Providing Electricity Increases Support for Data Centers: Evidence from a Survey Experiment

In recent years, there has been considerable debate over artificial intelligence and the data centers that power it. Among those debates is whether these data centers should have to provide their own power to mitigate their electricity usage and downstream effects on residential utility bills. We perform a survey experiment testing whether a proposed data center providing their own electricity can affect citizen support. The 3 X 3 factorial design conditioned on the source of electricity (wind, methane, nuclear) and the amount generated (50% of its usage, 100%, 150%). We find that providing any amount of electricity, compared to a control condition with no electricity provided, increased support for the proposed data center, regardless of the source. Path analysis models further show that the amount of electricity's effect is mediated through perceptions of increasing electricity bills. However, perceptions of air pollution did not mediate the impact of electricity source on project support.

In recent years, there has been considerable debate over artificial intelligence (AI) and the data centers (DCs) that power it. These debates center on several issues but can be categorized into material sources of opposition and symbolic sources of opposition. The former group centers on job creation and displacement (McClain and Park 2026), siting near homes and parks (Zhang et al. 2026), electricity price increases, and pollution, among other concerns. Symbolic opposition instead focuses on intangible outcomes from DCs, such as class conflict between local communities and AI companies and on the disengagement of city councils who site DCs without community consensus.

Attempts to mitigate material opposition have recently centered on requiring DCs to provide their own electricity, either with onsite generation or contributing to the municipal grids' power stocks (Henke 2025; Hytrek 2026; Moore 2026). This represents the "electricity hypothesis" of local DC opposition. Past work on "locally unwanted land uses" and extractive industries' resource usage has shown significant opposition to these types of projects, even when there are no pollutive externalities that result (Freudenburg and Pastor 1992; Popper 1985; Rudel, Roberts, and Carmin 2011; Steelman and Carmin 1998). Even though DCs themselves do not produce air pollution, the installation of natural gas power plants to power them, however, evidences that these externalities can and do result (Hiller 2025; Tabuchi 2026; Taft 2026). Should the electricity hypothesis hold, then DC support in local communities truly does hinge on utility prices and resulting pollution, with other factors only minorly shaping beliefs.

But electricity is likely not the only factor at play. Symbolic opposition may still exert significant influence. In this view, DC opposition represents the actions of local people fighting against AI at large, with DCs being a physical manifestation of this new technology and its increasing dominance over society, politics, and the economy. Illiberal democratic processes, with town councils approving DC development over local opposition, shows a "growth machine" that links DC developers and local governments, prioritizing economic growth and disregarding local community consensus (Molotch 1976). Rather than being purely about material outcomes, symbolic opposition sees DCs as more than their physical structures and local, material impacts. As such, there is nothing that can be done to affect support or opposition based on the design of a specific DC. While we do not test this second hypothesis, it is important to note the other possibilities that may affect and preclude the electricity hypothesis by acting underneath the surface.

Main Text

Data Centers' Electricity Usage: Three Questions

Underlying the electricity hypothesis and the broader DC - energy debate are three related questions. The first is what these facilities consume. The most comprehensive recent accounting estimates that U.S. data centers used approximately 176 TWh in 2023, or 4.4 percent of national electricity consumption, and projects 325 to 580 TWh by 2028, between 6.7 and 12.0 percent of forecast national consumption, with consumption expected to triple between 2024 and 2028 (Shehabi et al. 2024). The width of that range is itself informative: the projected load depends on GPU shipment volumes, operational practices, and cooling efficiency, none of which can be forecast confidently. Load growth of this magnitude, concentrated in specific service territories, is a genuine planning problem regardless of how the costs are ultimately allocated.

The second question is whether that consumption reaches residential ratepayers, and through what mechanism. This is the empirical claim the offset policy is designed to address, and the answer depends heavily on tariff design. Utilities and regulators have responded to large-load interconnection requests with a rapidly evolving set of instruments, minimum demand charges, contract-capacity floors, long-term commitments with exit fees, and marginal-cost pricing, specifically intended to assign costs to the customer causing them and to limit cross-subsidization by other ratepayers (Satchwell et al. 2025). Policies such as "BeYONCE" ("Bring Your Own New Clean Energy") work to offset residential effects by requiring new electric generation (Henke 2025). Whether these instruments succeed is contested and jurisdiction-specific. The important point for our purposes is that the causal chain from a data center's load to a household's bill is mediated by regulatory design and is neither automatic nor uniform.

The third question is how much these data centers affect local environmental quality, especially air pollution (Gour, Ortiz, and Maibach 2026). Recent reports of data centers constructing natural gas power plants have created alarm among activists, politicians, and the general public (Hiller 2025; Tabuchi 2026; Taft 2026). For example, earlier this year, the NAACP sued xAI over its Memphis-based, methane-powered data center, arguing that the pollution from the generation was both unlawful and increasing health risks (Kolodny 2026).

These three questions underlie our study. Regardless of the actual reality of their answers, however, it is important to understand how the public links these three questions' answers to whether a local data center should be cited in their neighborhood. Among Gallup respondents opposing local data centers, 50 percent cite excessive resource use generally, 18 percent name energy consumption specifically, 18 percent name water usage, and 15 percent name higher utility bills (Jones 2026). Political officials seeking to forge a middle ground between data centers' negative externalities and the tax and jobs benefits they promise have sought to implement rules requiring them to provide their own electricity, either to the municipal grid or as a micro-grid of its own (Hytrek 2026; Moore 2026).

We performed a survey experiment to assess the extent to which these requirements can affect public sentiment. We provided 1249 online survey respondents with a vignette asking about their favorability towards a proposed data center. The control condition did not receive more information, but the nine experimental conditions were promised that the data center developer would provide enough electricity to cover either 50%, 100%, or 150% of its usage, sourced from either wind power, natural gas, or nuclear power. This experiment directly targets whether proposals to require providing electricity will be effective at gathering public support. Table 1 shows the vignette conditions and a count in all ten of the condition group.

[Table 1 about here]

Providing Electricity Increases Data Center Support

Our first test is whether providing electricity generation increases support for the proposed data center. Figure 1 shows the average level of support between the two experimental conditions, with the amount of generation on the left and the source of that generation on the right. The top row shows the proportion in support, with black lines to indicate the 95% confidence interval. The bottom row compares the three experimental levels of the two conditions to the control. The point shows the difference in proportions, the curve shows the resampled distribution for the difference, and the black line shows the 95% confidence interval. As is readily evident, providing any amount of electricity generation was shown to increase support compared to the control group (left side). On the right side, we see that, while providing any generation led to support, those who were in the nuclear power condition offered less support.

[Figure 1 about here]

Mediation Effects: Electricity Prices and Air Pollution

We next wanted to test whether these effects of support are mediated through perceptions of electricity prices and air pollution. Figure 2 shows a path model for a structural equation model examining whether the effect of electricity generation is mediated through perceptions of the effect on electricity prices (decrease, no impact, or increase), and whether the effect of the source of generation is mediated through perceptions of the effect on air pollution (decrease, no impact, or increase). The amount generated was converted to an ordinal scale: control (0), 50% (1), 100% (2), 150% (3), while the source was measured as separate, dichotomous categories differentiated from the control. Respondents' overall impression of whether AI is good for society, a 4-point Likert, was included to predict support for building the DC as well as the perceived impacts to electricity prices and air pollution. Included in the model but not shown in the diagram are political party, political ideology, age, education, gender, race, and urbanity.

The model demonstrated excellent fit: $RMSEA = .011\,[90\%\,CI\!: .000 - .026]$, $TLI = .999$, $CFI = .989$. A covariance between perceptions of effects on electricity prices and air pollution was not originally included in the model but was added at the suggestion of modification indices. Further modification indices do not suggest significant changes to the model.

[Figure 2 about here]

On the left side of FIgure 2, it is shown that higher amounts of electricity generation lead to a decreased predicted effect on electricity prices ($b = -.154$, $z = -2.6$, $p = .009$), which in turn has a negative association with data center support ($b = -.199$, $z = -3.47$, $p = .001$). Even further, there is no significant direct effect from amount of generation to support ($b = .074, z = 1.35, p = .178$). The total effect of the amount generated to DC support was .104, and the indirect effect flowing through perceptions of electricity prices was .031; as such, the effect of generation amount was mediated 29.8% through the perceived impact on electricity prices.

Towards the right side of the figure, three boxes denoting the source of the proposed electrical generation (natural gas, wind, or nuclear power) are shown. Notably, none of them had a direct significant effect on support for the proposed data center. Further, while concerns over air pollution are negatively associated with support, there are no significant associations between any of the generation sources and perceived impacts on air pollution. The hypothesized mediation effect between source and air pollution is thus not supported.

Finally, respondents' overall perceptions of AI had significant and strong impacts. Perceptions that AI is generally good for society increased support for the proposed data center ($b = .531$, $z = 10.17$, $p < .001$) and whether it would affect electricity rates ($b = -.368$, $z = -6.29$, $p < .001$) and air pollution ($b = -.354$, $z = -5.3$, $p < .001$).

Discussion and Conclusion

As the data center (DC) debate rages on throughout local communities, it is important to understand how much different facets of opinion in these communities are based on concrete aspects of specific proposals versus symbolic manifestations of global AI politics. We fielded a survey experiment, asking online respondents whether they would be supportive of a data center being built in their community. We randomly assigned each respondent to be given information about whether the DC would provide electricity, how much, and from what source. We find that respondents whose proposed project included electricity generation had higher levels of support than those who did not. This effect was mediated by perceptions of how the DC would affect electricity prices. At the same time, the source of this generation was not a significant predictor of either DC support or whether it would increase air pollution.

This is an extremely relevant topic in the current American political discourse. Backlash to artificial intelligence has become a defining political cause, especially for younger cohorts (McClain and Park 2026). In our study, we targeted one source of opposition, electricity prices, as a component of material objections to DCs. However, symbolic opposition was not studied. AI as a technology does not provide accessible venues for the public to express objections. As such, DC opposition may be representing more generalized critiques of AI and its influence on illiberal democratic processes. DCs have become an important flash point in this conversation, with individuals who feel and foresee that AI has taken their jobs, harmed their communities, and displaced their professional standing in society given a tangible action to stand against it.

Beyond these symbolic aspects, though, many political leaders have adopted the “Bring Your Own New Clean Energy” stance we experimentally tested here, including state leaders in Illinois and Massachusetts and the President of the United States (Colman, Miller, and Long 2026; Healey-Driscoll Administration 2026; Henke 2025). The promise of energy is enough to increase support. Conditioning these projects with energy supply to offset costs can make local individuals more willing to back them. That said, attitudes on this issue are emerging and public perceptions are always evolving. It is very possible that opposition to data centers might grow beyond the point where project conditions can affect public support. This type of study warrants repeated testing to examine when this shift might occur. As for now, it appears that public attitudes are malleable on these new large scale data center projects, when varying conditions of how these proposed centers can operate in a local community.

Figures and Tables

Table 1. Summary of Vignette Conditions

Table 1a. *Number of Respondents*

	Control	Wind	Nuclear	Gas
Control	112			
50%		132	119	111
100%		126	130	129
150%		108	140	142

Table 1b. Conditions

Condition	Vignette Text
Control	Generative Technologies, a leading AI innovations company, is looking to build a state-of-the-art data center in your local community. How much do you support this data center being built?
Experimental	Generative Technologies, a leading AI innovations company, is looking to build a state-of-the-art data center in your local community. To offset its energy use, it will build a [wind farm/nuclear power plant/natural gas plant], which will cover [50%/100%/150%] of its energy electricity usage. How much do you support this data center being built?

Figure 1. Support for Data Centers between Experimental Conditions

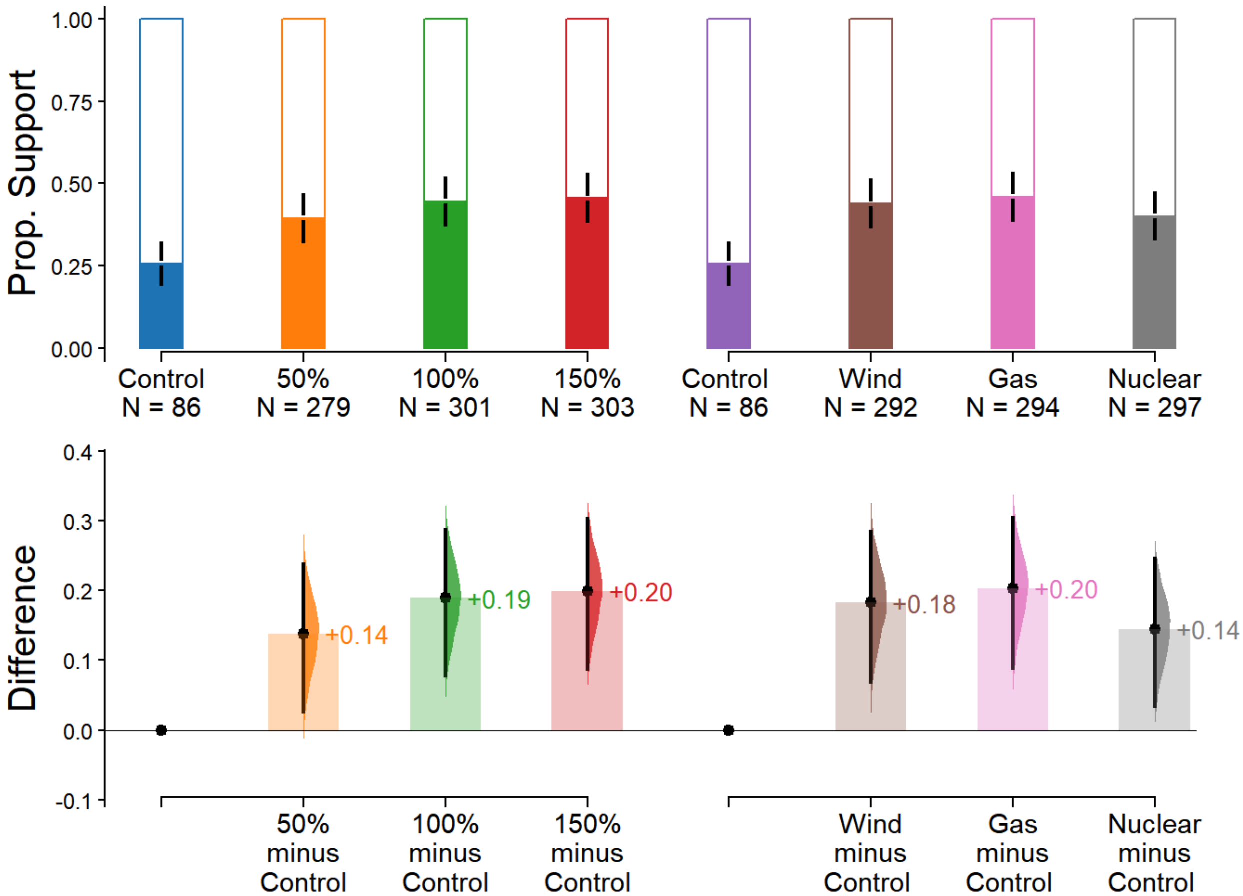

The top half shows proportion in support of the proposed data center across experimental conditions. Black lines represent 95% confidence intervals for the proportion. The bottom half shows the difference between each experimental group to the control group. Distributions reflect 2000 bootstrap comparisons and black lines reflect 95% intervals for those resamples. Figure created using dabestr (Ho et al. 2019; Lu et al. 2026).

Figure 2. Path Model for Testing Experimental Effects of Electricity Generation Amount and Source of Data Center Support

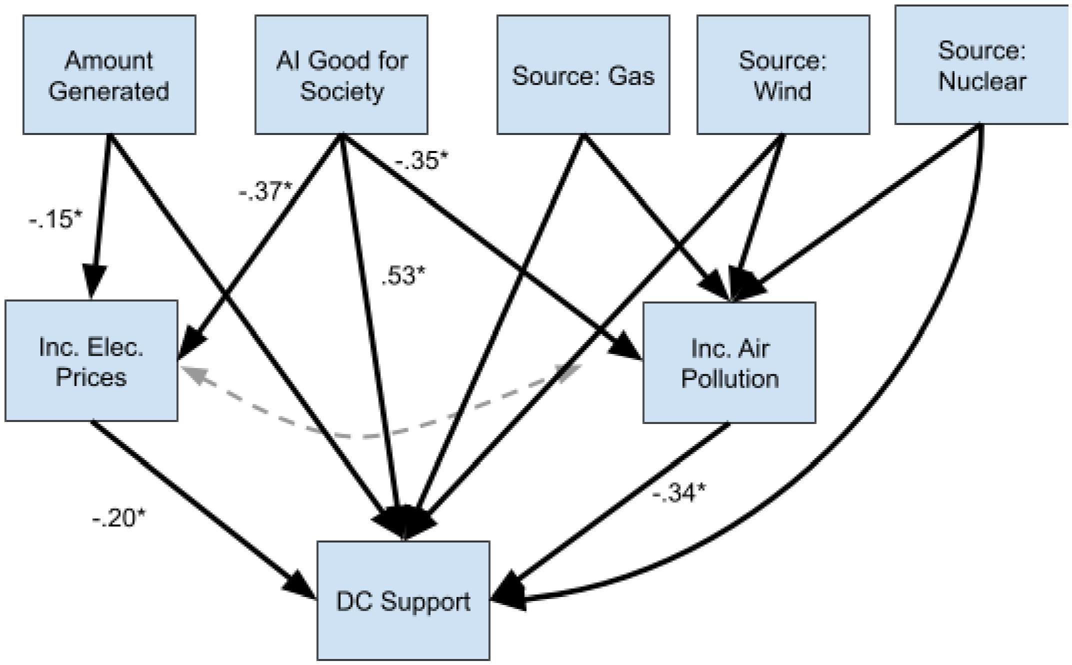

n=1,249. Coefficients with significance level above 0.1 not shown. *p<.01. Included in model to predict DC support but not shown: political affiliation, political ideology, gender, age, race, education, urbanity.
Model fit: $RMSEA = .011$ $[90\%\ CI: .000 - .026]$, $TLI = .999$, $CFI = .989$.
Covariance between perceived impacts: $b = 0.51, z = 12.6, p < .001$.

References

Chandler, Jesse, and Danielle Shapiro. 2016. “Conducting Clinical Research Using Crowdsourced Convenience Samples.” *Annual Review of Clinical Psychology* 12(1): 53–81. doi:10.1146/annurev-clinpsy-021815-093623.

Colman, Zack, Gabby Miller, and Katherine Long. 2026. “Trump and AI Leaders Tout His ‘Build Your Own Power Plant’ Pledge.” *POLITICO*. https://www.politico.com/news/2026/03/04/trump-and-ai-leaders-tout-his-build-your-own-power-plant-pledge-00812891 (August 3, 2026).

Coppock, Alexander, Thomas J. Leeper, and Kevin J. Mullinix. 2018. “Generalizability of Heterogeneous Treatment Effect Estimates across Samples.” *Proceedings of the National Academy of Sciences* 115(49): 12441–46. doi:10.1073/pnas.1808083115.

Freudenburg, William R., and Susan K. Pastor. 1992. “NIMBYs and LULUs: Stalking the Syndromes.” *Journal of Social Issues* 48(4): 39–61. doi:10.1111/j.1540-4560.1992.tb01944.x.

Gour, Neha, Luis Ortiz, and Edward Maibach. 2026. “Health Implications of the Rapid Rise of Data Centers in Virginia: An Exploratory Assessment.” *Frontiers in Climate* 8. doi:10.3389/fclim.2026.1648912.

Healey-Driscoll Administration. 2026. “Healey-Driscoll Administration Statement of Expectations for Responsible Data Center Development and Operations in Massachusetts.” https://www.mass.gov/info-details/healey-driscoll-administration-statement-of-expectations-for-responsible-data-center-development-and-operations-in-massachusetts.

Henke, Cole. 2025. “‘Bring Your Own New Clean Energy’: A Proposed Fix for Illinois’ Power Grid Shortcomings Draws Criticism from Manufacturers.” *WCIA.com*. https://www.wcia.com/news/bring-your-own-new-clean-energy-a-proposed-fix-for-illinois-power-grid-shortcomings-draws-criticism-from-manufacturers/ (September 7, 2026).

Hiller, Jennifer. 2025. “AI Data Centers, Desperate for Electricity, Are Building Their Own Power Plants.” *Wall Street Journal*. https://www.wsj.com/business/energy-oil/ai-data-centers-desperate-for-electricity-are-building-their-own-power-plants-291f5c81 (September 1, 2026).

Ho, Joses, Tayfun Tumkaya, Sameer Aryal, Hyungwon Choi, and Adam Claridge-Chang. 2019. “Moving beyond P Values: Data Analysis with Estimation Graphics.” *Nature Methods* 16(7): 565–66. doi:10.1038/s41592-019-0470-3.

Hytrek, Nikoel. 2026. “Lawmakers Seek to Regulate New Data Centers’ Power, Water Usage.” *Illinois Times*. https://www.illinoistimes.com/news/lawmakers-seek-to-regulate-new-data-centers-power-water-usage/ (February 17, 2026).

Jerit, Jennifer, and Jason Barabas. 2023. “Are Nonprobability Surveys Fit for Purpose?” *Public Opinion Quarterly* 87(3): 816–40. doi:10.1093/poq/nfad037.

Kennedy, Courtney, Nicholas Hatley, Arnold Lau, Andrew Mercer, Scott Keeter, Joshua Ferno, and Dorene Asare-Marfo. 2021. “Strategies for Detecting Insincere Respondents in Online Polling.” *Public Opinion Quarterly* 85(4): 1050–75. doi:10.1093/poq/nfab057.

Kolodny, Lora. 2026. “NAACP Sues Elon Musk’s xAI over Memphis Data Center Air Pollution.” *CNBC*. https://www.cnbc.com/2026/04/14/elon-musk-xai-memphis-data-centers.html (September 2, 2026).

Lu, Zinan, Jonathan Anns, Yishan Mai, Rou Zhang, Kahseng Lian, Nicole MynYi Lee, Shan Hashir, et al. 2026. “Getting over ANOVA: Estimation Graphics for Multi-Group Comparisons.” *Nature Methods*: 1–3. doi:10.1038/s41592-026-03187-7.

McClain, Colleen, and Eugenie Park. 2026. *Young Adults in the U.S. Are Increasingly Wary of AI, Concerned It Will Take Jobs*. Washington, DC: Pew Research Center. https://www.pewresearch.org/short-reads/2026/08/18/young-adults-in-the-us-are-increasingly-wary-of-ai-concerned-it-will-take-jobs/ (September 10, 2026).

Molotch, Harvey. 1976. “The City as a Growth Machine: Toward a Political Economy of Place.” *American Journal of Sociology* 82(2): 309–32. doi:10.1086/226311.

Moore, Daniel. 2026. “Data Centers Won Billions in Tax Breaks. Some States Are Balking.” *Bloomberg Tax*. https://news.bloombergtax.com/daily-tax-report-state/data-centers-won-billions-in-tax-breaks-some-states-are-balking (January 20, 2026).

Pasek, Josh. 2018. “anesrake: ANES Raking Implementation.” https://CRAN.R-project.org/package=anesrake.

Pew Research Center. 2026. “The National Public Opinion Reference Survey (NPORS).” *Pew Research Center*. https://www.pewresearch.org/methods/fact-sheet/national-public-opinion-reference-survey-npors/ (September 8, 2026).

Popper, Frank J. 1985. “The Environmentalist and the LULU.” *Environment* 27(2): 6–11, 37–40. doi:10.1080/00139157.1985.9933448.

Rosseel, Yves. 2012. “lavaan: An R Package for Structural Equation Modeling.” *Journal of Statistical Software* 48(2): 1–36. doi:10.18637/jss.v048.i02.

Rudel, Thomas K., J. Timmons Roberts, and JoAnn Carmin. 2011. “Political Economy of the Environment.” *Annual Review of Sociology* 37(1): 221–38. doi:10.1146/annurev.soc.012809.102639.

Satchwell, Andrew, Natalie Frick, Peter Cappers, Sanem Sergici, Ryan Hledik, Goksin Kavlak, and Glenda Oskar. 2025. *Electricity Rate Designs for Large Loads: Evolving Practices and Opportunities*. Lawrence Berkeley National Laboratory. Technical Brief. doi:10.2172/2502077.

Shehabi, Arman, Sarah Smith, Dale Sartor, Richard Brown, Magnus Herrlin, Jonathan Koomey, Eric Masanet, et al. 2024. *2024 United States Data Center Energy Usage Report*. Berkeley, CA: Lawrence Berkeley National Laboratory. doi:10.2172/1372902.

Steelman, Toddi A., and Joann Carmin. 1998. “Common Property, Collective Interests, and Community Opposition to Locally Unwanted Land Uses.” *Society & Natural Resources* 11(5): 485–504. doi:10.1080/08941929809381096.

Tabuchi, Hiroko. 2026. “Trump’s Vision for A.I. Dominance Comes With Major Air Pollution.” *The New York Times*. https://www.nytimes.com/2026/08/05/climate/data-centers-pollution-trump-ai-energy.html (August 11, 2026).

Taft, Molly. 2026. “Data Centers Are Driving a US Gas Boom.” *Wired*. https://www.wired.com/story/data-centers-are-driving-a-us-gas-boom/ (February 7, 2026).

Traylor, Frederic. 2025. “The Threat of AI Chatbot Responses to Crowdsourced Open-Ended Survey Questions.” *Energy Research & Social Science* 119: 103857. doi:10.1016/j.erss.2024.103857.

Traylor, Frederic, Steven Brechin, and Rachael Shwom. 2025. “Comparing the Effects of Information on Geoengineering Options Support Across Four Online Samples.” *Climatic Change* 178(12): 240. doi:10.1007/s10584-025-04029-9.

Westwood, Sean J. 2025. “The Potential Existential Threat of Large Language Models to Online Survey Research.” *Proceedings of the National Academy of Sciences* 122(47): e2518075122. doi:10.1073/pnas.2518075122.

Zhang, Fengxiu, Shuping Wang, Saba Siddiki, David W. S. Wong, and Jie Xu. 2026. “Balancing Growth and Impact: A Conjoint Analysis of Public Preferences for Data Center Policy.” *Social Science Quarterly* 107(5): e70196. doi:10.1111/ssqu.70196.

Methods

Data

Data for this study come from Rep Data, an online non-probability-based sample of Americans. Non-probability-samples are frequently used for experimentally testing new hypotheses; the low cost and ease and speed of data collection make it possible to adjust parameters when necessary (Chandler and Shapiro 2016). Despite this utility, there are valid concerns for being non-representative of the general public and for having illegitimate, ingenuine, and illogical responses to questions (Kennedy et al. 2021; Traylor 2025; Traylor, Brechin, and Shwom 2025). Recent studies also suggest they may face infiltration by respondents using LLMs to answer questions (Traylor 2025; Westwood 2025). However, in experimental research, they remain useful for piloting theories and finding trends that should exist among the public, even if noisy (Coppock, Leeper, and Mullinix 2018; Jerit and Barabas 2023). In this study, we therefore use this sample to test the popular assumption contained in our hypothesis (that data centter concerns are a result of electricity usage) and urge future work to develop and replicate the findings.

To bolster the validity of our data, we use four techniques. First, we used quota sampling to align the sample with basic US demographics.[1] Additionally, RepData screen responses for logical coherence, straightlining, speeding, and open-ended response patterns, duplicate IP addresses, ReCAPTCHA, and a proprietary screener for AI LLMs. Finally, we used raking to weight the data to align with American Community Survey national estimates. Data were weighted to the 2024 American Communities Survey (5 year estimates) on sex, race, Hispanic ethnicity, age, geographic region, education and interactions of race* sex, race * age, race * education, sex * education, and sex * age. They were also weighted on political party affiliation using the Pew Research Center's (2026) National Public Opinion Reference Survey. Weighing was performed using the *anesrake* package in R (Pasek 2018). Weighting led to a design effect of 1.8. All of these techniques improve our data quality and, paired with the experimental design of our study, should provide reliable and minimally biased estimates, even with nonprobability sampling.

Vignette Experiment Conditions

In this study, we perform a 3 x 3 factorial study, testing whether approval for a data center can hinge on whether it produces enough electricity to cover its usage and the source of that electricity. Respondents to the study were given a brief vignette about

[1] Gender: 48% male, 52% female or other. Race: 70% white, 14% Black, 16% other. Ethnicity: 18% Hispanic (any race), 82% non-Hispanic. Age: 33% 18-34, 30% 35-54, 37% 55 or older. Urban: 20% rural, 80% urban or suburban.

a potential data center being sited in their town. The control group received the following text:

> Generative Technologies, a leading AI innovations company, is looking to build a state-of-the-art data center in your local community.
> How much do you support this data center being built?

The experimental groups received an additional sentence saying that the company would provide electricity for its usage inserted in the middle:

> Generative Technologies, a leading AI innovations company, is looking to build a state-of-the-art data center in your local community. To offset its energy use, it will build a [wind far/nuclear power plant/natural gas plant], which will cover [50%/100%/150%] of its energy electricity usage.
> How much do you support this data center being built?

As such, there are ten total conditions: one control (no electricity information), three options for electricity amount (50%, 100%, 150%), and three options for electricity source (nuclear, natural gas, wind). Sample screenshots of the survey question, displaying different conditions, are shown in the online supplement.

<u>Outcomes and Variables</u>

Following the question, respondents were asked how favorable or unfavorable they were toward the construction of the data center. Options ranged from “dislike a great deal” (=-2) to “like a great deal” (=2). Respondents were then asked about whether the data center would increase (=1), decrease (=-1), or have no impact on (=0) their electric bills and on air pollution in their community.

At a later point in the survey, respondents were asked whether they generally believe artificial intelligence to be “very good for society” (=4) to “very bad for society” (=1). Demographic variables included in the model include political party, ideology, sex, race, education, age, urbanity, and geographic region. Coding for them are shown in the descriptive statistics included in the online supplement.

<u>Analysis</u>

Analysis proceeded in two steps. First, we performed pairwise comparisons, testing the effects of each experimental condition against the relevant control. The mean difference between these conditions was calculated using data analysis with bootstrap-coupled estimation (DABEST; Ho et al. 2019; Lu et al. 2026). Second, we test for whether there are differences in the perceptions of electricity use and air pollution as a result of the data center's associated electricity production. We ran a structural equation model that included these two factors as mediating the effects of electricity source and amount. Modeling was performed using lavaan in R (Rosseel 2012).