Electoral Backlash against Climate Policy:
A Natural Experiment on Retrospective Voting and Local Resistance to Public Policy

by

Leah C. Stokes

Master of Public Administration
Columbia University
New York, New York (2009)

Bachelor of Science
University of Toronto
Toronto, Ontario, Canada (2007)

Submitted to the Department of Political Science
in partial fulfillment of the requirements for the degree of

Master of Science in Political Science
at the
Massachusetts Institute of Technology

June 2015

Author Signature redacted
Department of Political Science
April 27, 2015

Certified by Signature redacted
Christopher Warshaw
Assistant Professor of Political Science
Thesis Supervisor

Accepted by Signature redacted
Andrea Campbell
Professor of Political Science
Chair, Graduate Program Committee

Electoral Backlash against Climate Policy:
A Natural Experiment on Retrospective Voting and Local Resistance to Public Policy

by

Leah C. Stokes

Submitted to the Department of Political Science
on April 27, 2015 in Partial Fulfillment of the Requirements
for the Degree of Master of Science in Political Science

ABSTRACT

Retrospective voting studies typically examine policies where the public has common interests. By contrast, climate policy has broad public support but concentrated opposition in communities where costs are imposed. This spatial distribution of weak supporters and strong, local opponents mirrors opposition to other policies with diffuse public benefits and concentrated local costs. I use a natural experiment to investigate whether citizens living in proximity to wind energy projects retrospectively punished an incumbent government because of its climate policy. Using both fixed effects and instrumental variable estimators, I identify electoral losses for the incumbent party ranging from 4-10%, with the effect persisting 3 km from wind turbines. Voters also discriminate by correctly punishing the level of government responsible for the policy, providing evidence that voters are informed. I conclude that the spatial distribution of citizens' policy preferences can affect democratic accountability and exacerbate political barriers to addressing climate change.

Thesis supervisor: Christopher Warshaw
Title: Assistant Professor of Political Science

Acknowledgements

Thank you first and foremost to my excellent advisors, Chris Warshaw and Danny Hidalgo, who both provided countless hours of support, guidance and encouragement. You are both model academics: humble, kind, brilliant, warm and open. Any student who finds you is lucky to have landed in such fertile soil for intellectual growth. I hope to model my own advising after your own.

Thank you also to Jens Hainmueller for inspiring a love of statistics, for first encouraging me to pursue this project, and for telling me to keep going until I struck oil. Thank you to Chad Hazlett for your support, wisdom and friendship, and for being the best math teacher I've ever had. Thank you also to Gina Bateson, Adam Berinsky, Andrea Campbell, Ken Oye, Kathy Thelen and Teppei Yamamoto for teaching me and helping me find a place in the department.

Thank you to Liqun Chen for first-rate assistance with mapping and for a wonderful research collaboration. Thank you Dan de Kadt for coding assistance and friendship. And finally, thank you to Rebecca Silverman, Kathryn Buggs and Nicole Goodman for additional assistance with this project.

Thank you to Laura Chirot and Alex Hertel-Fernandez for insightful feedback on many drafts, as well as your friendship. Your steadfast support of this research has buoyed my spirits on occasions too numerous to recall. Thank you to my other friends in the department for learning alongside me and welcoming me into your community: Krista Loose, Justin de Benedictis-Kessner, James Dunham, Kelly Pasolli, Sara Chatfield, Tess Wise, Chase Foster, Tom O'Grady, Yiqing Xu, Mike Sances, and Yue Hue.

Thank you also to the participants at ONSEP 2013, the UT Energy Forum 2014, APSA 2013, MPSA 2014, CPSA 2014, DUSP PhD workshop, GSWIP, Georgetown's McCourt School, Michigan's Ford School, University of Toronto's Political Science department, and UCSB's Political Science department for feedback on earlier drafts. Thank you also to several anonymous reviewers for helpful feedback and guidance.

Finally, thank you to Matto Mildenberger, for endless, patient support with framing, coding, data cleaning, and much more, not to mention numerous pep talks during moments of doubt. Your unfailing backing is the wind in my sails. You are the best partner anyone could hope for: an equal, brilliant, cheerleader. I dedicate this work, begun just before our marriage, to you.

Introduction

Democratic accountability depends on voters' ability to retrospectively evaluate public policy (Key, 1966; Fiorina, 1981; Page and Shapiro, 1992). If citizens are able to connect policy to politicians' actions during elections, then politicians should have an incentive to align policy with the public's preferences, enhancing welfare and democratic accountability (Erikson et al., 1993; Stimson et al., 1995; Wlezien, 1995). To date, studies have found some evidence of retrospective voting largely by examining cases where citizens have common interests, such as economic performance (Healy and Malhotra, 2013; Anderson, 2007). However, less attention has been paid to retrospective voting on public policies that divide citizens' preferences. Politicians are less likely to propose and enact policies with concentrated costs and diffuse public benefits (Arnold, 1992; Wilson, 1980; Olson, 1965). What is less appreciated is that when policies are enacted that concentrate costs spatially, policy losers can use electoral institutions to amplify their voice. Since voters cast their ballots where they live, spatially concentrated opponents can send a clear signal to politicians if they retrospectively punish the government. Crucially, this dynamic can lead to an accountability failure: if there is a concentrated, local opposition and a diffuse, supportive public, the government may receive a signal to adjust a policy that the majority supports.

Climate policy is one area where costs are imposed on specific communities causing citizens' policy preferences to be spatially distributed. For this reason, studying climate policy can provide insights into retrospective voting when preferences are divided. Climate change requires societies to transform their electricity system away from fossil fuels, such as coal and natural gas, towards low-carbon technologies, including wind and solar power (Caldeira et al., 2003). Public opinion polls consistently show majority support for addressing climate change and building renewable energy infrastructure in advanced democracies (Ansolabehere and Konisky, 2014; Aldy et al., 2012; Lachapelle et al., 2012). However, citizens living near divisive facilities, such as wind turbines, often resist these projects. Similarly, carbon-dependent communities, such as coal mining areas, will lose from climate policy.

Consequently, transforming the electricity system to address climate change entails a distributional dilemma: although the benefits are global, the negative effects associated with the solution—renewable energy infrastructure in people's backyards or job losses in fossil fuel dependent communities—are local. This paper examines whether spatially concentrated citizens punish the government because of a climate policy's local cost imposed through wind energy projects.

Energy and environmental policy are not common domains for studying retrospective voting, in part due to an assumption that these issues may not be salient during political campaigns. But, the public has shaped energy policy through elections for over a century, beginning with referenda in the 1890s (Hirsh, 1999). This trend continued with ballot initiatives, spanning a 1935 vote in Redondo Beach, California that banned oil development (Smith, 2002) to recent climate policy votes in California, Michigan and Colorado during the 2000s. Similarly, in communities affected by energy infrastructure, climate policy can mobilize citizens. Many individuals believe that wind energy projects create local negative externalities, including lowered house prices and noise, alongside environmental, aesthetic and perceived health impacts. When residents mount opposition against local projects they are often characterized as NIMBYs - people who say "Not In My Back Yard" to new infrastructure (Rabe, 1994; van der Horst, 2007; Michaud et al., 2008).[1] Resistance to locally divisive projects with diffuse public benefits is not unique to energy infrastructure. The same pattern is found when governments site airports, housing projects, roads, subways, hospitals, jails, and waste facilities (Aldrich, 2008). However, causally determining whether citizens punish incumbent governments for controversial facilities is difficult due to selection bias in project location and insufficient survey data (Smith, 2002). By contrast, the design of a recent climate policy provides a natural experiment that can be used to test whether citizens retrospectively punish governments for policies with concentrated local costs and

[1] In practice, citizens may not only be thinking about their own self interest when they resist projects, but may generally oppose putting infrastructure in anyone's backyard. Citizens may also be opposed to the process surrounding siting decisions, the impacts projects have on ecosystems, or have other concerns.

diffuse public benefits.

In 2009, the Liberal government in Ontario, Canada expanded its ambitious climate policy, passing the Green Energy Act. This policy allowed corporations, communities and individuals to build wind turbines and other renewable energy projects throughout the province, signing long-term contracts with the government to sell their energy (Stokes, 2013; Yatchew and Baziliauskas, 2011). After implementation, support for the policy was divided spatially. In 2010, wind energy remained extremely popular in Ontario, with 90% of respondents in a representative poll supporting wind energy in their region.[2] However, the policy also led over 50 anti-wind groups to mobilize against the incumbent Liberal Party across the province, staging protests outside campaign events, posting signs and unsuccessfully attempting to block projects (Stokes, 2013). In the Fall 2011 election, the incumbent Liberal Party government lost its majority in the provincial legislature by one seat. The media attributed this loss to popular dissatisfaction in rural areas with the climate policy, which was a major issue during the campaign.[3] After the election, the government responded to local anti-wind opponents by modifying the policy, placing a de facto moratorium on new wind turbines.

The policy design meant that between 2006 and 2011, developers were able to build wind projects where they were able to secure land. Since the policy provided a significant return on investment, companies acted quickly to develop projects in locations with the best wind resources to maximize their profits. The policy's design eliminated communities' ability to select into or out of receiving wind energy infrastructure by taking away local planning authority. Further, private actors, not the government, chose project locations. The Ontario case thus offers two strategies to identify the causal effect of local wind projects on voting behavior. First, exploiting the fact that there was no selection bias in wind energy project siting decisions allows for causal estimation with fixed effects. Second, to the degree that wind resources are orthogonal to political geography, wind speed can be used as an

[2] Ipsos Reid poll, June 25-30 2010.

[3] See for example, John Spears. "Wind turbines churn rural votes." *Toronto Star*, 2011; Michael Warren. "How Ontario's Liberals can win back rural Ontario." *Toronto Star*, 2013.

instrument for turbine location, offering a complementary identification strategy.

Using this natural experiment, this study investigates whether Ontario citizens retrospectively punished the incumbent government for nearby wind turbines. Through an original panel data set that locates all proposed and operational wind turbines, merged with voting and census data at the precinct level, I use a variety of estimators to examine differences in incumbent vote share and turnout in precincts with and without wind turbines. For precincts with a proposed or operational turbine, the vote share for the governing party declined by 4-10% compared to similar precincts without a turbine. I demonstrate the effect with both fixed effects and instrumental variable estimators, and the results are robust to many alternative specifications. Further, using a measure of distance to turbines, I directly estimate the geographic scale of NIMBY effects: the effect persists in precincts 3 km away from proposed or operational turbines. Taken together, these results suggest 6,050 votes were lost in the 2011 provincial election due to citizens' dissatisfaction with wind turbines. Most of this shift in votes came from existing voters changing their votes, rather than from new voters being mobilized. To test whether citizens' behaviors were informed, I also check whether the same citizens punished the same party nationally over the same time period, despite the fact that the national Liberal Party had no jurisdiction over the climate policy. I find little vote share loss at the national level, suggesting citizens' were informed: they only punished the level of government responsible for the climate policy.

More broadly, the paper provides insight into political accountability when citizen's policy preferences are divided across space, with implications for other kinds of policy. To whom is the government accountable: mobilized local opponents or the quiet but largely supportive public? The paper's findings pose two key problems for democratic governance and accountability. First, I demonstrate that a small group of spatially concentrated citizens with intensely held preferences are able to create incentives for politicians to abandon policy, bucking the preferences of 90% of the public. This asymmetry in mobilization can cause public policy to reflect a vocal and local minority, rather than the diffuse majority's preferences. Second, even if the policy remains stable, projects may simply be relocated to another community's backyard with lower political mobilization, less wealth and lower social capital, creating clear equity problems for accountability to local opponents. This dynamic raises further questions about who benefits from accountability if individuals and communities with greater resources are more likely to mobilize and express their preferences to the government. Substantively, the paper also highlights important but under appreciated mass politics associated with addressing climate change, given the spatial distribution of winners and losers.

The Distributional Politics of Climate Policy and Divisive Facilities

Although climate change is often viewed as a global collective action problem (Ostrom, 2010; Stern, 2006), it is also a distributional challenge with implications for domestic politics. Most research on domestic climate change politics has examined interest groups in carbon-intensive industries with concentrated stakes in status quo technologies. For example, representatives from US Congressional districts with emissions-intensive industries are less likely to support climate legislation (Cragg et al., 2013). Climate policies may also redistribute benefits to renewable energy companies through subsidies (Aklin and Urpelainen, 2013; Urpelainen, 2012). Less attention has been paid to the distributional politics associated with public support for climate policy (Aldy et al., 2012; Bechtel and Scheve, 2013). A small literature focuses on public perceptions of renewable energy given higher costs (Ansolabehere and Konisky, 2014). However, climate policy also imposes costs on communities through job losses and negative externalities from divisive renewable energy infrastructure. The political ramifications of these climate policy costs are less appreciated.

Like other divisive facilities, energy infrastructure typically imposes concentrated, local costs on proximate communities, even as these projects provide benefits to society broadly

(Ansolabehere and Konisky, 2009; Boudet and Ortolano, 2010). Thus, when asked about energy policy, people are concerned about the local rather than the global impacts (Ansolabehere and Konisky, 2014). For citizens, the local negative impacts from energy infrastructure are immediate, personal and visible, whereas the global climate consequences are in the future and harder to understand (Krohn and Damborg, 1999). This means that people are more likely to mobilize against an energy development in their backyard—whether a wind turbine or nuclear power plant—than mobilize to support climate policy broadly. In the case of fossil fuel infrastructure such as pipelines and power plants, there are both local and global pollutants. Here, local citizens' interests are aligned with addressing climate change. However, with renewable energy, local citizens interests often run counter to climate policy. Further, both wind and solar energy need more land area than conventional electricity sources (Carley, 2009). Consequently, these technologies will be distributed across the landscape exacerbating the political challenge of addressing climate change.

These patterns of local resistance are not limited to energy infrastructure or climate policy. Communities also commonly resist other divisive facilities with diffuse public benefits and concentrated local costs, including airports, subways, roads, housing projects, trailer parks, and jails (Rabe, 1994; Clingermayer, 1994). These locally unwanted projects fall into two categories: those that have potentially harmful environmental and health impacts for the local community, and those that do not (Schively, 2007). Energy projects generally fall into the former category. But with renewable energy, particularly wind projects, harmful health impacts are not supported by peer-reviewed, scientific studies (Bolin et al., 2011; Crichton et al., 2014). Despite increasing demand for divisive facilities due to consumption and population growth, locally contentious projects remain difficult to site, with governments largely relying on coercion (Aldrich, 2008).

What are the political effects when divisive facilities are forced on communities? Politicians may worry about siting public bads in their district, but it is less clear whether citizens are able to connect their discontent with nearby divisive facilities to government policy, and punish incumbents at the ballot box. There are two broad views on the relationship between citizens' public policy perceptions and vote choice. One view holds that voters are myopic and following elites' lead with little knowledge or interest in specific policies (Campbell et al., 1960; Converse, 1964; Zaller, 1992; Bartels, 1996; Achen and Bartels, 2012; Lenz, 2012). According to this perspective, we should not expect citizens to understand public policy enough to punish the government for an unwanted climate policy. A second view argues that the public, in aggregate, holds policy knowledge, demonstrates retrospective voting and exhibits lasting electoral responsiveness to policy (Key, 1966; Fiorina, 1981; Page and Shapiro, 1992; Lupia, 1994; Johnson et al., 2005; Bechtel and Hainmueller, 2011; Erikson and Stoker, 2011). In this view, it is conceivable that citizens could understand climate policy and punish the government for imposing costs on their communities.

Studying the political consequences of local resistance to divisive facilities has proven difficult. National public opinion surveys do not include enough respondents living near projects to examine whether nearby citizens' beliefs differ from those living farther away (Smith, 2002). When there is sufficient data, selection bias in project location may make places with and without projects incomparable. Large energy infrastructure, such as coal or nuclear plants, are often strategically placed in communities with more minorities, that are poorer or have lower social capital (Mohai et al., 2009; Aldrich, 2008). Finally, many studies rely on mail in surveys, case studies, interviews or newspaper data (Baxter et al., 2013; McAdam and Boudet, 2012; Warren et al., 2005). While these studies are able to identify dynamics surrounding political mobilization, they are less able to estimate the extent of resistance to projects.

As a consequence of these identification challenges, research on local resistance to energy infrastructure has produced divergent estimates (Wolsink, 2007). Some researchers find little resistance related to proximity (Wolsink, 2000), others find that nearby communities are more supportive (Michaud et al., 2008; Baxter et al., 2013; Krohn and Damborg, 1999), while still others find proximity is associated with lower support for projects, including wind turbines (Jacquet, 2012; Swofford and Slattery, 2010). It is also not clear whether people oppose divisive facilities for short periods of time and come to accept their presence in the community once built (van der Horst, 2007; Devine-Wright, 2005; Wolsink, 2007); or, if discontent and political mobilization continues after projects are operational.

This paper overcomes these identifications challenges through a natural experiment that effectively assigned divisive facilities to communities based on an exogenous factor. This policy allows for an examination of whether citizens living in proximity to divisive facilities disproportionately punished the government in a subsequent election. Further, the number of projects developed in this case creates a large enough sample size to facilitate statistical inference. As a result, Ontario's climate policy provides an improved empirical setting to examine whether citizens punish incumbents for placing controversial infrastructure in their backyards.

Ontario's Climate and Renewable Energy Policy

In Ontario's 2003 election, the centrist Liberal Party ran on a platform to phase out coal. At the time, coal accounted for around one-fifth of the electricity mix, making this a significant commitment to addressing climate change. After they won the election, the Liberals enjoyed majority control of government for 8 years. During this period, the Liberals enacted a 'feed-in tariff' (FIT) policy to promote wind energy development. This policy aimed to facilitate the coal phase-out, reduce greenhouse gas emissions, and create jobs in the province.

Feed-in tariffs are a climate policy that focuses on subsidizing new, low-carbon technologies, rather than imposing costs on existing fossil fuel technologies. A FIT sets a price per unit of energy, given a certain technology (i.e. wind), and offers long term contracts to private sector project developers. FITs are used to drive investment in new and more expensive technologies, particularly renewable energy including wind, solar and geothermal (Mitchell et al., 2006). The idea is to create investor certainty, to increase renewable energy projects and reduce prices through innovation. FITs have been used extensively in Germany, the United Kingdom, Spain and over 80 other jurisdictions worldwide.

In 2006, the Liberal government passed a small-scale FIT policy. In 2009, they expanded this policy through the Green Energy Act. Crucially, this law preempted much of the Planning Act, taking away decision making authority from local communities. Similarly to a prominent hydraulic fracturing law in Pennsylvania, the law took away local jurisdiction to approve or reject projects (Hill and Knott, 2010; Malin, 2014). As a result, communities could neither select into nor out of receiving a wind project. This decision was made because of concerns that local anti-wind activists would block projects through municipal governments. As the former Premier, the chief executive for Ontario, stated:

> "To leave decision making, about setbacks for example, to the local authorities would be to effectively keep the [renewable energy] sector out of Ontario...it would have driven a stake through the heart of the FIT program before it even got out the door. The local politics are just so difficult to manage. So the [Mayor] says, 'it would be good for me and for the municipal finances, but I can understand the neighbours don't need that hassle; we'll just tell them to go elsewhere.' And then word gets out that Ontario is not really serious. Because the proponent says, 'We've been sent to 7 different municipalities so far and we can't get in because local opposition has mounted and its just not possible.' So to be realistic about these things, what you've got to do is have a single, province-wide standard."
>
> — Dalton McGuinty, Premier of Ontario (2003-2013)[4]

Since the FIT price was set based on estimated project costs plus a significant return on investment, many companies acted quickly to develop projects. Between 2006 to 2012, project proponents were able to build wind turbines wherever they were able to secure land,

[4]Interview with author, December 6, 2013.

often by signing agreements with farmers and other land owners. After 2009, municipalities were not able to block proposals if local citizens disapproved of the proposed infrastructure. Only one project was stopped or cancelled between 2006 and 2012, and this was because of financing rather than local opposition.[5] Ultimately, government efforts to promote wind energy through the FIT were successful: despite widespread protests, Ontario signed contracts by 2011 for around 5,500 MW of wind energy, equivalent to building around 10 coal plants or 5 nuclear reactors.

During a Fall 2011 election, the incumbent Liberal government lost their majority in the provincial legislature by one seat, a loss the media attributed to popular dissatisfaction with the wind turbines. This judgment was reasonable: by the 2011 election, every district with a wind turbine had at least one anti-wind group. These groups engaged actively with the political process, staging protests outside Liberal Party campaign events and posting signs throughout their communities. They also successfully encouraged 78 municipalities to pass resolutions against wind turbines, even though such resolutions had no legal effect. Media coverage of wind protests in Ontario also grew between 2003 and 2011, suggesting significant political mobilization. As one columnist summarized: "[M]assive corporate wind farms have been built in rural communities against their wishes...Ten of the 18 seats lost in 2011 were in rural ridings targeted by anti-wind coalitions. They made the election a referendum about halting wind farm development and restoring local decision-making in their communities."[6] The Liberals' two rival parties, the center-right Progressive Conservatives and the left-wing New Democrats, both committed to eliminating the FIT policy if elected. As a result, there was a sharp and well-publicized difference between the policy commitments of the incumbent Liberal Party and their political opponents.[7] Thus, citizens faced a clear choice at the polls: support the incumbents and their climate policy, or vote for an opposition party.

[5]A moratorium on offshore wind projects was put in place during this period, affecting 2 projects. However, these projects are outside the scope of this study, which only examines onshore turbines.

[6]Michael Warren. "How Ontario's Liberals can win back rural Ontario." *Toronto Star*, 2013.

[7]See Supporting Information (SI) for more qualitative evidence of wind project resistance.

Data, Research Design, and Methodology

This paper uses a natural experiment to identify the causal effect of proximity to wind energy projects on voting behavior. Natural experiments depend on exogeneity of treatment assignment to ensure that treatment effect are unbiased. In this case, evidence suggests that communities were not able to select into nor out of receiving a wind project because the policy took away local planning authority. Consequently, communities could not block projects. Similarly, it is unlikely that communities were able to select into receiving a turbine. Since the policy offered a flat rate to project developers, developers would make the greatest profit from building projects in sites with the highest wind resources. For these reasons, I argue that wind project treatment assignment in Ontario between 2006 through mid-2013 was plausibly exogenous to political boundaries: wind farms were proposed and built where wind resources were strongest, and wind resources are conditionally orthogonal to the political preferences of voters.

In addition, this paper uses two strategies to identify the causal effect of proximity to wind energy developments on voting behavior: fixed effects estimators and instrumental variable estimators. These two estimators are complimentary in that they rely on different assumptions to identify the same causal effect. The fact that both estimation strategies provide similar results provides additional certainty that observed changes in political behavior can be causally attributed to nearby wind energy developments.

Fixed effects estimators rely on variation within a given political unit over time to identify causal effects. Fixed effects estimators do not thus rely on the presence of cross-sectional, as-if random assignment of turbines to precincts. Instead, the main identification assumption is parallel trends in the treated units' potential outcomes, absent treatment, and the control units. By contrast, the instrumental variable estimator uses a cross-sectional approach to compare different units. Wind speed cubed is the instrument used to predict turbine placement. This second approach is a robustness check on the fixed effects specifications, because it addresses any selection bias in wind turbine location. With this broad overview of the identification strategy, the specific data, research design and methodology can be introduced in greater detail.

Data

In Ontario the main political unit is a provincial electoral district, also called a riding. Each district elects a provincial representative. To ensure plausible control units, a subset of Ontario's districts where we could reasonably expect wind turbines to be developed was created by excluding all districts without any wind energy proposals. Consequently, the sample includes 26 districts where wind projects were proposed or operational, out of 107 total districts, or just over a quarter of province's total land area (Figure 1). In this way, major cities such as Toronto, are excluded from the data, since these are implausible locations for wind projects. Ontario districts are made up of precincts which have an average of 350 voters.[8] The unit of analysis for this study is the precinct, where electors cast their vote for their provincial representative at a nearby polling place. All 6,186 precincts within these 26 districts are included in the data set, providing both treated and control units. At this political scale, data are available for the 2003, 2007 and 2011 elections.

There were some changes in the number and size of precincts between 2003, 2007 and 2011. As a result, the panel relies on matching 2007 precincts with the 2003 and 2011 precincts that have the largest overlapping area. Ultimately, these boundaries changes were small and administrative, and did not affect a district's electorate, but rather, created minor alterations in where individual citizens cast their vote.[9] National elections data are also available over a similar time period at the precinct level. Three national elections were used because they correspond most closely to the provincial elections: 2004, 2008 and 2011. With the national data, the same approach was used to match each federal precinct to a 2007 provincial precinct.[10]

[8]Elections Ontario calls precincts 'polling divisions.' Precincts in the sample area vary from 4 to 1722 electors.

[9]See Supporting Information.

[10]In 114 cases, federal data was missing for turnout or liberal vote share. These cases were set at the mean

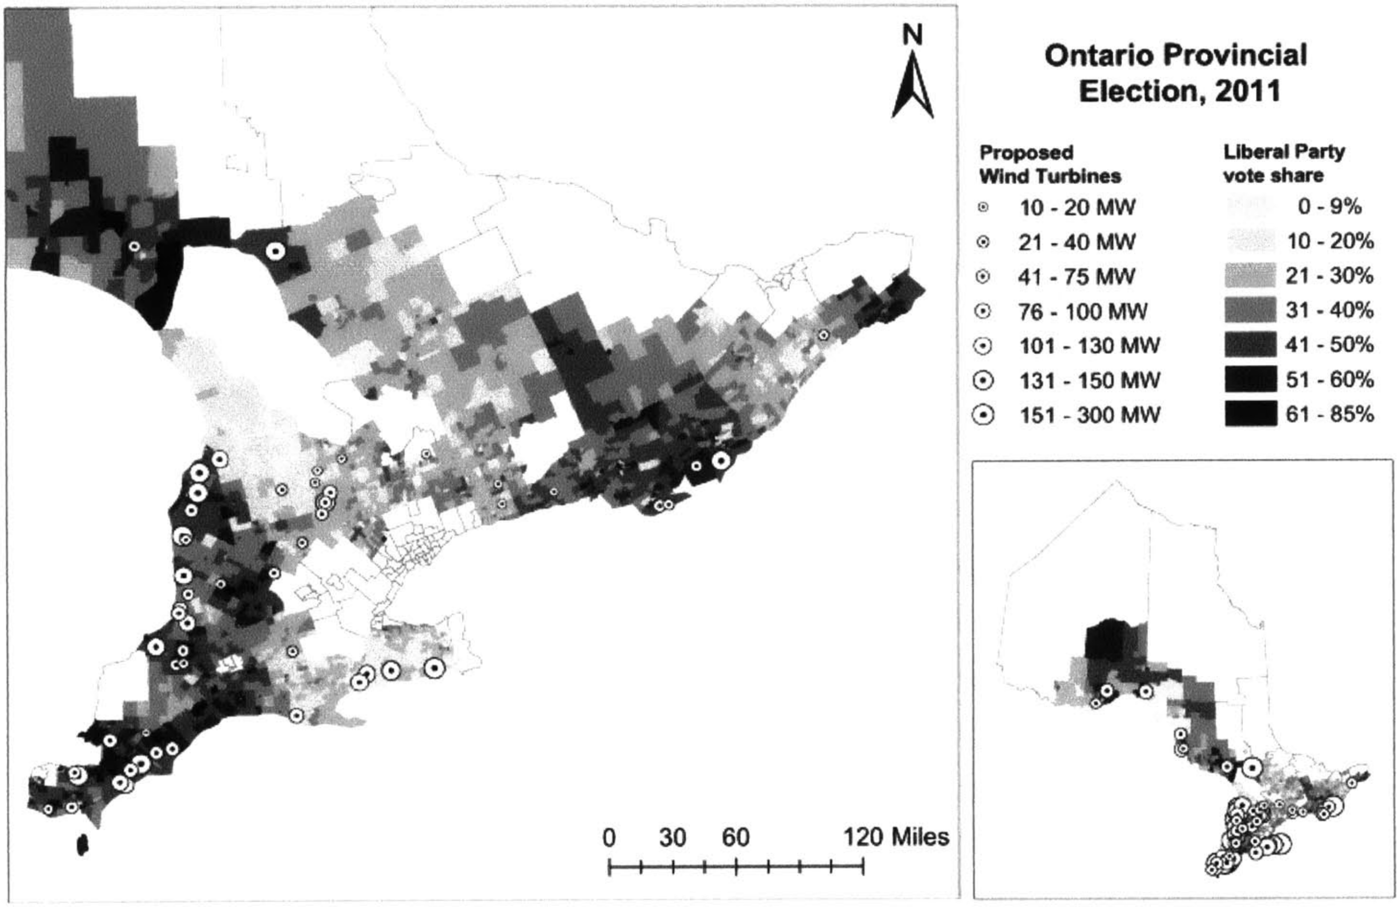

Figure 1: Governing party's vote share in 2011 by precinct in 26 districts, and proposed wind projects, by scale of project. The map also shows the districts excluded from the sample in white.

Census data are important for assessing pre-treatment balance and to allow the models to include controls for potential time-varying confounders. Census data from Statistics Canada are available in 2001, 2006 and 2010 at the dissemination area. With an average of 640 people, a dissemination area is around twice the size of a precinct. Dissemination areas do not track precinct boundaries. For this reason, each precinct's census data was estimated using a geographically-weighted sum.[11] This approach introduces minor measurement error

level of each variable for that year. In a few cases where turnout was greater than 100%, due to errors in official data, turnout was set to 100%.

[11] For each precinct overlapping census dissemination areas were identified. The census values for each precinct were estimated by weighting the sum of each variable based on the amount of overlap between the precinct and each dissemination area. This approach assumes that census variables are evenly distributed

into the panel's census covariates.

Treatment Variable and Instrument

This study uses original panel data on wind energy projects in Ontario. A wind project is defined as any onshore wind turbine development proposed or operational in the province before the 2011 election with a size greater than 10 MW. This size is used in the FIT program to distinguish between macro-contracts (industrial-scale facilities) and micro-contracts. Data was gathered using the Ontario Power Authority's (OPA) website of active wind projects and the Independent Electricity System Operator's list of wind projects pending grid connection.[12] These lists were supplemented with public Project Description from project developers that show each project's location, technology and timeline. Information was collected for 65 industrial-scale wind projects across the province, totaling 2,753 proposed or operational turbines. Figure 1 shows a map of the wind turbine data overlaid with the provincial, precinct-level Liberal vote share in 2011. This map demonstrates that wind projects are geographically clustered, likely because of wind resource availability and electricity grid infrastructure, but that they are located in districts with both high and low levels of support for the incumbent party.

All turbines were geolocated and assigned to a precinct.[13] Each turbine was also classified as proposed or operational before the 2007 and 2011 elections. From these data, two binary treatment variables were constructed. The first identifies precincts that had an active

across space. This is a plausible assumption, particularly for variables that do not appear infrequently, because the census units are small. This approach was validated at the municipal level, where original census data are available, with very similar results compared to the actual census data. Some census data were censored by Statistics Canada. When more than 25% geographically of the census data in a precinct was censored, that variable was set to NA. When less than 25% of the data was censored, the missing area for that covariate was set at the mean level for that precinct.

[12] I assume that if the project is not listed in the OPA's list of projects with a FIT contract, the project is not considered proposed to communities at a given moment in time.

[13] For four projects, only a highly bounded geographic region for the project was shown in project description documents. In these cases, the turbines were randomly placed within the project's boundaries. If this approach affects treatment assignment, it would tend to increase the area coded as treated, since it caused all precincts within the project's geographic boundaries to be coded as treated. Given precincts' small size, this approach should not result in miscoding.

proposal or an operational turbine within its boundaries. The second identifies precincts with an operational wind turbine within its boundaries. There were 184 precincts that were treated with a proposal and 52 with an operational turbine by 2011, out of 6,186 precincts. In addition, each precinct was classified in terms of its distance to the nearest turbine in 1 km increments.

Wind speed is theoretically orthogonal to precinct boundaries but predicts the placement of wind turbine locations. As a robustness check, wind power data, which is approximately wind speed in meters per second to the third power, was used as an instrumental variable for turbine location. Wind power data was collected from Land Information Ontario, which provides this data for wind developers. This data gives the wind resource strength throughout the province at a 100 m_2 spatial resolution. The average wind power resources for each precinct was calculated through a spatial merge.

Fixed Effects Models

The first empirical strategy relies on a panel data set analyzed using fixed effects estimators. The core model includes both unit-specific fixed effects at the precinct level and time-period fixed effects for each election year. This basic specification is given by:

$$Y_{it} = \alpha_0 + \gamma_i + \delta_t + \alpha D_{it} + \epsilon_{it} \tag{1}$$

where Y_{it} is the vote share for the Liberal Party, or the voter turnout, in precinct i in time t; α_0 is the intercept; γ_i is the precinct fixed effect; δ_t is the election year fixed effect and αD_{it} is the treatment effect of a turbine being proposed or operational within that precinct's boundaries. This approach, which uses precincts as the unit of analysis, gives equal weight to every precinct in the 26 districts with at least one wind turbine proposal. These models assume that $Cov(\epsilon_{it}, D_{it}) = 0$.

The main threat to causal inference using fixed effects estimators is time-varying confounders—

variables that change differentially in the treated and control groups. Thus, for the average treatment effect to be interpreted causally, the parallel trends assumption must hold. To provide evidence to support the parallel trends assumption, the mean level of the outcome variable by year is plotted below, showing parallel trends pre-treatment (Figure 2). In addition, census data are included to control for observable potential time-varying confounders. In addition alternate specifications subset or reweight the data, and some models include precinct linear time trends. If the parallel trends assumption holds, even in the absence of perfect covariate balance between the treated and control units, the effect is identified.

Instrumental Variable Model

Although qualitative evidence suggests that treatment assignment was orthogonal to political boundaries, as a robustness check the effect can also be estimated using wind power, which is approximately wind speed to the third power, as an instrument to directly address endogeneity concerns. This empirical strategy uses a cross-section of the 2011 data rather than the panel used above because the instrument, wind power, does not vary over time. Only turbine proposals were used in this model because the timing of specific wind turbines becoming operational is unrelated to wind power, and instead of function of the initial proposal date. In addition, these data do not include precincts that were already treated prior to the 2007 election, defined as being within 3 km of a proposal, nor precincts that are missing wind data because they are too small. The instrument is estimated using two-stage least squares regression:

$$\text{First stage: } D_{it} = \pi_0 + \delta_i + \pi_1 Z_i + \gamma^\top X_i + \epsilon_i \tag{2}$$

$$\text{Second stage: } Y_i = \alpha_0 + \delta_i + \alpha_1 D_{it} + \theta^\top X_i + \eta_i \tag{3}$$

Where D_{it} is the treatment variable, predicted by Z_i, the wind power in that unit, along with a district fixed effect, δ_i, and covariates, X_i. In the second stage, Y_i is the change in the Liberal Party vote share in that precinct between the 2007 and 2011 elections, which is predicted by the reduced form of the first stage and the same district fixed effects and covariates. Using the change in vote share allows the model to incorporate how the precincts are varying overtime and is equivalent to including a precinct fixed effect.

This model relies on the following assumptions: ignorability ($Cov[\epsilon_i, Z_i] = 0$), the exclusion restriction ($Cov[\eta_i, Z_i] = 0$), relevance ($\pi_1 \neq 0$) and monotonicity ($\pi \geq 0$ or $\pi \leq 0$ for all i). While most of these assumptions are not directly testable, the results section and SI provide evidence for why they are plausible in this case. Unlike the fixed effects approach, this estimation strategy relies on the assumption that wind speed varies as-if randomly. It also assumes that wind speed influences the location of wind-turbines.

Using matching to pre-process the data before using an IV estimator may strengthen the instrument and correct for biases from confounders if the instrument departs from as-if random assignment (Keele and Morgan, 2013). To improve the balance in the sample, Mahalanobis distance matching was first used to pair each treated unit with a control unit based on observable characteristics. The data was balanced on four variables: the average home price pre-treatment in 2006 (log), the population with a university degree (%), median income (log), and population density (log). These variables were chosen because they predict the outcome variable; thus, controlling for them should work against finding an effect. These matched pairs were then put into the instrumental variable regression with flexible geographic controls to compare the change in vote share in nearby places with and without a proposed turbine. Geographic controls, in both the first and second stage, included longitude, latitude, both variables squared and their interaction. Since many projects are near the Great Lakes, due to wind resources, the model also controls for distance to these lakes and this variable squared.

Results

This paper examines whether people living near wind turbines voted against the incumbent provincial government to punish them for a climate policy creating incentives for wind infrastructure in their community. This analysis consistently finds an effect size of around -4-8% for proposed wind turbines and -7-10% for operational wind turbines on the governing party's vote share. Figure 2 shows the trend in the governing party's vote share over time. Given the slopes before the majority of units were treated, between 2007 and 2011, the parallel trends assumption necessary for inference using fixed effects estimators is likely plausible. For proposals, there is excellent balance without matching on the outcome variable between treated and control units pre-treatment, in 2003.

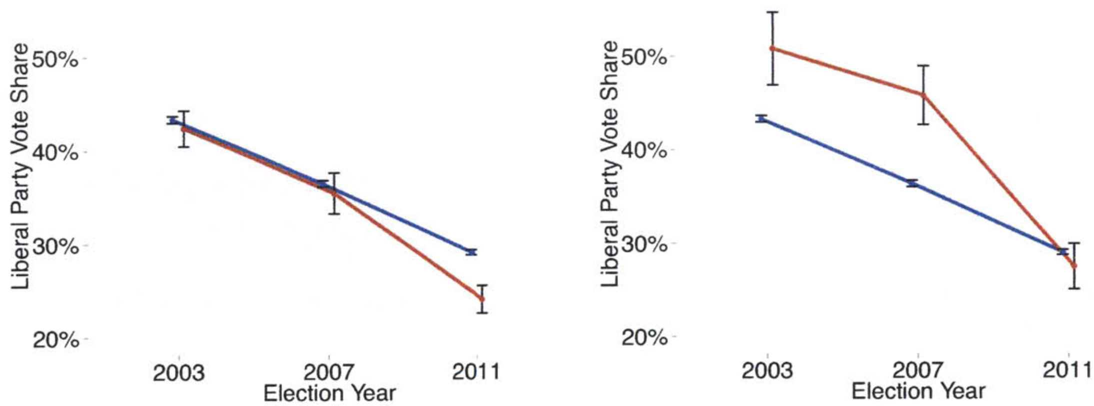

Figure 2: The average governing Liberal Party's vote share in treated (red) and control (blue) precincts with proposed (left) or operational (right) wind turbines, for each provincial election between 2003 to 2011. The error bars show 95% confidence intervals.

Fixed Effects Models

The treatment effect of having a proposed or operational wind turbine within a precinct on the governing party's vote share is first estimated using fixed effects estimators with election year and precinct fixed effects (Table 1). The government's vote share went down by 4% in precincts with a proposal for a turbine and 10% in precincts with an operational turbine, with both point estimates statistically significant.[14] However, through examining balance tables, it is clear that towns within these 26 districts, where turbines could not plausibly be placed, are still in the sample as control units. This likely makes the control and treatment units somewhat incomparable. As a result, the balance on pre-treatment observable census data for treated and control units in 2003 is poor on all covariates; still, the pre-treatment dependent variable is notably well balanced for places with and without proposed turbines.

In an attempt to achieve better balance on the covariates across the treated and control units, the data was subset to rural areas with a population density of 400 people per km^2 or less in 2007, corresponding to Statistics Canada's definition. Since all of the treated units are rural, none are dropped. For both the proposal and operational treatment variables, the models using the rural precincts subset return a very similar point estimate as when the data includes all precincts as control units. Namely, these rural models find a 5% decline in precincts with a proposed turbine and a 10% decline in precincts with an operating turbine in support for the governing, provincial Liberal Party.

Third, I used entropy balancing, a form of preprocessing that reweights each unit based on its covariates so that the treated and control groups have the same moments, for example the mean and variance. Like matching, entropy balancing attempts to achieve better balance between treated and control units on observable covariates to reduce model dependence (Hainmueller, 2012). Using the 2003, pre-treatment data, the data was weighted on population in the precinct (log), population density (log), median income (log), average home value (log), unemployment rate, percentage housing ownership, percentage of people with a university degree, percentage of immigrants and each term squared. To avoid dropping data, units missing covariate information were assigned the mean level for that variable for that year.[15] As Table 1 shows, the results from using entropy balancing are very similar to

[14]These results are also robust to precinct linear time trends being included in the model, with the operational treatment variable significant at the standard 0.05 p-value level and the proposal treatment variable significant at the 0.10 p-value level.

[15]This should not affect the results, since only 1-4% of the data are missing for each census variable, and

Table 1: Effects of Wind Turbines on Incumbent Party Vote Share in Precincts

	Turbine Proposal					
	All Precincts (ATT)	All Precincts (ATT)	Rural Precincts (ATT)	Rural Precincts (ATT)	Balanced Precincts (ATT)	Balanced Precincts (ATT)
Proposed Turbine	-0.042^{***}	-0.039^{***}	-0.048^{***}	-0.046^{***}	-0.050^{***}	-0.050^{***}
(Treatment)	(0.009)	(0.009)	(0.009)	(0.009)	(0.011)	(0.011)
Population with		0.084^{***}		0.055^{*}		-0.069
University Degree (%)		(0.018)		(0.024)		(0.078)
Population Density		0.006^{***}		0.007^{***}		0.002
(log)		(0.001)		(0.001)		(0.004)
Unemployment Rate		0.001^{*}		0.000		0.001
		(0.000)		(0.000)		(0.002)
Median Income (log)		$0.013^{\dagger}$		0.008		0.022
		(0.007)		(0.009)		(0.022)
Immigrant		0.074^{**}		0.084^{*}		0.047
Population (%)		(0.027)		(0.038)		(0.084)
N Treated	184	184	184	184	184	184
N Control	6002	6002	2985	2985	6002	6002
Fixed Effects	Y	Y	Y	Y	Y	Y
	Turbine Operational					
Operational Turbine	-0.096^{***}	-0.092^{***}	-0.099^{***}	-0.098^{***}	-0.084^{***}	-0.080^{***}
(Treatment)	(0.014)	(0.014)	(0.014)	(0.014)	(0.015)	(0.014)
Population with		0.084^{***}		0.057^{*}		0.136
University Degree (%)		(0.018)		(0.024)		(0.102)
Population Density		0.006^{***}		0.007^{***}		0.007
(log)		(0.001)		(0.001)		(0.004)
Unemployment Rate		0.001^{*}		0.000		-0.001
		(0.000)		(0.000)		(0.002)
Median Income (log)		$0.013^{\dagger}$		0.008		0.006
		(0.007)		(0.009)		(0.035)
Immigrant		0.075^{**}		0.088^{*}		-0.028
Population (%)		(0.027)		(0.038)		(0.082)
N Treated	52	52	52	52	52	52
N Control	6134	6134	3117	3117	6134	6134
Fixed Effects	Y	Y	Y	Y	Y	Y

Robust standard errors, clustered at precinct level. Intercepts are not reported.

Significant at $^{\dagger}p < .10$; $^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$

the results using the full sample and the rural subset.[16]

All three models were also run including census covariates as controls, to address potential time-varying confounders. Covariates included population with a university degree, average home value, population density, the unemployment rate and median income. As Table 1 shows, including these covariates does not substantially change the results. Together, these convergent results suggest that there was a statistically significant and substantive decline in the governing Liberal Party's vote share because of proposed and operational turbines. Operational turbines appear to provoke the ire of a greater proportion of the population than proposals, likely because they are now visible, and therefore, more of the community knows about the development and its impacts.

How Large is the Backyard?

To examine how far the effect persists over space, I also examined whether vote share declines occur in precincts near turbines. This approach also addresses concerns that the treatment effect may be biased towards zero because of spillovers to nearby precincts. Spillovers are likely because people can see turbines for several kilometers. Further, Ontario regulations had a requirement for turbines to be setback a minimum of 550 to 1,500 meters from homes, suggesting that wind turbines' perceived negative effects persist for some distance.

Figure 3 shows the treatment effect within precincts grouped into 1 km increments away from proposed or operational turbines. These buffers can be visualized as 1 km wide rings, with the turbine at the center; for each group, the buffer moves 1 km further out. If the 1 km buffer touches any part of a precinct, it enters into that group for distance from a proposed or operational turbine. Each treated unit only enters into one group, so that the group of precincts 2 km away from the turbines does not include precincts 1 km away from the turbines. Further, when estimating the effect for each group, the sample excludes units

indeed running entropy balancing with these units dropped leads to similar results.

[16]While it is not reported here, Mahalanobis distance matching was also used and resulted in similar estimates for places with proposed turbines. The operational turbine treatment effect was estimated to be smaller, at around 5%.

less than 6 km away from the turbines as controls, to eliminate spillovers when estimating each group's treatment effect.

The point estimate for each treated group was calculated using fixed effects estimators, with time period and unit-specific fixed effects, and the standard errors were clustered at the precinct-level. The error bars in Figure 3 show 95% confidence intervals, which are large in part because the sample size for each group is quite small (43-199 treated units). For both proposed and operational wind turbines, vote share declines can be detected in precincts up to 3km away from the turbines. At 4 and 5 km away from the turbines, the effects are essentially indistinguishable from zero.

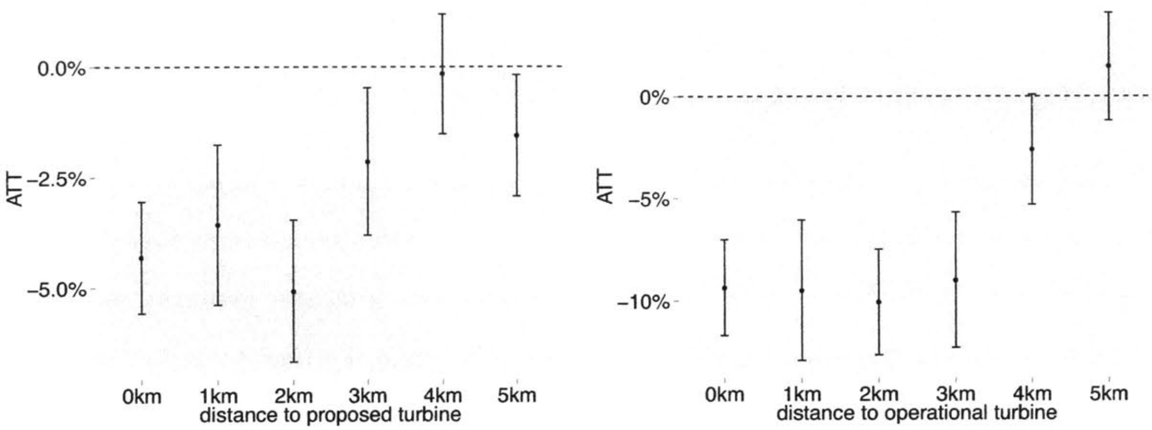

Figure 3: Estimates of the decline in the Liberal Party vote share using fixed effects estimators. Each group represents precincts that are a given distance from treatment, from 0 km (in precinct) to 5 km away from a proposed or operational turbine.

These results suggest that people voting in precincts that are several kilometers away from projects may still perceive negative effects on their communities. Further, it suggests a higher upper bound on the estimate of the government's loss in votes, since defining treatment as voting in a precinct within 3 km of a proposed or operational turbine increases the treated group's size. From a theoretical perspective, the detectable decline in vote share as a function of distance suggests the NIMBY effect is related to proximity to unwanted infrastructure. The effect is also similar to another study, which found effects within 1-5 km away from wind projects, although it only examined variation in 5 km increments (Warren et al., 2005).

Instrumental Variable Model

The following model uses wind power as an instrument for wind turbine proposals. This estimation strategy relies on cross-sectional matched-pairs from 2011 because wind speed does not vary over time, precluding the use of the panel. Here, the outcome variable is the change in the governing party's vote share in that precinct in 2011 compared to 2007, in effect a precinct fixed effect. Given the preceding section's results on how far the effect persists, these models examine all precincts within 3 km of a proposed turbine.

As Table 2 shows, wind turbine location is highly predicted by wind power. Using flexible geographic controls and a district fixed effect, these models suggest that the vote share decline was 8% in places with a turbine proposal. This effect is robust to including more controls and using different specifications. Compared to the previous estimates for votes within 3km of a proposed turbine, this slightly higher effect is plausible, given a potential downward bias on the fixed effects estimates due to measurement error caused by the data only being available at the aggregate, precinct level. Importantly, this same specification does not return significant results if the Liberal vote share in 2003, pre-treatment, is used as the outcome variable. This robustness check suggests the relationship is not a spurious correlation between wind resources and political boundaries.[17]

[17]The SI provides further information on balance and robustness checks.

Table 2: Wind Power as an Instrument for Turbines' Effect on Change in Governing Party's Vote Share, 2007 to 2011

	First Stage	Second Stage
Average Wind Power (log)	0.760***	
	(0.091)	
Proposed Turbine within 3 km of Precinct, 2011		−0.077**
		(0.026)
N Treated	354	354
N Control	354	354
District Fixed Effects	Y	Y
Other Geographic Controls	Y	Y
F-statistic on instrument	69	

Significant at $^{\dagger}p < .10$; $^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$

Mobilization versus Vote Swings

Thus far, we have seen strong evidence that wind turbines led to decline in support for the incumbent party. But this evidence does not tell us whether the opposition came from existing Liberal party voters switching their votes, or from new voters drawn to the polls as a result of the issue. It is possible that organized anti-wind groups were able to mobilize non-voters to vote against the government because of proposed and operational wind turbines. By 2011, there were over 50 active anti-wind groups in the province, who posted signs across communities that explicitly connected the Liberal Party policy to the wind projects. Did these groups bring new voters to the polls, who otherwise would not have voted, or did they convince previous Liberal supporters to change their votes?

Examining precincts with and without turbine proposals suggests that these groups were able to mobilize some new voters to go to the polls. As Table 3 shows, in precincts with an active proposal, voter turnout went up by 2%. Given that turnout was 49% across the province in the 2011 election, this change represents a substantively large increase. This increase in turnout accounts for about half of the fixed effects estimate of the change in the government's vote share due to proposals (5%). Given this result holds for places up to with 3 km away from a turbine proposal, it suggests communities were mobilized against proposals and that this mobilization accounted for some of the change in the government's vote share.

Table 3: Effect of Wind Turbine Status on Voter Turnout

	Proposal in Precinct	Proposal within 3km	Operational in Precinct	Operational within 3km
Voter Turnout	0.024***	0.016***	0.019	0.017†
	(0.007)	(0.004)	(0.015)	(0.009)
N Treated	184	492	52	145
N Control	6002	5694	6134	6041
Fixed Effects	Y	Y	Y	Y

Robust standard errors, clustered at precinct level
significant at $^{\dagger}p < .10$; $^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$

But, in places with an operational turbine, turnout effects were not statistically significant, although the lack of significance may be attributable in part to the smaller treated group for precincts with an operational turbine. Nevertheless, the turnout effect compared to the overall estimated change in the government's vote share (10%) suggests that the majority of the decline in places with an operational wind turbine occurred through existing voters changing their vote to punish the government for unwanted infrastructure, rather than through anti-wind groups mobilizing new voters. Together, these results suggest that, while mobilization of new voters occurred, most of the effect is likely attributable to existing voters changing their votes to punish the government.

Informed Backlash? National Election Results

While the provincial election results demonstrate that citizens living near wind projects punished the governing Liberal Party, we might also want to know whether voters were informed. On the one hand, their behavior appears to be based on policy knowledge, since the provincial Liberal Party enacted the policy and citizens living near unwanted wind infrastructure punished them at the polls. However, citizens could simultaneously be punishing other politicians despite their lack of control over the policy. Ontario municipal elections are non-partisan; however, we can examine whether the federal Liberal Party was also punished for wind energy developments. If citizens also punished the federal Liberal Party, we could interpret this as uninformed behavior.

Using the same estimator in Equation 1, I examine the federal election results over the same time period. The outcome variable is the federal Liberal Party's vote share, the treatment variable is whether there is a proposed or operational turbine in the precinct, and both unit and time period fixed effects are included in the model. As Table 4 shows, the presence of a proposed or operational turbine in a precinct does not lead to a statistically significant decline in the federal Liberal Party's vote share. However, when the sample is expanded to include precincts within 3km of a wind turbine, the effect is statistically significant and the point estimate is stable at a 1% decline for precincts with a proposed turbine and 2% for precincts with an operation turbine.[18] However, note that the magnitude of the effect is about one-quarter the size of the effect in the provincial models. This suggests that while a portion of citizens may be indiscriminately punishing both the provincial and federal parties, most citizens can discriminate and only punish the level of government responsible for the policy. This result suggests voters are informed about the climate policy.

[18]Although the results are not reported here, when census controls are included in these national models the results are stable.

Table 4: Effect of Wind Turbine Status on Federal Liberal Party

	Proposal in Precinct	Proposal within 3km	Operational in Precinct	Operational within 3km
Federal Liberal Party Vote Share	−0.009	−0.013*	−0.023	−0.023**
	(0.010)	(0.005)	(0.018)	(0.009)
N Treated	184	492	52	145
N Control	6002	5694	6134	6041
Fixed Effects	Y	Y	Y	Y

Robust standard errors, clustered at federal precinct level
significant at $^{\dagger}p < .10$; $^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$

Conclusion

Using a natural experiment, I demonstrate that voters pay attention to climate policy and can retrospectively punish incumbent governments for nearby facilities they perceive as harming their communities. Taking the average treatment effect across all models, wind turbine proposals were associated with a 5% decline and operational wind turbines with a 10% decline in the incumbent provincial government's vote share. Across all models the results remain negative and statistically significant, suggesting the null hypothesis, that there was no change in incumbent support as a consequence of nearby wind turbines, is unlikely. Further, the results provide evidence for the NIMBY hypothesis: proximity to infrastructure leads local communities to oppose developments. Unlike previous research, an effect was found after projects became operational, suggesting communities may continue to mobilize after divisive facilities are built. This result is intuitive: when projects become operational more citizens know about the project and experience its impacts. Finally, voters were informed, generally punishing the provincial Liberals but not their federal counterparts over the same time period.

Substantively, using the average point estimates and the average voter turnout per precinct, these results suggest the Liberal government lost around 6,050 votes during the

2011 election.[19] This means that wind turbines may not have changed any individual district's outcome, causing the government to lose its majority. Further, the public at large remained overwhelmingly supportive of the policy. A 2010 poll found more than 85% of Ontarians agreed that their local government should encourage wind developments, that they would support wind energy in their community, and that wind turbines have less health impacts than other energy technologies. Yet, the retrospective punishment that occurred in communities within 3 km of a turbine sent the government a different signal. Supporters and opponents' spatial distribution and asymmetry in mobilization meant that the government was able to hear policy opponents' views more clearly than policy supporters. Likely as a result of anti-wind opponents mobilizing and punishing them at the polls in the 2011 election, the Liberal government froze the wind policy in 2012.

More generally, this case demonstrates retrospective voting against a government in an understudied policy domain. It is a reminder of the importance of examining a more diverse set of policy areas in political behavior research. Examining a policy area that creates divided preferences allows for an exploration of retrospective voting dynamics when preferences are divided. Some retrospective voting research have examined policy that divided citizens' preferences, for example Erikson & Stoker's (2011) study of the the Vietnam draft lottery. However, this study examines a case of retrospective voting when costs are imposed and preferences divided based on location. This pattern mirrors other types of beneficial but locally opposed facilities, such as housing projects, airports and waste incinerators. If this study is replicated with other kinds of divisive facilities or public policies with spatially concentrated impacts, close attention to treatment assignment and potential selection bias in host communities will be essential.

When the government responds to locally-bounded opposition with policy changes, despite majority support, this can create a democratic accountability failure. Further norma-

[19] We do not know how many people in other political districts (i.e. urban districts) voted for the Liberal Party because they passed the climate policy; I have no way to measure that effect with this approach. Although the hypothesis is that only nearby communities are likely to change their vote because of this policy.

tive concerns are raised because vocal opponents to local divisive projects tend to be white, older, more educated, wealthier and more likely to vote (Mansfield et al., 2001; Walsh et al., 1997). Since these individuals will be more likely to mobilize to oppose projects, this raises questions about the local community's preferences. Communities with higher wealth and social capital, and fewer minorities, are also more likely to oppose divisive projects (Mohai et al., 2009; Aldrich, 2008). If projects are simply moved to poorer or minority communities when local resistance occurs, there are clear equity concerns with accountability to local opponents. More broadly, future work could examine potential failures in democratic accountability when small scale but spatially concentrated electoral groups are mobilized in support of public policy.

Addressing climate change will require sustained political support on the part of both politicians and the public (Schenk and Stokes, 2013). Climate policy will impose costs on the public, whether through unwanted renewable energy in Ontario citizens' backyards, or job losses in West Virginian coal communities. The challenge will be maintaining stable climate policy in the face of intense opposition from local opponents, even while a diffuse majority supports the policy. As renewable energy facilities are sited in communities' backyards around the world, opposition to specific projects is likely. Similar to locating other kinds of divisive facilities, policymakers may need to engage citizens more during project development, to build trust, address concerns about fairness, and if necessary, require revenue sharing provisions to build support (Rabe, 1994; Warren and McFadyen, 2010; Boudet and Ortolano, 2010). Future research could empirically examine variation in policy design, to see whether the way the government structures the process or compensates the community quells local opposition to divisive facilities. Experimenting with policy designs that lead to greater political acceptance of renewable energy will be crucial to minimizing political barriers to effectively addressing climate change.

References

Achen, Christopher H. and Larry M. Bartels (2012). Blind Retrospection: Why Shark Attacks Are Bad For Democracy. Unpublished Manuscript.

Aklin, Michaël and Johannes Urpelainen (2013). Political Competition, Path Dependence, and the Strategy of Sustainable Energy Transitions. *American Journal of Political Science 57*(3), 643–658.

Aldrich, Daniel P. (2008). *Site fights: Divisive facilities and civil society in Japan and the West.* Ithaca: Cornell University Press.

Aldy, Joseph E., Matthew J. Kotchen, and Anthony A. Leiserowitz (2012). Willingness to pay and political support for a US national clean energy standard. *Nature Climate Change 2*(8), 596–599.

Anderson, Christopher J. (2007). The End of Economic Voting? Contingency Dilemmas and the Limits of Democratic Accountability. *Annual Review of Political Science 10*(1), 271–296.

Ansolabehere, Stephen and David M. Konisky (2009). Public Attitudes Toward Construction of New Power Plants. *Public Opinion Quarterly 73*(3), 566–577.

Ansolabehere, Stephen and David M. Konisky (2014). *Clean and Cheap: How Americans Think About Energy in the Age of Global Warming.* Cambridge: MIT Press.

Arnold, R. Douglas (1992). *The Logic of Congressional Action.* New Haven: Yale University Press.

Bartels, Larry M. (1996). Uninformed votes: Information effects in presidential elections. *American Journal of Political Science 40*(1), 194–230.

Baxter, Jamie, Rakhee Morzaria, and Rachel Hirsch (2013). A case-control study of support/opposition to wind turbines: Perceptions of health risk, economic benefits, and community conflict. *Energy Policy 61*, 931–943.

Bechtel, Michael M. and Jens Hainmueller (2011). How Lasting Is Voter Gratitude? An Analysis of the Short- and Long-Term Electoral Returns to Beneficial Policy. *American Journal of Political Science 55*(4), 852–868.

Bechtel, Michael M. and Kenneth F. Scheve (2013). Mass support for global climate agreements depends on institutional design. *PNAS 110*(34), 13763–8.

Bolin, Karl, Gösta Bluhm, Gabriella Eriksson, and Mats E. Nilsson (2011). Infrasound and low frequency noise from wind turbines: exposure and health effects. *Environmental Research Letters 6*(3), 035103.

Boudet, Hilary Schaffer and Leonard Ortolano (2010). A Tale of Two Sitings: Contentious Politics in Liquefied Natural Gas Facility Siting in California. *Journal of Planning Education and Research 30*(1), 5–21.

Caldeira, Ken, Atul K. Jain, and Martin I. Hoffert (2003). Climate sensitivity uncertainty and the need for energy without CO2 emission. *Science 299*(5615), 2052–2054.

Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes (1960). *The American Voter.* Chicago: University of Chicago Press.

Carley, Sanya (2009). Distributed generation: An empirical analysis of primary motivators. *Energy Policy 37*(5), 1648–1659.

Clingermayer, James C. (1994). Electoral Representation, Zoning Politics, and the Exclusion of Group Homes. *Political Research Quarterly 47*(4), 969–984.

Converse, Philip E. (1964). The Nature of Belief Systems in Mass Publics. In D. E. Apter (Ed.), *Ideology and Discontent*, pp. 206–61. New York: Free Press.

Cragg, Michael I., Yuyu Zhou, Kevin Gurney, and Matthew E. Kahn (2013). Carbon Geography: the Political Economy of Congressional Support for Legislation Intended To Mitigate Greenhouse Gas Production. *Economic Inquiry 51*(2), 1640–1650.

Crichton, Fiona, George Dodd, Gian Schmid, Greg Gamble, and Keith J Petrie (2014). Can expectations produce symptoms from infrasound associated with wind turbines? *Health psychology 33*(4), 360–4.

Devine-Wright, Patrick (2005). Beyond NIMBYism: towards an integrated framework for understanding public perceptions of wind energy. *Wind Energy 8*(2), 125–139.

Erikson, Robert S. and Laura Stoker (2011). Caught in the Draft: The Effects of Vietnam Draft Lottery Status on Political Attitudes. *American Political Science Review 105*(02), 221–237.

Erikson, Robert S., Gerald C. Wright, and John P. McIver (1993). *Statehouse democracy: Public opinion and policy in the American states.* Cambridge: Cambridge University Press.

Fiorina, Morris P. (1981). *Retrospective Voting in American National Elections.* New Haven: Yale University Press.

Hainmueller, Jens (2012). Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies. *Political Analysis 20*(1), 25–46.

Healy, Andrew and Neil Malhotra (2013). Retrospective Voting Reconsidered. *Annual Review of Political Science 16*(1), 285–306.

Hill, Stephen D. and James D. Knott (2010). Too close for comfort: Social controversies surrounding wind farm noise setback policies in Ontario. *Journal of Renewable Energy Law and Policy 2*, 153–168.

Hirsh, Richard F. (1999). *Power Loss: The Origins of Deregulation and Restructuring in the American Electric Utility System.* Cambridge, MA: MIT Press.

Jacquet, Jeffrey B. (2012). Landowner attitudes toward natural gas and wind farm development in northern Pennsylvania. *Energy Policy 50*, 677–688.

Johnson, Martin, Paul Brace, and Kevin Arceneaux (2005). Public Opinion and Dynamic Representation in the American States: The Case of Environmental Attitudes. *Social Science Quarterly 86*(1), 87–108.

Keele, Luke and Jason Morgan (2013). Stronger Instruments By Design. Unpublished Manuscipt.

Key, V. O. (1966). *The Responsible Electorate: Rationality in Presidential Voting, 19361960.* Cambridge, MA: Harvard University Press.

Krohn, Soren and Steffen Damborg (1999). On Public Attitudes Towards Wind Power. *Renewable energy 16*, 954–960.

Lachapelle, Erick, Christopher P. Borick, and Barry Rabe (2012). Public Attitudes toward Climate Science and Climate Policy in Federal Systems: Canada and the United States Compared. *Review of Policy Research 29*(3), 334–357.

Lenz, Gabriel S. (2012). *Follow the Leader?: How Voters Respond to Politicians' Policies and Performance.* Chicago: University Of Chicago Press.

Lupia, Arthur (1994). Shortcuts Versus Encyclopedias: Information and Voting Behavior in California Insurance Reform Elections. *American Political Science Review 88*(1), 63–76.

Malin, Stephanie (2014). Theres no real choice but to sign: neoliberalization and normalization of hydraulic fracturing on Pennsylvania farmland. *Journal of Environmental Studies and Sciences 4*(1), 17–27.

Mansfield, Carol, George Van Houtven, and Joel Huber (2001). The Efficiency of Political Mechanisms for Siting Nuisance Facilities: Are Opponents More Likely to Participate than Supporters? *The Journal of Real Estate Finance and Economics 22*(2/3), 141–161.

McAdam, Doug and Hilary Boudet (2012). *Putting Social Movements in their Place: Explaining Opposition to Energy Projects in the United States, 2000-2005.* Cambridge University Press.

Michaud, Kristy, Juliet E. Carlisle, and Eric R.A.N. Smith (2008). Nimbyism vs. environmentalism in attitudes toward energy development. *Environmental Politics 17*(1), 20–39.

Mitchell, C., D. Bauknecht, and P. M. Connor (2006). Effectiveness through risk reduction: a comparison of the renewable obligation in England and Wales and the feed-in system in Germany. *Energy Policy 34*(3), 297–305.

Mohai, Paul, David Pellow, and J. Timmons Roberts (2009). Environmental Justice. *Annual Review of Environment and Resources 34*(1), 405–430.

Olson, Mancur (1965). *The Logic of Collective Action: Public Goods and The Theory of Groups.* Cambridge: Harvard University Press.

Ostrom, Elinor (2010). Polycentric systems for coping with collective action and global environmental change. *Global Environmental Change 20*(4), 550–557.

Page, Benjamin I. and Robert Y. Shapiro (1992). *The rational public: Fifty years of trends in Americans' policy preferences.* Chicago: The University of Chicago Press.

Rabe, Barry G. (1994). *Beyond NIMBY: Hazardous waste siting in Canada and the United States.* Cambridge: Cambridge University Press.

Schenk, Todd and Leah C. Stokes (2013). The power of collaboration: Engaging all parties in renewable energy infrastructure development. *IEEE Power and Energy Magazine 11*(3), 56–65.

Schively, Carissa (2007). Understanding the NIMBY and LULU Phenomena: Reassessing Our Knowledge Base and Informing Future Research. *Journal of Planning Literature 21*(3), 255–266.

Smith, Eric R.A.N. (2002). *Energy, the environment, and public opinion.* Lanham: Rowman & Littlefield Publishers.

Stern, Nicholas H. (2006). *Stern Review: The Economics of Climate Change.* London: HM Treasury.

Stimson, James A., Michael B. MacKuen, and Robert S. Erikson (1995). Dynamic representation. *American Political Science Review 89*(3), 543–565.

Stokes, Leah C. (2013). The politics of renewable energy policies: The case of feed-in tariffs in Ontario, Canada. *Energy Policy 56*, 490–500.

Swofford, Jeffrey and Michael Slattery (2010). Public attitudes of wind energy in Texas: Local communities in close proximity to wind farms and their effect on decision-making. *Energy Policy 38*(5), 2508–2519.

Urpelainen, Johannes (2012). How do electoral competition and special interests shape the stringency of renewable energy standards? *Environmental Economics and Policy Studies*, 23–34.

van der Horst, Dan (2007). NIMBY or not? Exploring the relevance of location and the politics of voiced opinions in renewable energy siting controversies. *Energy Policy 35*(5), 2705–2714.

Walsh, Edward J., Rex Warland, and D. Clayton Smith (1997). *Don't burn it here: Grassroots challenges to trash incinerators.* University Park: Pennsylvania State University Press.

Warren, Charles R., Carolyn Lumsden, Simone O'Dowd, and Richard V. Birnie (2005). Green On Green: Public perceptions of wind power in Scotland and Ireland. *Journal of Environmental Planning and Management 48*(6), 853–875.

Warren, Charles R. and Malcolm McFadyen (2010). Does community ownership affect public attitudes to wind energy? A case study from south-west Scotland. *Land Use Policy 27*(2), 204–213.

Wilson, James Q. (1980). The Politics of Regulation. In *The Politics of Regulation*, pp. 357–394. New York: Basic Books.

Wlezien, Christopher (1995). The Public as Thermostat: Dynamics of Preferences for Spending. *American Journal of Political Science 39*(4), 981–1000.

Wolsink, Maarten (2000). Wind power and the NIMBY-myth: institutional capacity and the limited signicance of public support. *Renewable energy 21*, 49–64.

Wolsink, Maarten (2007). Wind power implementation: The nature of public attitudes: Equity and fairness instead of backyard motives. *Renewable and Sustainable Energy Reviews 11*(6), 1188–1207.

Yatchew, Adonis and Andy Baziliauskas (2011). Ontario feed-in-tariff programs. *Energy Policy 39*(7), 3885–3893.

Zaller, John (1992). *The Nature and Origins of Mass Opinion.* Cambridge: Cambridge University Press.

Supporting Information (SI)

Qualitative evidence of wind project resistance

In addition to the quantitative analyses described in the main text, significant qualitative work was undertaken as part of this research project. Twenty interviews were conducted with politicians, advocates and opponents of the Ontario FIT policy between 2010 and 2013. In addition, anti-wind group's websites were analyzed to understand the extent of political resistance to wind turbines. People living near wind projects were unhappy with these projects and blamed the provincial Liberal Party's climate policy (The Green Energy and Green Economy Act) for catalyzing these projects in their communities.

As Table 5 demonstrates, every single district with a wind project had an active community group working against wind projects in that district. In many communities, anti-wind groups posted signs, intervened at developer meetings, and organized protests. These activities were cataloged online through anti-wind websites. In addition, these groups actively used websites and social media to reach other communities members and anti-wind groups across the province. Two large umbrella organizations, Wind Concerns Ontario and Ontario Wind Resistance, informally coordinated local community group's efforts by sharing detailed information on the climate policy, posting news and distributing signs. The political signs they produced often referenced the Liberal Party, the Premier and the Green Energy Act by name. In this way, these activist groups were informing citizens about the provincial climate policy.

Unsurprisingly, media coverage of these political activities was also significant and grew alongside the anti-wind movement. Figure 4 shows, the number of newspaper articles in Canada with the three terms "Ontario", "wind energy" and "protest" between January 1 and December 31 in each year. In total, there were 211 articles over this time period suggesting significant media attention to public mobilization around the issue. This Figure also shows that issue grew in prominence over time, reaching a peak in 2010 and 2011.

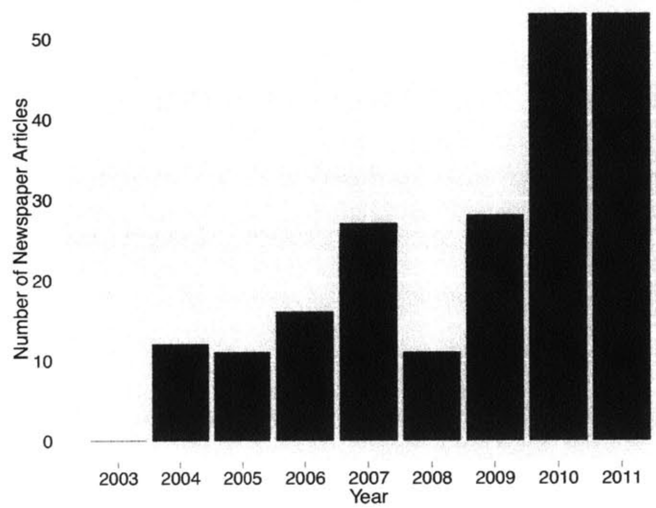

Figure 4: The number of newspaper articles in Canada by year with three terms: "wind energy", "Ontario" and "protest". Data was complied through LexisNexis Academic.

Although citizens were not voting in a referendum on wind turbines, their votes against the government can nevertheless be interpreted as a punishment of the government because of its climate policy. Both of the main rival parties, the Progressive Conservatives and New Democrats, stated they would change the policy if elected. Although the models are not reported here, using the same fixed effects specification as Equation 1 in text, both the Progressive Conservatives and New Democrats received statistically significant increases in vote share in places with proposed and operational turbines. In other words, they were rewarded at the polls for stating they would change the policy if elected.

By contrast, the smaller Green Party, which does not have any seats in the legislature, supported the climate policy. Like the Liberal Party, the same fixed effects models showed they had significant albeit lesser losses in vote share. These results are consistent with the voters near wind projects rewarding the parties that stated they would change the policy and punishing the parties that stated they would keep the policy in place. This evidence further suggests citizens near turbines were responding to the policy.

Table 5: Anti-Wind Groups in each District with a Proposed or Operational Turbine

District	**Number of Groups**	**Names of Groups**
Algoma-Manitoulin	2	Lake Superior Action Research Conservation Manitoulin Coalition for Safe Energy Alternatives
Bruce-Grey-Owen Sound	7	Wind Concerns Bruce Central Bruce Grey Wind Concerns Ontario Bruce Peninsula Wind Turbine Action Group Wind Concerns Meaford West Grey Residents Against Industrial Turbines Preserve Grey Highlands Blue Highlands Citizen Coalition
Chatham-Kent-Essex	1	Chatham Kent Wind Action Group
Dufferin-Caledon	3	Oppose Belwood Wind Farm Grand Valley Group East Garafraxa Coalition
Durham	1	Clarington Wind Concerns
Elgin-Middlesex-London	1	West Elgin Residents Opposing Wind Turbines
Essex	2	Rural Grubby's Blog (Harrow) Citizens Against Lake Erie Wind Turbines
Haldimand-Norfolk	2	Haldimand Wind Concerns Norfolk Victims of Industrial Wind Turbines
Haliburton-Kawartha Lakes-Brock	1	Manvers Wind Concerns
Huron-Bruce	7	Bluewater Against Turbines Saugeen Turbine Operation Policy Ashfield Colborne Wawanosh Against Industrial Turbines Huron East Against Turbines Central Huron Against Turbines Huron-Kinloss Against Lakeside Turbines Ripley Area Victims Group
Kingston and the Islands	1	Wolfe Island Residents for the Environment
Lambton-Kent-Middlesex	3	Lambton County Concerned Citizens Plympton-Wyoming We're Against Wind Turbines Conservation of Rural Enniskillen Middlesex Lambton Wind Action Group
Lanark-Frontenac-Lennox and Addington	3	Beckwith Responsible Wind Action Group Amherst Island Wind Info Protect Amherst Island
Leeds-Grenville	1	South Branch Wind Opposition Group
Niagara West-Glanbrook	2	Wainfleet Wind Action Group West Lincoln/Glanbrook Wind Action Group

Northumberland-Quinte West	2	Alliance for the Protection of Northumberland Hills Northumberland Citizens for Wind Awareness (Port Hope)
Oxford	2	East Oxford Community Alliance East Zorra-Tavistock Township Wind Concerns
Parry Sound-Muskoka	1	Ontarians for Responsible Wind - Georgian Bay chapter
Perth-Wellington	1	West and East Perth Against Turbines
Prince Edward-Hastings	1	CCSAGE
Simcoe-Grey	4	Clearview Wind Issues Clearview Warning About Industrial Turbines Coalition of Residents of Tiny Township Blue Highlands Citizen Coalition
Stormont-Dundas-South Glengarry	1	South Branch Wind Opposition Group
Thunder Bay	1	Nor'Wester Escarpment Protection Committee
Welland	2	Oppose Belwood Wind Farm Stop Mapleton Wind Farm
York-Simcoe	1	Innisfil Windwatchers

Polling precinct boundary shifts

As the paper outlined, there were some boundary changes at the precinct level across all years. Primarily, units were broken up in 2007, and then recombined in 2011. The data set has 6,186 unique precincts in 2011, 5,233 in 2007 and 4,316 in 2003. Largely, this did not affect who was voting for each representative, as the changes were primarily made below the district level. Further, these were administrative changes made by Elections Ontario, an independent election agency.

To create a panel, each 2007 unit was matched with its closest match in 2011 and 2003. The overlap between units was extremely high, as Figure 5 shows. Thus, 2011 and 2003 units were largely broken into several smaller units in 2007, rather than creating entirely new precincts. However, this approach means that some 2003 and 2011 precincts enter the data set more than once because 2007 precincts were often half the size. This is preferable than matching to either the 2011 or 2003 boundaries, because 2007 data would be dropped. The results are the same if all units with replicated entries in either 2003 or 2011 are dropped.

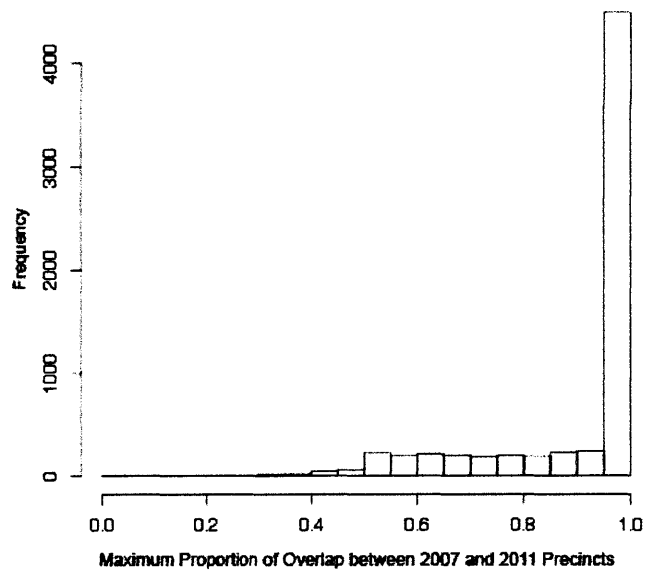

Figure 5: The extent of precinct boundary shifts between 2007 and 2011. For each 2007 precinct, the proportion of overlap with the closest 2011 precinct is shown. Most of the density falls between 0.95 and 1, with a mean of 0.91 and a median of 0.99. This suggests that most units in the panel closely correspond, year to year.

Balance checks for fixed effects models

Although fixed effects estimators create estimates based on within unit variation over time, we nevertheless should be concerned about balance between the treated and control units. In the pre-treatment period, balance on the outcome variable is excellent, with very similar means and variance among the treated and control units (p-value: 0.368, KS test: 0.444, see Figure 6). However balance in the pre-treatment covariates, before reweighting with entropy balancing, is not as good (see Table 6). Substantively the means for many of the variables look similar. However, the large sample size allows detection of even small imbalances between the treated and control units. For example the mean value of a home is only \$6,000 different and median income is less than \$700 different between the treated and control groups, substantively small differences. Still, the balance on population density is more problematic, suggesting urban areas within the rural districts are entering the sample as control units, as is expected. This is the motivation for the rural subset, presented in the main text.

Despite this poor balance on all covariates, when these variables are included as controls in the fixed effects regressions they do not change the results substantially, as Table 1 in the main text clearly shows. In addition, entropy balancing was used to equate the first and second sample moments on population in the precinct (log), population density (log), median income (log), average home value (log), unemployment rate, percentage housing ownership, percentage of people with a university degree, percentage of immigrants and each term squared. After balancing, the means for these variables are equivalent in the treated and control groups. The estimates when these weights are used in the fixed effects models are nearly equivalent as the results from the other models, as Table 1 shows. Together, these results suggest that despite initial imbalances on observable covariates, adjusting for these covariates does not substantially alter the conclusion. Ultimately, confounders must be related to both the treatment and outcome variables to affect the results.

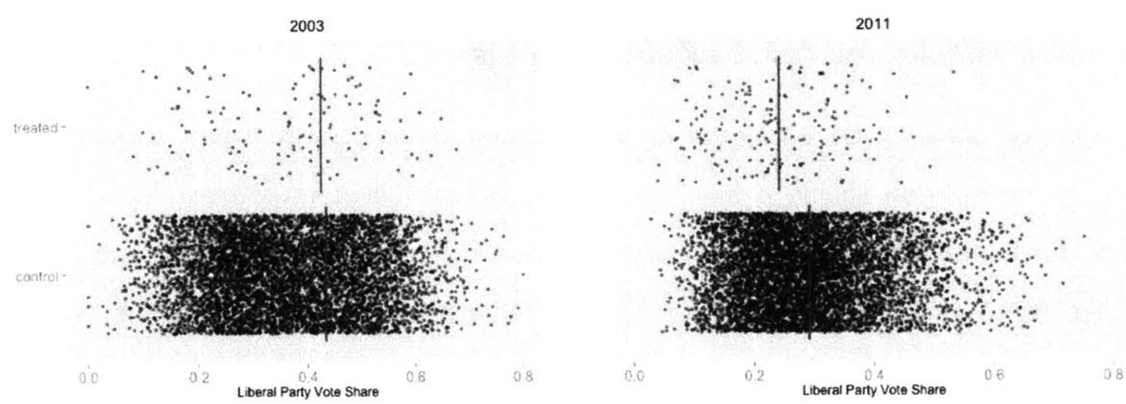

Figure 6: The Liberal Party vote share in each precinct in 2003 and 2011 by treatment status.

Table 6: Balance Table for All Precincts with and without a Wind Turbine Proposal, Pre-treatment (2003)

Variable	Treated N=354	Control N=354	T-test p-value	KS test p-value
Liberal Party vote share (%)	0.424	0.433	0.368	0.444
Population with University Degree (%)	0.093	0.123	0.000	0.000
Median Income (log)	10.020	10.055	0.009	0.006
Average value of home, 2006	159863	153933	0.019	0.000
Population density (log)	2.269	4.782	0.000	0.000

Robustness checks for instrumental variable

The instrumental variable models use a smaller data set with 354 matched pairs. As the main text explains, this is done to reduce model dependence. The subsequent balance across the treated and control groups is excellent, as Table 7 shows.

Table 7: Balance Table for Instrumental Variables Model, Precincts with and without a Wind Turbine Proposal, 2011

Variable	Treated N=354	Control N=354	T-test p-value	KS test p-value
Population with University Degree (%)	0.126	0.126	0.960	0.998
Median Income (log)	10.310	10.310	0.982	0.982
Average value of home, 2006 (log)	12.293	12.292	0.980	1.000
Population density (log)	3.540	3.602	0.642	0.608

Several robustness checks were performed on the instrumental variable models. First, I performed a placebo test on the matched pairs to see whether the instrument, wind power or wind speed to the third power, predicts the Liberal Party vote share in the pre-treatment period (2003). If some confounder exists that drives both the instrument and Liberal Party support, this would be problematic. However we would also expect such a confounder to create an apparent relationship between the instrument and Liberal Party vote share even in the pre-treatment period. I used wind speed to predict future wind treatment in the first stage, and the fitted values to predict 2003 Liberal Party support in the second stage. This placebo test also used the same specification as given in the main text, namely the same flexible geographic controls and the district fixed-effect. The result was insignificant meaning that the instrument does not predict vote share for the Liberal Party pre-treatment. This result provides us some assurance that there is not a confounder driving both wind speed and Liberal Party support.

It is not possible to test the exclusion restriction—the assumption that the effect of the instrument on the outcome variable only occurs through the treatment—directly. Still, plotting the residuals of a regression that uses to same treatment, geographic controls and fixed effects specification to predict the change in the Liberal Party vote share allows us to perform a visual inspection. Ideally, these residuals should be unrelated to the instrument. We can see this visually in Figure 7, which plots these residuals against the instrument. The lowess line is flat, suggesting that after controlling for treatment, geographic factors and fixed effects, there is no residual relationship between the outcome and the instrument. This data set is the same as the one in text, namely it only includes the 354 matched pairs.

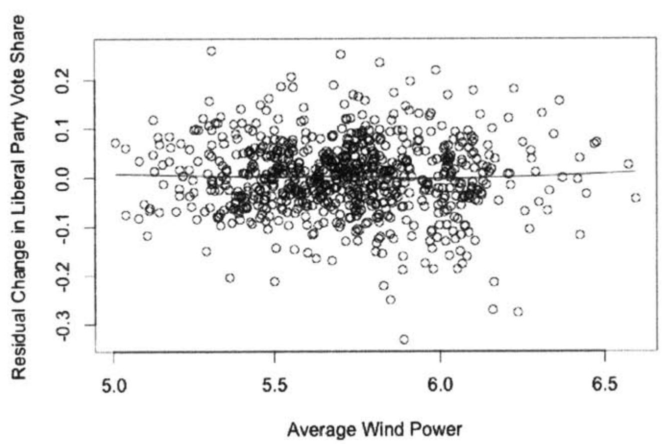

Figure 7: This figure shows that the residuals from predicting the outcome variable (Y) conditional on treatment (D), flexible geographic covariates and district fixed effects (X) are not related to the instrument, average wind power. This is seen by the flat lowess line.