duration only	-0.17	[-1.02, +0.68]	0.7	26,007
25	Policy-documented	Corrupted removed	Simpson	Duration only	-0.16	[-1.00, +0.69]	0.72	25,996

#	Sample	Cleaning	Metric	Effort	Effect (%/s.d.)	95% CI	*P*	*n*
26	Policy-documented	≥2 birders	Shannon	Explicit effort	-0.04	[-1.36, +1.28]	0.95	14,510
27	Policy-documented	As published	Simpson	Explicit effort	+0.09	[-0.75, +0.93]	0.83	26,007
28	Policy-documented	Corrupted removed	Simpson	Explicit effort	+0.10	[-0.74, +0.94]	0.81	25,996
29	Policy-documented	As published	Shannon	Explicit effort	+0.13	[-0.97, +1.23]	0.81	26,007
30	Policy-documented	Corrupted removed	Shannon	Explicit effort	+0.14	[-0.96, +1.24]	0.8	25,996
31	All counties	As published	Simpson	Explicit effort	+0.39	[-0.41, +1.20]	0.34	46,371
32	All counties	Corrupted removed	Simpson	Explicit effort	+0.40	[-0.40, +1.20]	0.33	46,352
33	All counties	As published	Shannon	Explicit effort	+0.49	[-0.54, +1.53]	0.35	46,371
34	All counties	Corrupted removed	Shannon	Explicit effort	+0.49	[-0.54, +1.53]	0.35	46,352
35	All counties	As published	log richness	Explicit effort	+0.54	[-2.21, +3.29]	0.7	45,606
36	All counties	Corrupted removed	log richness	Explicit effort	+0.57	[-2.18, +3.32]	0.69	45,587

Robustness variants of the preferred specification (corrupted removed, policy-documented, explicit effort):

Variant	Metric	Effect (%/s.d.)	95% CI	*P*	*n*
Province-clustered SEs	Shannon	+0.14	[-1.40, +1.68]	0.86	25,996
Province-clustered SEs	log richness	-1.06	[-5.96, +3.83]	0.67	25,234
Period 2014–2019	Shannon	-3.39	[-7.38, +0.60]	0.096	11,405
Period 2020–2023	Shannon	+0.03	[-1.12, +1.18]	0.96	34,496
Period 2020–2023	log richness	+0.36	[-2.52, +3.23]	0.81	33,904

350 **Table S2. Diagnostic and negative-control regressions.** All models include county and
351 year-month fixed effects; standard errors clustered by county.

Regression (county and year-month FE, county-clustered)	β(PSI)	*P*	*n*	Note
ln(birdwatchers) on PSI, full sample	-0.0410	2.1e-11	46,371	
ln(birdwatchers) on PSI, policy-documented counties	-0.0078	0.166	26,007	81% attenuation
ln(birdwatchers) on next calendar-year PSI, conditional on current PSI	-0.0510	1.71e-06	30,426	pre-trend
Realized PV area on PSI, capacity imputed as in the authors' code	-0.0175	0.268	46,371	54.6% imputed zeros
Realized PV	-0.0541	0.0949	21,054	not positive

Regression (county and year-month FE, county-clustered)	β(PSI)	*P*	*n*	Note
area on PSI, county-years with recorded capacity				either
ln(nightlight) on PSI	-0.0017	0.237	46,371	clean negative control
Shannon on PSI, county-specific linear trends added	-0.0079	0.102	46,371	headline non-significant
Shannon on PSI + ln(birdwatchers)	+0.0032	0.31	46,371	effect null under effort control
Richness (Poisson FE): Poisson FE, no effort (authors' logic)	-0.0233	6.29e-08	45,606	-9.2%/s.d.
Richness (Poisson FE): + offset log(birdwatchers) [CPUE]	+0.0151	0.00568	45,606	+6.4%/s.d.
Richness (Poisson FE): offset + policy-documented counties	+0.0055	0.318	25,245	+2.3%/s.d.
Richness (Poisson FE): offset + ln(hours) + policy-documented	+0.0029	0.553	25,245	+1.2%/s.d.