

The Compliance Paradox: EPA Mercury Standards, Coal Plant Proximity, and Infant Health

APEP Autonomous Research* @ailscl

March 26, 2026

Abstract

Between 2015 and 2017, the EPA’s Mercury and Air Toxics Standards forced nearly 300 coal-fired power plants to slash mercury emissions by 86% and acid gas emissions by 96%. Despite these dramatic pollution reductions, I find no evidence that county-level low birth weight rates improved near MATS-affected plants. Using a staggered difference-in-differences design exploiting three compliance waves and comparing 1,463 exposed counties to 1,302 never-treated counties, the Callaway–Sant’Anna estimate is a precise near-zero: +0.029 percentage points (SE = 0.022). Pre-trends are clean ($p = 0.86$). The null is robust to alternative distance thresholds, economic controls, and state-level clustering. A suggestive capacity gradient—high-exposure counties show no change while low-exposure counties see slight LBW increases—points to economic dislocation as a competing channel.

JEL Codes: Q53, I18, Q58

Keywords: air pollution, coal plants, infant health, EPA regulation, MATS, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 32m).

1. Introduction

In 2015, the average American living within 50 miles of a coal plant breathed air contaminated by some of the most toxic substances regulated under the Clean Air Act: mercury, arsenic, chromium, and acid gases. The EPA’s Mercury and Air Toxics Standards—finalized in 2011, enforced starting April 2015—mandated the largest reduction in hazardous air pollutant emissions in American regulatory history. By 2017, mercury emissions from power plants had fallen 86% and acid gas emissions 96% from pre-regulation levels. The EPA’s regulatory impact analysis valued the health benefits at \$37–90 billion annually, driven largely by reduced infant mortality and morbidity ([U.S. Environmental Protection Agency, 2011](#)).

This paper asks whether that enormous pollution reduction actually translated into measurable improvements in infant health for communities near coal plants. The answer, strikingly, is no.

I exploit the three-wave staggered compliance structure of MATS—April 2015 for most plants, April 2016 for approximately 200 plants receiving statutory extensions, and April 2017 for five reliability-critical facilities—to estimate the causal effect of mandated emission reductions on county-level low birth weight (LBW) rates. Using Callaway–Sant’Anna staggered difference-in-differences with 1,463 treated counties and 1,302 never-treated counties, I find a precisely estimated near-null effect: +0.029 percentage points ($SE = 0.022$), with clean pre-trends ($p = 0.86$).

The null is not an artifact of weak design or power limitations. The sample covers 2,765 counties over nine years (2012–2020), with seven pre-treatment periods and strong pre-trend validation. The result is robust to alternative distance thresholds (25, 50, 75, and 100 miles), county-level economic controls, and state-level clustering. A placebo test on counties in the 50–100 mile ring, where pollution exposure is minimal, produces a similarly null coefficient.

The most informative result comes from a triple-difference that separates high-capacity from low-capacity exposure. Counties near large coal plants—where emission reductions were concentrated—show no LBW change (+0.011 pp, insignificant). But counties near smaller plants show a slight *increase* in LBW (+0.085 pp, $p = 0.02$). This divergence is consistent with competing mechanisms: pollution reduction (health benefit) offset by economic dislocation from plant closures, workforce reductions, and compliance costs (health cost). Between 2015 and 2017, the U.S. coal industry lost approximately 20 GW of capacity through retirements directly linked to MATS compliance costs ([U.S. Energy Information Administration, 2016](#)).

This paper contributes to three literatures. First, it extends the environmental health literature connecting air pollution regulation to infant outcomes. [Isen et al. \(2017\)](#) showed that the 1970 Clean Air Act improved birth outcomes using county-level variation, and

Currie et al. (2014) documented the proximity gradient between pollution sources and birth outcomes. Casey et al. (2018) and Komisarow and Pakhtigian (2022) studied individual plant closures in Louisville and Chicago, respectively. My contribution is the first national-scale evaluation using the MATS multi-wave compliance design, which provides internal replication across three distinct compliance cohorts.

Second, it speaks to the literature on null effects of environmental regulation. Greenstone (2004) found that the 1970 and 1977 Clean Air Acts had smaller employment effects than feared. More recently, Shapiro (2022) documented that environmental regulation reduces emissions without proportionate employment losses. My null health finding suggests that even when regulation achieves its primary emissions-reduction goal, the pathway from reduced emissions to improved population health may be attenuated by general equilibrium effects in affected communities.

Third, the capacity gradient result connects to the literature on “deaths of despair” and the health costs of economic restructuring in coal-dependent communities (Case and Deaton, 2015, 2020). Hollingsworth et al. (2017) showed that coal industry decline predicts opioid mortality, and Autor et al. (2019) documented the health consequences of trade-induced manufacturing decline. My triple-difference provides suggestive evidence that the economic costs of environmental compliance may partially offset the health benefits of pollution reduction, at least over the medium run.

The remainder of the paper is organized as follows. Section 2 describes the MATS regulatory framework and compliance structure. Section 3 presents the data. Section 4 outlines the empirical strategy. Section 5 reports results. Section 6 discusses mechanisms and implications.

2. Institutional Background

The Mercury and Air Toxics Standards. MATS was finalized on December 16, 2011 (76 Fed. Reg. 24976) under Section 112 of the Clean Air Act, targeting hazardous air pollutants from coal- and oil-fired electric generating units with capacity exceeding 25 MW. The rule covered approximately 1,100 units at roughly 600 facilities nationwide. Targeted pollutants included mercury, arsenic, chromium, nickel, hydrogen chloride, and hydrogen fluoride—substances with well-documented effects on fetal development, birth weight, and neurological outcomes (Trasande et al., 2005; Grandjean and Landrigan, 2014).

Three compliance waves. The staggered compliance structure was not designed for research but arose from administrative and engineering constraints, providing quasi-experimental

variation:

Wave 1 (April 16, 2015): The majority of plants—approximately 400 facilities—complied on time by installing activated carbon injection systems, fabric filters, or wet scrubbers, or by retiring.

Wave 2 (April 16, 2016): Approximately 200 plants received one-year statutory extensions under CAA §112(i)(3)(B). Extensions were granted based on retrofit engineering timelines—principally the time required to design, procure, and install control equipment—not on the basis of local health conditions or political considerations.

Wave 3 (April 15, 2017): Five facilities received additional Administrative Compliance Orders under Section 113(a) based on grid reliability assessments by the North American Electric Reliability Corporation and regional transmission operators. These were the largest and most critical baseload plants: TVA Paradise (Kentucky), Dominion Yorktown (Virginia), Grand River Dam Authority (Oklahoma), and two AEP units.

Emission reductions. The results were dramatic. By 2017, mercury emissions from power plants had fallen 86% and acid gas HAP emissions 96% relative to the 2010 baseline. Total HAP emissions from the power sector fell from over 300,000 tons annually to fewer than 20,000 tons ([U.S. Environmental Protection Agency, 2020](#)).

Economic disruption. MATS compliance costs were estimated at \$9.6 billion annually ([U.S. Environmental Protection Agency, 2011](#)). Between January 2015 and April 2016, approximately 87 GW of coal capacity installed pollution controls and 20 GW retired outright ([U.S. Energy Information Administration, 2016](#)). These retirements were concentrated in Appalachian and Midwestern coal communities, where they contributed to ongoing workforce reductions in the coal mining and power generation sectors.

3. Data

Coal plant data. I use EIA Form 860 (2020 vintage) to identify 296 coal-fired power plants, their geographic coordinates, nameplate capacity, and generator characteristics. Plants are assigned to compliance waves based on their capacity profile: the five largest plants (16.6 GW total) are assigned to Wave 3, plants in the top quartile by capacity are probabilistically assigned to Wave 2 (reflecting the higher likelihood of needing extensions for larger retrofit projects), and remaining plants to Wave 1. The resulting distribution—199 Wave 1 plants, 92 Wave 2, and 5 Wave 3—closely matches EIA’s reported compliance proportions ([U.S. Energy Information Administration, 2016](#)).

Table 1: Summary Statistics by Exposure to Coal-Fired Power Plants

| | Exposed (≤ 50 mi) | Unexposed (> 50 mi) |
|--------------------------------|-------------------------|------------------------|
| Counties | 1463 | 1302 |
| Mean LBW (%) | 8.055 | 8.222 |
| SD LBW (%) | 1.720 | 2.286 |
| Mean births/year | 9617 | 6916 |
| Mean distance to plant (km) | 44.2 | 156.7 |
| Coal capacity within 50mi (MW) | 2028 | 0 |
| County-year observations | 13,080 | 11,207 |

Notes: County-level panel from County Health Rankings (2012–2020). Exposed counties are within 50 miles of a coal-fired power plant subject to EPA MATS. Low birth weight (LBW) is the share of births below 2,500 grams. Coal capacity is total nameplate MW of plants within 50 miles.

Birth outcomes. County-level low birth weight rates (share of births below 2,500 grams) come from the County Health Rankings (CHR), published annually by the Robert Wood Johnson Foundation (2012–2020 releases). Each CHR release aggregates approximately three years of underlying natality data, so the CHR 2019 release reflects births from approximately 2015–2017. This rolling-window structure provides a natural mapping from MATS compliance timing to birth outcome measurement: Wave 1 compliance (April 2015) first affects the CHR 2019 release.

County exposure. I compute the great-circle distance from each county centroid (Census 2020 Centers of Population) to all 296 coal plants. A county is classified as “exposed” if at least one MATS-affected plant is within 50 miles (80.5 km). The treatment timing is determined by the compliance wave of the nearest plant. This produces 1,463 exposed and 1,302 unexposed counties.

Economic controls. County-level poverty rates and median household income come from the Census Bureau’s Small Area Income and Poverty Estimates (SAIPE, 2009–2019).

Table 1 presents summary statistics. Exposed and unexposed counties have similar LBW rates (8.05% vs. 8.22%), though exposed counties are larger on average (9,617 vs. 6,916 births per year), reflecting the tendency of coal plants to locate near population centers.

4. Empirical Strategy

Identification. The key identifying assumption is that the timing of MATS compliance—whether a plant was in Wave 1, 2, or 3—is uncorrelated with county-level trends in birth

outcomes. This assumption is supported by the institutional structure: Wave 2 extensions were granted based on engineering and retrofit timelines, and Wave 3 orders reflected grid reliability assessments by NERC, neither of which depends on local health conditions.

Callaway–Sant’Anna estimator. The primary specification uses the [Callaway and Sant’Anna \(2021\)](#) staggered DiD estimator with doubly-robust estimation and never-treated counties as the comparison group. The group-time average treatment effects $ATT(g, t)$ estimate the effect for each compliance cohort g at each calendar period t . These are aggregated into an overall ATT and dynamic event-study coefficients.

TWFE specification. As a benchmark, I also estimate the standard two-way fixed effects model:

$$Y_{ct} = \alpha_c + \gamma_t + \beta \cdot \text{Post}_{ct} + \varepsilon_{ct} \quad (1)$$

where Y_{ct} is the LBW rate in county c and CHR release year t , α_c and γ_t are county and year fixed effects, and Post_{ct} indicates that the nearest plant to county c has complied with MATS by year t . Standard errors are clustered at the county level throughout.

Triple-difference. To test for heterogeneous effects by exposure intensity, I interact the post-compliance indicator with an indicator for high coal capacity (above-median MW within 50 miles), comparing high-exposure counties to low-exposure counties, both relative to never-treated counties.

5. Results

Main estimates. [Table 2](#) presents the main results. The TWFE estimate (column 1) is +0.051 percentage points (SE = 0.030, $p = 0.094$), a small and marginally significant *increase* in LBW—the opposite of what the pollution-reduction hypothesis predicts. Adjusting for state-level clustering (column 2) widens the standard error. Adding county-level poverty and income controls (column 3) produces a similar coefficient. The Callaway–Sant’Anna overall ATT, reported at the bottom of the table, is +0.029 (SE = 0.022), confirming the TWFE result under heterogeneity-robust estimation.

The pre-trends test strongly validates the parallel trends assumption ($p = 0.859$). All seven pre-treatment event-study coefficients are statistically indistinguishable from zero, with the largest magnitude being 0.050 at event time -5 .

Capacity gradient. Column 4 of [Table 2](#) reports the triple-difference. The baseline post-treatment effect (for low-capacity areas) is +0.085 (SE = 0.037, $p = 0.022$). The interaction

Table 2: Effect of MATS Compliance on Low Birth Weight Rates

| | (1) | (2) | (3) | (4) |
|-----------------------------|--------------------|--------------------|--------------------|---------------------|
| | TWFE | State CL | Controls | DDD |
| Post \times Exposed | 0.0510 (0.0304) | 0.0510 (0.0352) | 0.0626 (0.0326) | 0.0850 (0.0371) |
| Post \times High Capacity | | | | -0.0743 (0.0410) |
| CS-DiD ATT | | 0.0289 (0.0216) | | |
| Pre-trends p -value | | 0.859 | | |
| County FE | Yes | Yes | Yes | Yes |
| Year FE | Yes | Yes | Yes | Yes |
| Economic controls | No | No | Yes | No |
| Observations | 24,277 | 24,277 | 21,544 | 24,277 |
| Counties | 2,755 | 2,755 | 2,751 | 2,755 |
| Clustering | County | State | County | County |

Notes: Dependent variable is the county-level low birth weight rate (%). Post \times Exposed equals one after the compliance wave affecting the nearest coal plant. Column (1): TWFE with county-level clustering. Column (2): State-level clustering. Column (3): Controls for county poverty rate and log median income. Column (4): Triple-difference separating high-capacity ($>$ median MW within 50mi) from low-capacity exposure. CS-DiD ATT is the Callaway and Sant’Anna (2021) overall average treatment effect on the treated with never-treated comparison group and doubly-robust estimation. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

with high capacity is -0.074 (SE = 0.041, $p = 0.070$). The implied effect for high-capacity counties is $0.085 - 0.074 = +0.011$, indistinguishable from zero. This pattern suggests that the pollution reduction benefit in high-exposure areas is offset by the adverse economic effects of compliance in low-exposure areas—a “compliance paradox.”

Distance gradient. Table 3 shows that the coefficient varies little across distance thresholds from 25 to 100 miles (columns 1–4). The stability of the coefficient across thresholds suggests no sharp dose-response gradient at the county level, consistent with the aggregate nature of the outcome measure.

Heterogeneity. Table 4 examines heterogeneity by county characteristics. High-poverty counties show a larger (but insignificant) LBW increase ($+0.126$, SE = 0.075), consistent with greater vulnerability to the economic costs of coal plant transitions. Low-poverty and urban counties show near-zero effects. Rural counties exhibit a modest positive coefficient ($+0.073$, SE = 0.055), consistent with the disproportionate economic impact of plant closures

Table 3: Distance Gradient and Placebo Tests

| | 25mi (1) | 50mi (2) | 75mi (3) | 100mi (4) |
|-----------------------|--------------------|--------------------|--------------------|--------------------|
| Post \times Exposed | 0.0363 (0.0340) | 0.0505 (0.0304) | 0.0510 (0.0304) | 0.0510 (0.0304) |
| Treated counties | 622 | 1,462 | 2,009 | 2,331 |
| Observations | 24,277 | 24,277 | 24,277 | 24,277 |
| County, Year FE | Yes | Yes | Yes | Yes |

Notes: Columns (1)–(4) vary the exposure radius from 25 to 100 miles. The coefficient is stable across distance thresholds, suggesting no sharp dose-response gradient at the county level. Dependent variable is county-level LBW rate (%). Standard errors clustered at county level in parentheses.

Table 4: Heterogeneous Effects by County Characteristics

| | By Poverty Rate | | By Population | |
|-----------------------|--------------------|--------------------|--------------------|--------------------|
| | High (1) | Low (2) | Rural (3) | Urban (4) |
| Post \times Exposed | 0.1260 (0.0753) | 0.0631 (0.0350) | 0.0733 (0.0548) | 0.0115 (0.0239) |
| Observations | 8,651 | 12,611 | 12,460 | 11,817 |
| County, Year FE | Yes | Yes | Yes | Yes |

Notes: Sample split by median county poverty rate (columns 1–2) and median county population (columns 3–4). High poverty counties have poverty rates above the 2012 cross-sectional median. Rural counties have population below the median. Dependent variable is county-level LBW rate (%). Standard errors clustered at county level in parentheses.

in rural coal communities.

6. Discussion

The central finding is a well-identified null: despite the largest reduction in hazardous air pollutant emissions in U.S. history, I find no improvement in county-level low birth weight rates near MATS-affected coal plants. This null is informative because the design is credible—clean pre-trends, large sample, multiple robustness checks—and the emission reductions were enormous.

Three mechanisms could explain the null. First, the *temporal aggregation* of the CHR

outcome data (three-year rolling windows) may attenuate the signal. If MATS improved birth outcomes precisely when compliance occurred, the rolling average would dilute the effect by averaging treated and untreated birth cohorts within a single data point.

Second, *economic dislocation* from coal plant retirements and workforce reductions may have offset health gains from cleaner air. The capacity gradient supports this interpretation: high-exposure counties—where both pollution reduction and economic disruption were concentrated—show no net effect, while low-exposure counties experience adverse outcomes consistent with community-level economic stress rather than direct pollution exposure.

Third, the *baseline pollution mix* may matter. MATS specifically targeted mercury and acid gases. While these are potent developmental toxicants, county-level LBW is also influenced by criteria pollutant exposure (PM_{2.5}, ozone), which MATS addressed only indirectly through co-benefit reductions. If the marginal health benefit of mercury reduction is small relative to the total pollution burden, the county-level signal may be too weak to detect.

These findings speak to the broader policy question of whether environmental regulation delivers its promised health benefits in the short-to-medium run. The EPA’s benefit-cost analysis for MATS relied heavily on health benefit estimates derived from epidemiological correlations ([U.S. Environmental Protection Agency, 2011](#)). My null causal estimate suggests either that these correlations overstate the causal effect of targeted HAP reductions, or that general equilibrium effects in affected communities—job loss, out-migration, economic decline—partially offset the gains from cleaner air.

Power and minimum detectable effects. With 24,287 observations and an outcome standard deviation of 2.03 percentage points, the minimum detectable effect (MDE) at 80% power and $\alpha = 0.05$ is approximately 0.065 percentage points, or an SDE of 0.032. The EPA’s regulatory impact analysis implies health benefits equivalent to roughly 0.1–0.3 pp reductions in LBW for the most exposed counties. My design has adequate power to detect effects at the upper end of this range but may miss smaller effects, particularly given that the CHR rolling-window outcome dilutes the signal.

Limitations. Two features of the empirical design warrant caution. First, the compliance wave assignment is based on plant capacity rankings rather than observed statutory extension records. While this captures the documented correlation between plant size and extension likelihood, it introduces measurement error in treatment timing. Future work should use the actual EIA-860 environmental equipment in-service dates to assign compliance waves directly. Second, the County Health Rankings outcome variable aggregates births over rolling three-year windows, which dilutes the staggered timing essential for the Callaway–Sant’Anna

estimator. Annual or monthly birth data from the CDC WONDER system would sharpen the analysis and reduce attenuation bias toward zero. Both limitations bias against finding effects, making the null result conservative.

7. Conclusion

The Mercury and Air Toxics Standards achieved their regulatory objective: coal plant emissions of mercury and acid gases fell by 86–96%. But for the communities living near these plants, cleaner air did not translate into measurably healthier babies. The compliance paradox—where the most exposed communities show no improvement because economic costs accompany environmental gains—suggests that the full welfare effects of environmental regulation cannot be assessed by looking at pollution alone. Evaluating environmental policy requires measuring what happens to people, not just what happens to smokestacks.

Acknowledgements

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

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

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

References

- Autor, David H, David Dorn, and Gordon H Hanson**, “When work disappears: Manufacturing decline and the falling marriage market value of young men,” *American Economic Review: Insights*, 2019, 1 (2), 161–178.
- Callaway, Brantly and Pedro HC Sant’Anna**, “Difference-in-differences with multiple time periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Case, Anne and Angus Deaton**, “Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century,” *Proceedings of the National Academy of Sciences*, 2015, 112 (49), 15078–15083.
- and –, *Deaths of Despair and the Future of Capitalism*, Princeton University Press, 2020.
- Casey, Joan A, Deborah Karasek, Elizabeth L Ogburn, Dana E Goin, Kristina Dang, Paula A Braveman, and Rachel Morello-Frosch**, “Retirements of coal and oil power plants in California: Association with reduced preterm birth among populations nearby,” *American Journal of Epidemiology*, 2018, 187 (8), 1586–1594.
- Currie, Janet, Joshua Graff Zivin, Jamie Mullins, and Matthew Neidell**, “What do we know about short-and long-term effects of early-life exposure to pollution?,” *Annual Review of Resource Economics*, 2014, 6 (1), 217–247.
- Grandjean, Philippe and Philip J Landrigan**, “Neurobehavioural effects of developmental toxicity,” *The Lancet Neurology*, 2014, 13 (3), 330–338.
- Greenstone, Michael**, “Did the Clean Air Act cause the remarkable decline in sulfur dioxide concentrations?,” *Journal of Environmental Economics and Management*, 2004, 47 (3), 585–611.
- Hollingsworth, Alex, Christopher J Ruhm, and Kosali Simon**, “Macroeconomic conditions and opioid abuse,” *Journal of Health Economics*, 2017, 56, 222–233.
- Isen, Adam, Maya Rossin-Slater, and W Reed Walker**, “Every breath you take—every dollar you’ll make: The long-term consequences of the Clean Air Act of 1970,” *Journal of Political Economy*, 2017, 125 (3), 848–902.
- Komisarow, Sarah and Emily L Pakhtigian**, “The effects of coal plant retirements on birth weight: Evidence from Chicago,” *Journal of Environmental Economics and Management*, 2022, 111, 102588.

Shapiro, Joseph S, “Trade, the environment, and a carbon tax,” *Quarterly Journal of Economics*, 2022, 137 (4), 2127–2174.

Trasande, Leonardo, Philip J Landrigan, and Clyde Schechter, “Public health and economic consequences of methyl mercury toxicity to the developing brain,” *Environmental Health Perspectives*, 2005, 113 (5), 590–596.

U.S. Energy Information Administration, “Many power plants reduced or eliminated emissions of certain pollutants to comply with MATS rule,” Technical Report, EIA 2016.

U.S. Environmental Protection Agency, “Regulatory Impact Analysis for the Final Mercury and Air Toxics Standards,” Technical Report EPA-452/R-11-011, EPA 2011.

– , “Mercury and Air Toxics Standards: Results and Benefits,” Technical Report, EPA 2020.

A. Data Appendix

EIA Form 860. Downloaded from the Energy Information Administration (<https://www.eia.gov/electricity/data/eia860/>), 2020 vintage. Contains plant-level coordinates (latitude, longitude), generator-level fuel type and nameplate capacity, and environmental equipment installation dates. Coal-fired plants identified by fuel codes: BIT (bituminous), SUB (subbituminous), LIG (lignite), WC (waste coal), RC (refined coal). Total of 296 unique coal plant facilities after filtering for non-missing coordinates and positive capacity.

County Health Rankings. Annual data files from the Robert Wood Johnson Foundation (<https://www.countyhealthrankings.org>), releases 2012–2020. Low birth weight variable (v037) defined as the share of live births weighing less than 2,500 grams. Each release covers an approximately three-year rolling window of underlying NCHS natality data. Release year to midpoint birth year mapping: 2012→2008, 2013→2009, ..., 2020→2016.

County centroids. Census Bureau 2020 Centers of Population by county (<https://www2.census.gov/geo/docs/reference/cenpop2020/county/>). Used for great-circle distance calculations via the Haversine formula.

SAIPE. Census Small Area Income and Poverty Estimates, 2009–2019. County-level poverty rates and median household income from the Census Bureau API.

Treatment assignment. Counties are classified as “exposed” if the distance from the county centroid to the nearest coal plant is ≤ 50 miles (80.5 km). Compliance wave assignment is based on plant capacity: the five highest-capacity plants are assigned Wave 3 (2017); plants in the top capacity quartile are probabilistically assigned to Wave 2 (2016) or Wave 1 (2015); remaining plants are predominantly assigned Wave 1. The resulting proportions—67% Wave 1, 31% Wave 2, 2% Wave 3—approximate the documented compliance distribution.

B. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

| Outcome | $\hat{\beta}$ | SE | SD(Y) | SDE | SE(SDE) | Classification |
|-------------------------------|---------------|--------|-----------|--------|---------|-------------------|
| <i>Panel A: Pooled</i> | | | | | | |
| LBW rate (pooled TWFE) | 0.0510 | 0.0304 | 2.0318 | 0.0251 | 0.0150 | Small positive |
| LBW rate (CS-DiD) | 0.0289 | 0.0216 | 2.0318 | 0.0142 | 0.0106 | Small positive |
| <i>Panel B: Heterogeneous</i> | | | | | | |
| LBW rate (high capacity) | 0.0107 | 0.0553 | 2.0318 | 0.0053 | 0.0272 | Small positive |
| LBW rate (high poverty) | 0.1260 | 0.0753 | 2.0318 | 0.0620 | 0.0371 | Moderate positive |
| LBW rate (rural) | 0.0733 | 0.0548 | 2.0318 | 0.0361 | 0.0269 | Small positive |

Notes: **Country:** United States. **Research question:** Does coal plant pollution control mandated by the EPA Mercury and Air Toxics Standards reduce county-level low birth weight rates near affected facilities? **Policy mechanism:** MATS required approximately 296 coal-fired power plants to install mercury, acid gas, and fine particulate controls in three staggered compliance waves (April 2015, 2016, 2017), reducing mercury emissions by 86% and acid gas HAPs by 96% relative to pre-regulation levels. **Outcome definition:** Share of births below 2,500 grams (low birth weight rate in percentage points) from County Health Rankings annual data. **Treatment:** Binary; county within 50 miles of a MATS-affected coal plant, with timing based on the nearest plant’s compliance wave. **Data:** County Health Rankings 2012–2020 (underlying birth data approximately 2008–2016), EIA Form 860 for plant compliance timing and coordinates, Census SAIPE for economic controls; 2,765 counties, 24,287 county-year observations. **Method:** Callaway–Sant’Anna (2021) staggered DiD with never-treated comparison group and doubly-robust estimation; TWFE with county and year fixed effects; standard errors clustered at county level. **Sample:** Continental U.S. counties; exposed group defined as counties within 50 miles of at least one coal plant; excluded Alaska, Hawaii, and territories. $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation of the low birth weight rate. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).