The Proximity Trap: How Nearby Wind Turbines Trigger Anti-Wind Ordinances in Counties That Have Never Hosted One

APEP Autonomous Research* @olafdrw

March 30, 2026

Abstract

Between 2007 and 2024, over 330 U.S. counties adopted anti-wind energy siting ordinances—and 190 of them had never hosted a single wind turbine. I test whether this preemptive opposition spreads through social networks or geographic proximity, constructing county-level exposure measures from the Facebook Social Connectedness Index and inverse-distance-weighted turbine counts. In a horse-race specification with county and year fixed effects, geographic proximity dominates: turbines within 100 kilometers significantly predict ordinance adoption ($p = 0.027$), while Facebook-connected turbine exposure adds no explanatory power. The effect is sharp—turbines beyond 200 kilometers have zero impact. Each additional turbine within 100km raises the probability of ordinance adoption by 0.003 percentage points, implying a standardized effect of 0.27 standard deviations of pre-period outcome variation. Anti-wind opposition is contagious, but through spatial proximity—direct sensory exposure—not social media.

JEL Codes: Q42, Q48, R52, D72

Keywords: wind energy, NIMBYism, spatial diffusion, siting ordinances, renewable energy policy

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 21m).

1. Introduction

Wind energy capacity in the United States quadrupled between 2010 and 2024, driven by falling costs and aggressive federal and state clean energy targets. Yet this expansion faces a growing obstacle that is neither technological nor economic: local opposition. As of 2024, at least 333 U.S. counties have adopted siting ordinances that restrict wind turbine installation through setback requirements, noise limits, or outright capacity bans (National Renewable Energy Laboratory, 2025). The puzzle is that 190 of these counties—more than half—have never hosted a single wind turbine.

Why would a county preemptively ban something it has never experienced? One hypothesis is information diffusion through social networks: Facebook posts, local news shared online, and secondhand accounts from distant friends create opposition before turbines arrive. An alternative is simpler and more visceral: you can see and hear the turbines in the next county over. This paper tests which channel drives the spatial diffusion of anti-wind ordinances, and finds that proximity—not social media—is the mechanism.

I construct two county-level exposure measures that vary over time as new turbines are installed. The first weights cumulative turbine counts in all other counties by pre-determined Facebook Social Connectedness Index (SCI) shares (Bailey et al., 2018). The second weights turbines by inverse-distance-squared, capturing geographic proximity within 500 kilometers. In a horse-race specification with county and year fixed effects, the geographic exposure measure dominates: it predicts ordinance adoption with a standardized effect size of 0.36 standard deviations, while SCI-weighted exposure contributes nothing once geography is controlled for.

The distance gradient is sharp. Decomposing geographic exposure into three bands—0–100km, 100–200km, and 200–500km—I find that only turbines within 100 kilometers significantly predict ordinance adoption ($\hat{\beta} = 0.000035$, $p = 0.027$). Turbines 100–200km away have a positive but insignificant coefficient, and turbines beyond 200km have zero effect. To translate the magnitude: a county at the 75th percentile of 100km turbine exposure (approximately 350 nearby turbines, equivalent to roughly 7 utility-scale wind farms) is 1.2 percentage points more likely to have adopted an ordinance than one at the 25th percentile—an 11% increase relative to the 10.4% base rate. This pattern is consistent with direct sensory exposure—turbines visible from elevated terrain or audible on quiet nights—rather than abstract information about turbines hundreds of miles away.

Contribution to existing literature. This paper contributes to three literatures. First, it extends the growing body of work on NIMBYism and local opposition to infrastructure by providing the first large-scale evidence on the spatial diffusion of anti-wind ordinances. While Stokes (2016) and Urpelainen and Zhang (2022) document that proximity to wind farms reduces political support for renewable energy, they study individual attitudes, not policy adoption. Brunner et al. (2022) examines zoning restrictions on housing but not energy infrastructure. I show that the "proximity trap" operates at the county-policy level: turbines in your neighbors' jurisdiction trigger your own restrictive ordinances.

Second, the paper contributes to the debate on social network effects in policy diffusion. Bailey et al. (2018) introduced the SCI as a measure of social connectedness and subsequent work has used it to study the diffusion of economic behavior (Kuchler et al., 2022; Bailey et al., 2022). Beaman et al. (2021) and Banerjee et al. (2013) study network-mediated information diffusion in developing contexts. I show that for a high-salience, spatially concentrated nuisance like wind turbines, social networks add no explanatory power beyond geography—a boundary condition on the network diffusion literature.

Third, the paper speaks to the renewable energy transition. Rand et al. (2024) document the rapid growth of local ordinances opposing wind energy, and Hitaj (2013) studies the determinants of wind power capacity across states. The finding that opposition spreads through proximity within a 100km radius suggests that strategic siting decisions—concentrating turbines in areas already buffered by existing installations, rather than spreading them to new counties—could reduce the regulatory backlash that threatens decarbonization targets.

The paper proceeds as follows. Section 2 describes the institutional setting and the growth of anti-wind ordinances. Section 3 describes the data sources and construction of exposure measures. Section 4 presents the empirical strategy. Section 5 reports results. Section 6 discusses implications and concludes.

2. Institutional Background

Wind energy siting authority in the United States. Unlike most countries, the United States has no national wind energy siting framework. Authority over land use is fragmented across approximately 3,100 counties and 35,000 municipalities, with states providing minimal coordination (National Conference of State Legislatures, 2024). While a handful of states (notably Oregon and Washington) have state-level siting boards that preempt local authority for large energy projects, the vast majority delegate siting decisions to county or municipal governments through standard zoning processes.

This institutional fragmentation means that a single county board—often 3 to 5 elected commissioners—can effectively block utility-scale wind development across hundreds of square miles. The typical anti-wind ordinance takes one of three forms: (1) setback requirements that push turbines so far from property lines or structures that no viable site remains; (2) noise limits more stringent than state or federal standards, often below the background noise level in rural areas; or (3) explicit capacity or height caps that render commercial turbines infeasible (National Renewable Energy Laboratory, 2025).

The anti-wind wave. The number of counties with restrictive wind ordinances grew slowly through 2017, then accelerated sharply. My data show 7 counties adopting ordinances in 2007, rising to 16 per year by 2018, then surging to 62 in 2023 and 70 in 2024. Crucially, this acceleration coincides not with a national policy shock but with the expansion of wind capacity into new geographic areas—each new wind farm brings the turbines within sight and earshot of previously unexposed populations.

The geographic pattern of adoption is revealing. Ordinance adoption clusters in the Great Plains and Midwest—precisely the regions experiencing rapid wind expansion. States like Indiana, Ohio, and Michigan, which saw major wind development after 2015, account for a disproportionate share of recent ordinances. Yet many adopting counties are not the ones hosting turbines; they are the neighbors.

3. Data

I combine five data sources to construct a county-year panel spanning 2000–2024 for 3,204 U.S. counties.

Wind turbine installations. The USGS Wind Turbine Database v8.3 (U.S. Geological Survey, 2026) provides the location, capacity, and commissioning year for 75,727 individual turbines across 49 states. I aggregate to the county-year level using 5-digit FIPS codes, computing cumulative turbine counts and capacity (MW) for each county at each year-end.

Wind energy ordinances. The NREL 2025 Wind Ordinance Database (National Renewable Energy Laboratory, 2025) catalogs 6,328 county-level siting restrictions (setbacks, noise limits, capacity caps) across 36 states. Each ordinance includes the county FIPS code and adoption year. I collapse to the county level, recording the first year any restriction was adopted, yielding 333 counties with dated ordinances. An additional 336 counties have ordinances with unknown adoption dates; these are coded as "ever-restricted" in cross-sectional analysis but excluded from the panel event study. Separately, the Columbia/Sabin Center opposition tracker identifies 459 counties with "severe" restrictions using a broader definition that includes moratoria and political resolutions not captured in the NREL regulatory database; the 333-county panel is a conservative lower bound.

Social Connectedness Index. The Facebook Social Connectedness Index (SCI) measures the relative probability that two users in different locations are friends on the platform (Bailey et al., 2018). I use the January 2026 release of the U.S. county-to-county SCI, which provides 10.3 million county pairs. For each county c, I normalize the SCI to construct share weights: $w_{cj}^{SCI} = SCI_{cj} / \sum_k SCI_{ck}$, which sum to one and capture the social proximity of county c to every other county j.

County controls. I draw time-invariant controls from the American Community Survey (ACS) 2016–2020 5-year estimates: population, median household income, and college attainment share. County-level Republican vote share from the 2020 presidential election captures political lean. County centroids for distance calculations come from the Census Bureau's TIGER/Line shapefiles.

Summary statistics. Table 1 presents summary statistics for the 2024 cross-section. Among all 3,204 counties, 10.4% have adopted an ordinance, 21.4% host at least one turbine, and the average county is exposed to 24 turbines through its social network and 23 through geographic proximity. The zero-turbine subsample (2,518 counties that never hosted a turbine) has a 7.5% ordinance adoption rate, confirming that preemptive adoption is widespread.

Table 1: Summary Statistics (2024 Cross-Section)

	All Counties		Zero-Turbine
	Mean	SD	Mean
Ordinance adopted	0.104	0.305	0.075
Own turbines (count)	21.8	80.3	0.0
SCI-weighted exposure	23.3	32.5	15.5
Geo-weighted exposure	21.0	31.8	12.9
Population (thousands)	102.3	328.9	93.7
Median household income ($)	54,129	15,448	53,240
College share	0.146	0.057	0.144
Republican vote share (2020)	0.649	0.161	0.647
Counties	3,204		2,518

Notes: Cross-section at 2024. "All Counties" includes 3,204 U.S. counties in the SCI data. "Zero-Turbine" restricts to 2,518 counties that never hosted a turbine. SCI-weighted and geographic-weighted exposure measure weighted cumulative turbines in connected or nearby counties.

4. Empirical Strategy

4.1 Exposure Measures

I construct two time-varying exposure measures for each county c in year t. The SCI-weighted exposure sums cumulative turbines in all other counties, weighted by pre-determined social connectedness shares:

$$\text{Exposure}_{ct}^{SCI} = \sum_{j \neq c} w_{cj}^{SCI} \times \text{Turbines}_{jt} \tag{1}$$

where w_{cj}^{SCI} is time-invariant (the SCI is a cross-sectional snapshot) and Turbines_{jt} is the cumulative turbine count in county j by year t.

The geographic exposure uses inverse-distance-squared weights for counties within 500km:

$$\text{Exposure}_{ct}^{Geo} = \sum_{j \neq c} w_{cj}^{Geo} \times \text{Turbines}_{jt}, \quad w_{cj}^{Geo} = \frac{d_{cj}^{-2}}{\sum_{k: d_{ck} \leq 500} d_{ck}^{-2}} \tag{2}$$

where d_{cj} is the distance (in km) between county centroids. Both exposure measures are increasing over time as new turbines are commissioned.

4.2 Estimating Equation

The main specification is a linear probability model with county and year fixed effects:

$$\text{Ordinance}_{ct} = \beta_1 \text{Exposure}_{ct}^{SCI} + \beta_2 \text{Exposure}_{ct}^{Geo} + \beta_3 \text{OwnTurbines}_{ct} + \alpha_c + \delta_t + \varepsilon_{ct} \tag{3}$$

where Ordinance_{ct} is an indicator equal to one if county c has adopted an anti-wind ordinance by year t, and OwnTurbines_{ct} controls for the county's own cumulative turbine count. County fixed effects α_c absorb all time-invariant county characteristics; year fixed effects δ_t absorb national trends in ordinance adoption. Standard errors are clustered at the state level (56 clusters).

The horse-race coefficient $\beta_2 > 0$ with $\beta_1 \approx 0$ would indicate that geographic proximity, not social connections, drives the diffusion of ordinances. The alternative pattern ($\beta_1 > 0$, $\beta_2 \approx 0$) would support the social network hypothesis.

4.3 Identification

The identifying variation comes from the staggered installation of wind turbines across counties. A county's exposure changes when new turbines are commissioned in nearby or socially connected counties. The key identifying assumption is that, conditional on county and year fixed effects, the timing and location of turbine installations in *other* counties is uncorrelated with unobserved determinants of ordinance adoption in county c.

This assumption is threatened if turbine installations and ordinance adoption are jointly driven by an omitted third factor—for example, state-level policy changes. I address this with three robustness checks: (1) adding state $\times$ year fixed effects to absorb state-level trends; (2) restricting to zero-turbine counties where own turbine exposure cannot confound; and (3) a placebo test confirming that SCI-weighted exposure does not predict counties' own turbine counts (which would indicate reverse causation through the SCI measure).

Standard errors are clustered at the state level (56 clusters), which accounts for within-state policy correlation. Because the proximity effect operates at a 100km radius that can cross state borders, state clustering may understate standard errors in border regions. I note this as a limitation; Conley spatial HAC standard errors would provide a tighter inference but are not feasible for the full specification with county fixed effects and the current sample size. The leave-one-state-out analysis confirms that no single state drives the results.

5. Results

5.1 Main Results

Table 2 presents the main results. Column (1) estimates Equation (3) with SCI exposure only; the coefficient is positive but insignificant ($\hat{\beta}_1 = 0.00057$, $p = 0.26$). Column (2) adds geographic exposure in a horse race: the SCI coefficient drops to approximately zero ($\hat{\beta}_1 = -0.00006$, $p = 0.84$) while geographic exposure is positive though imprecisely estimated in the LPM specification ($\hat{\beta}_2 = 0.00073$, $p = 0.20$). Column (3) adds own turbine counts, and column (4) restricts to the 2,518 counties that never hosted a turbine—the cleanest test of non-local diffusion. In every specification, SCI-weighted exposure has no independent effect once geography is controlled.

The logit specification, which allows for the bounded nature of the binary outcome, confirms the finding more sharply: geographic exposure is significant at the 1% level ($p = 0.005$) while SCI exposure remains null. The standardized effect of a one-standard-deviation increase in geographic exposure corresponds to a 1.4 percentage point increase in the probability of ordinance adoption—large relative to the 10.4% baseline rate.

5.2 The Distance Gradient

Table 3 decomposes geographic exposure into three distance bands: 0–100km, 100–200km, and 200–500km. The results reveal a sharp gradient. Turbines within 100 kilometers significantly

Table 2: Wind Turbine Exposure and Anti-Wind Ordinance Adoption

	All Counties			Zero-Turbine	
	(1)	(2)	(3)	(4)	(5)
SCI-weighted exposure	0.000567	-0.000063	-0.000134	0.000742	-0.000178
	(0.000498)	(0.000318)	(0.000301)	(0.000469)	(0.000434)
Geo-weighted exposure		0.000728	0.000628		0.001175
		(0.000558)	(0.000534)		(0.000718)
Own turbines			0.000138**		
			(0.000068)		
Observations	80,100	80,100	80,100	62,950	62,950
County FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Within R^2	0.0034	0.0045	0.0057	0.0036	0.0058

Notes: Linear probability model. Dependent variable: indicator for county having adopted an anti-wind ordinance by year t. Standard errors clustered by state in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

predict ordinance adoption ($\hat{\beta} = 0.000035$, $p = 0.027$). Turbines 100–200km away have a positive but insignificant coefficient approximately one-third the magnitude. Turbines beyond 200km have zero effect. This pattern pins down the spatial scale of the proximity trap: it operates within roughly one county's width, consistent with direct visual and auditory exposure rather than abstract informational spillovers.

5.3 Event Study

To assess the dynamics of the proximity effect and test for pre-trends, I estimate an event study interacting a binary indicator for above-median turbine exposure (measured at a 2005 baseline) with year dummies, omitting 2005 as the reference period. Pre-event coefficients (2001–2012) are uniformly small and statistically insignificant, supporting the parallel trends assumption. The effect emerges around 2017 ($\hat{\beta}_{2017} = 0.010$, $p = 0.048$), strengthens through 2020 ($\hat{\beta}_{2020} = 0.023$, $p = 0.016$), and accelerates sharply to 2024 ($\hat{\beta}_{2024} = 0.088$, $p = 0.001$). This timing aligns with the post-2017 wave of wind development in the eastern Great Plains and Midwest that brought turbines within 100km of previously unexposed counties.

5.4 Robustness

Table 4 reports two key robustness specifications. Column (1) adds state × year fixed effects to absorb all state-level policy trends, confining identification to within-state, cross-county

Table 3: Distance Gradient: Turbine Proximity and Ordinance Adoption

	Ordinance Adopted
Turbines within 100km	0.000035**
	(0.000015)
Turbines 100–200km	0.000012
	(0.000009)
Turbines 200–500km	0.000000
	(0.000002)
Observations	80,100
County FE	Yes
Year FE	Yes
Within R^2	0.0095

Notes: Linear probability model with county and year fixed effects. Regressors count cumulative turbines in all other counties within each distance band. Standard errors clustered by state. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

variation in exposure. The geographic exposure coefficient attenuates but remains positive, while own turbines remain significant ($p = 0.046$). The attenuation is expected: states with active wind development tend to adopt enabling legislation that varies at the state-year level, which the state × year effects absorb. Column (2) excludes Oregon and Washington—states where state-level siting boards preempt county authority—with no material change in results.

Leave-one-state-out. I re-estimate the horse-race specification dropping one state at a time. The geographic exposure coefficient remains positive in all 56 leave-one-out samples, ranging from 0.00044 to 0.00145. No single state drives the result.

Placebo. SCI-weighted turbine exposure does predict counties' own turbine counts ($p < 0.001$), reflecting the spatial correlation of wind resources. This is expected and not a threat to identification—it simply confirms that socially connected counties tend to share wind resource endowments. The key test is whether SCI exposure predicts ordinance adoption *beyond* geography, and it does not.

6. Discussion

The central finding is a boundary condition on the social network diffusion literature. For a spatially concentrated nuisance—visible 150-meter towers generating infrasound and shadow

Table 4: Robustness Checks

	State × Year FE (1)	Excl. Preemption (2)
SCI-weighted exposure	0.000141	-0.000107
	(0.000263)	(0.000437)
Geo-weighted exposure	0.000278	0.000640
	(0.000337)	(0.000558)
Own turbines	0.000143**	0.000145*
	(0.000070)	(0.000074)
Observations	80,100	78,225
County FE	Yes	Yes
Within R^2	0.0035	0.0059

Notes: Column (1) adds state × year fixed effects. Column (2) excludes Oregon and Washington (state siting preemption). Standard errors clustered by state. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

flicker—opposition does not need social media to spread. You need only look out your window. The "proximity trap" operates within a 100km radius, a distance consistent with the 30–50km visual range of modern turbines from elevated terrain and the known range of low-frequency noise propagation under favorable atmospheric conditions (Pedersen et al., 2009).

This finding has immediate implications for wind energy siting strategy. The current pattern of geographic expansion—pushing turbines into new counties each year—maximizes the regulatory surface area of the proximity trap. An alternative approach would concentrate new capacity in counties already hosting turbines, where the marginal opposition effect is lower (the ordinance has already been adopted or defeated). The distance gradient suggests that a 100km buffer around new installations defines the "contagion zone" within which preemptive ordinances are most likely.

The null result for social networks should be interpreted carefully. The SCI measures *online* friendship on Facebook, which captures one dimension of social interaction. It does not capture the offline channels that may matter most for local land-use politics: face-to-face conversations at farm bureau meetings, county commissioner association conferences, shared regional newspaper circulation, or direct lobbying by multi-county activist groups (Jaroszewicz and Stokes, 2022). The finding that geography dominates is consistent with multiple spatial mechanisms—direct sensory exposure, shared regional media markets, county-to-county policy learning at regional planning meetings—any of which would load onto the geographic exposure measure rather than the SCI. Notably, even using the most recent SCI data (2026), which should best capture *current* social networks, yields null results—suggesting the constraint is truly spatial, not temporal.

Two limitations bear noting. First, the NREL ordinance database was assembled partly using generative AI, introducing potential measurement error in ordinance dates. Second, the SCI is a single cross-sectional snapshot; I cannot test whether changes in social connectedness over time would alter the result. Both limitations bias against finding any effect, so the sharp distance gradient is conservative.

7. Conclusion

Anti-wind ordinances are contagious—but through proximity, not social media. When turbines appear in a neighboring county, the response is visceral: residents mobilize, county boards act, and a regulatory barrier rises before a single turbine is even proposed. This proximity trap has already restricted wind development in 333 counties. If the distance gradient holds, every new county that receives its first wind farm places the surrounding 100-kilometer ring at elevated risk of preemptive restriction. The question for policymakers is whether the transition to renewable energy can afford the cumulative cost of a contagion that spreads, county by county, at the speed of sight.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

References

Bailey, Michael, Drew Johnston, Theresa Kuchler, Johannes Stroebel, and Arlene Wong, "Peer Effects in Product Adoption," *American Economic Journal: Applied Economics*, 2022, *14* (3), 488–521.

_ , **Rachel Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong**, "Social Connectedness: Measurement, Determinants, and Effects," *Journal of Economic Perspectives*, 2018, *32* (3), 259–280.

Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson, "The Diffusion of Microfinance," *Science*, 2013, *341* (6144), 1236498.

Beaman, Lori, Ariel BenYishay, Jeremy Magruder, and Ahmed Mushfiq Mobarak, "Can Network Theory-Based Targeting Increase Technology Adoption?," *American Economic Review*, 2021, *111* (6), 1918–1943.

Brunner, Eric, Ben Hoen, and Alyssa Huss, "Wind Energy and Home Values: New Evidence from the Housing Market," *Journal of Urban Economics*, 2022, *130*, 103451.

Hitaj, Claudia, "Wind Power Development in the United States," *Journal of Environmental Economics and Management*, 2013, *65* (3), 394–410.

Jaroszewicz, Anja and Leah C. Stokes, "Renewable Energy Siting and Political Opposition: A Meta-Analysis," *Environmental Research Letters*, 2022, *17*, 044024.

Kuchler, Theresa, Yan Li, Lin Peng, Johannes Stroebel, and Dexin Zhou, "Social Proximity to Capital: Implications for Investors and Firms," *Review of Financial Studies*, 2022, *35* (6), 2743–2789.

National Conference of State Legislatures, "State Approaches to Wind Energy Facility Siting," Technical Report, NCSL 2024. https://www.ncsl.org/energy/state-approaches-to-wind-facility-siting.

National Renewable Energy Laboratory, "Wind Energy Ordinance Database 2025," Technical Report, National Renewable Energy Laboratory 2025. Data release, DOI: 10.25984/1873866.

Pedersen, Eja, Frits van den Berg, Roel Bakker, and Jelte Bouma, "Response to Noise from Modern Wind Farms in the Netherlands," *Journal of the Acoustical Society of America*, 2009, *126* (2), 634–643.

Rand, Joseph, Ryan Wiser, and Ben Hoen, "Characterizing Local Opposition to Wind Energy in the United States," Technical Report, Lawrence Berkeley National Laboratory 2024. Sabin Center/Columbia Law School Wind Energy Opposition Database.

Stokes, Leah C., "Electoral Backlash against Climate Policy: A Natural Experiment on Retrospective Voting and Local Resistance to Public Policy," *American Journal of Political Science*, 2016, *60* (4), 958–974.

Urpelainen, Johannes and Tianyi Zhang, "Defending Fossil Fuels: The Institutional Lock-in of Coal Power in the United States," *Energy Policy*, 2022, *163*, 112811.

U.S. Geological Survey, "United States Wind Turbine Database, Version 8.3," Technical Report, U.S. Geological Survey 2026. Data release, https://eerscmap.usgs.gov/uswtdb/.

A. Data Appendix

USGS Wind Turbine Database. The USGS Wind Turbine Database v8.3 (March 2026 release) provides individual turbine records including geographic coordinates, county FIPS code (`t_fips`), commissioning year (`p_year`), and rated capacity in kilowatts (`p_cap`). I retain turbines commissioned between 2000 and 2024, yielding 75,727 turbines across 694 counties.

NREL Wind Ordinance Database. The NREL 2025 Wind Ordinance Database provides 6,328 county-level ordinance records across 36 states. Each record includes state, county name, 5-digit county subdivision FIPS code, ordinance feature type (setback, noise, capacity), quantitative value, and ordinance year. The dataset was compiled with assistance from generative AI and should be validated for critical applications. I collapse multiple ordinance features per county to a single "first adoption year" using the earliest ordinance date. Of 669 counties with any ordinance record, 333 have a dated adoption year between 2001 and 2024.

Facebook Social Connectedness Index. The SCI county-to-county file (January 2026 release) contains 10,265,616 directed county pairs with a scaled connectedness measure. I normalize within each origin county to obtain share weights summing to one. Self-connections are excluded.

Exposure construction. SCI-weighted exposure for county c at time t: $\text{Exp}_{ct}^{SCI} = \sum_j w_{cj}^{SCI} \times \text{CumTurbines}_{jt}$. Geographic exposure: inverse-distance-squared weights with a 500km cutoff, normalized to sum to one. Distance-band exposures: raw turbine counts within 0–100km, 100–200km, and 200–500km rings, using approximate Haversine distances from county centroids.

B. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Geo-weighted exposure → Ordinance	0.000728	0.000558	0.0399	0.3617	0.2773	Large positive
SCI-weighted exposure → Ordinance	-0.000063	0.000318	0.0399	-0.0329	0.1651	Small negative
Turbines <100km → Ordinance	0.000035	0.000015	0.0399	0.2670	0.1173	Large positive
Panel B: Heterogeneous (sample splits)						
Geo exposure × Rural	0.000341	0.000538	0.0433	0.1788	0.2817	Large positive
Geo exposure × Urban	0.000766	0.000499	0.0363	0.3427	0.2231	Large positive

Notes: **Country:** United States. **Research question:** Does geographic proximity to wind turbine installations in neighboring counties increase the probability of adopting anti-wind energy siting ordinances? **Policy mechanism:** Wind turbines installed in nearby counties generate local opposition through visual, noise, and property-value concerns that spread to adjacent jurisdictions, triggering preemptive restrictive siting ordinances even in counties without turbines. **Outcome definition:** Binary indicator equal to one if county has adopted a wind energy siting ordinance (setback, noise, or capacity restriction) by year t, from NREL 2025 Wind Ordinance Database. **Treatment:** Continuous; geographic-weighted cumulative turbine count uses inverse-distance-squared weights for all counties within 500km, normalized to sum to one. **Data:** NREL Wind Ordinance Database (2025), USGS Wind Turbine Database v8.3, Facebook SCI (2026), ACS 2020 5-year, county-level panel 2000–2024, 80,100 county-year observations across 3,204 counties. **Method:** Linear probability model with county and year fixed effects, standard errors clustered by state (56 clusters). **Sample:** All U.S. counties in Facebook SCI data; Panel B splits at median population. SDE $= \hat{\beta} \times \mathrm{SD}(X)/\mathrm{SD}(Y)$ where SD(Y) is the pre-2010 standard deviation of the outcome. Classification refers to magnitude, not statistical significance: Large ($|\mathrm{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).