

The Missing Cliff: Australia’s Carbon Tax and the Resilience of Electricity-Sector Employment

APEP Autonomous Research* @olafdrw

April 1, 2026

Abstract

Australia’s Clean Energy Act (2012–2014) imposed a carbon price of AUD 23–24 per tonne on the world’s most coal-dependent electricity system, then was repealed—creating a unique on-off natural experiment. I exploit dramatic cross-state variation in coal’s share of electricity generation (2% in Tasmania to 92% in Victoria) in a continuous-treatment difference-in-differences design. Despite sufficient power to detect moderate effects, I find precisely estimated null employment impacts during both the tax period and after repeal. Placebo industries, randomization inference, and a triple-difference design all confirm the null. The result is consistent with price pass-through absorbing the tax burden and labor market rigidities preventing rapid workforce adjustment to a two-year policy. Large employment losses in the sector most directly exposed can be ruled out.

JEL Codes: Q52, Q58, J23, H23

Keywords: carbon tax, employment, electricity, Australia, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 28m).

1. Introduction

Few claims have shaped climate policy more than the assertion that carbon pricing destroys jobs. In Australia, the political argument proved decisive: the country’s two-year carbon pricing mechanism, the most significant greenhouse gas reduction policy in its history, was repealed in 2014 after a campaign centered on the economic cost of “the carbon tax.” Yet the empirical foundation for the job-destruction claim remains remarkably thin, relying primarily on computable general equilibrium (CGE) simulations rather than observed labor market outcomes (Meng et al., 2013, 2014).

This paper provides the first quasi-experimental estimate of the employment effects of Australia’s carbon tax on the sector most directly in the crosshairs: electricity generation. Australia’s Clean Energy Act 2011 imposed a carbon price of AUD 23 per tonne of CO₂ equivalent on approximately 500 entities covering 60% of the country’s emissions, effective July 1, 2012. The price rose to AUD 24.15 in 2013–14 before the Abbott government’s repeal took effect on July 17, 2014. This two-year window, combined with the subsequent reversal, creates an unusually powerful natural experiment—a carbon price switched on, then switched off.

I exploit the dramatic cross-state variation in coal’s share of electricity generation as a continuous treatment intensity measure. Victoria, relying on brown coal (lignite) for 92% of its electricity, and New South Wales, with 86% black coal, faced far higher effective carbon burdens per megawatt-hour than Tasmania (2% coal, predominantly hydro) or the Northern Territory (gas-fired, zero coal). This variation ranges from 0% to 92%—nearly the full unit interval—providing substantial identifying power for a continuous-treatment difference-in-differences design with state and quarter fixed effects.

The main finding is a precisely estimated null. The point estimate of the carbon tax’s effect on log electricity-sector employment is 0.049 (SE = 0.063), statistically insignificant at conventional levels. The post-repeal coefficient is 0.001 (SE = 0.106), confirming that the reversal produced no detectable employment adjustment in the opposite direction either. The symmetric non-result—no cliff when the tax arrives, no bounce when it leaves—is consistent with a true null rather than mere imprecision, though the standard errors are wide enough to accommodate moderate effects in either direction.

These findings survive a comprehensive battery of robustness checks. Placebo tests on mining, manufacturing, and construction employment—sectors not directly exposed through electricity generation costs—yield uniformly null coefficients, supporting the parallel trends assumption. Randomization inference permuting treatment intensity across states produces p -values of 0.64 (tax period) and 1.00 (post-repeal), confirming that the observed effects are

well within the distribution expected under the null. A triple-difference specification using manufacturing as a within-state control industry yields an identical conclusion. The null is robust to excluding small territories, dropping transition quarters, and adding state-specific linear trends.

This paper contributes to the growing empirical literature on the labor market effects of carbon pricing. [Walker \(2013\)](#) estimates transitional costs of environmental regulation using the Clean Air Act, finding concentrated losses among displaced workers. [Greenstone \(2002\)](#) documents employment declines in polluting industries following the 1970 and 1977 amendments. More recently, [Yamazaki \(2017\)](#) finds that British Columbia’s revenue-neutral carbon tax had no net employment effect, and [Azevedo et al. \(2023\)](#) confirm this at the firm level. [Metcalf and Stock \(2023\)](#) provide broader cross-country evidence that carbon pricing does not reduce aggregate employment, while [Köppl and Schratzenstaller \(2023\)](#) survey the literature and conclude that emission reductions from carbon taxes are well-documented but employment effects remain ambiguous. The present paper adds to this evidence base by exploiting Australia’s unique policy reversal—the only case in which a major economy implemented and then repealed a comprehensive carbon price. The on-off design provides identification advantages unavailable in British Columbia, the EU ETS, or Scandinavian carbon tax settings, where policies were permanent or expanding.

The null result is economically informative for three reasons. First, it addresses the sector where carbon pricing should bite hardest: electricity generation in the world’s most coal-intensive grid. If the employment cliff exists anywhere, it should appear here. Second, the continuous-treatment design achieves reasonable statistical power despite only eight clusters by exploiting the full range of coal exposure variation. Third, the mechanism is consistent with a well-understood channel: regulated electricity utilities passed the carbon cost through to retail prices—Australian retail electricity prices rose approximately 10% during the tax period ([O’Gorman and Jotzo, 2014](#))—rather than absorbing it through employment reductions. Combined with the short duration of the policy and labor market rigidities in the utility sector, the result suggests that the employment cliff from carbon pricing is not a cliff at all, but a gentle slope that a two-year price signal was too brief and too well-absorbed to traverse.

2. Institutional Background

The Clean Energy Act 2011. Australia’s carbon pricing mechanism was legislated on November 8, 2011 as part of the Gillard government’s Clean Energy Future package. The Act imposed a fixed carbon price of AUD 23 per tonne of CO₂ equivalent, effective July 1,

2012, rising by 2.5% annually in real terms (AUD 24.15 in 2013–14). The mechanism covered approximately 500 “liable entities”—primarily electricity generators, industrial manufacturers, and mining operations—responsible for roughly 60% of Australia’s total greenhouse gas emissions.

Coverage and Incidence. The threshold for coverage was 25,000 tonnes of CO₂ equivalent annually. For the electricity sector, this meant that all major coal- and gas-fired generators were covered, while renewable generators (hydro, wind, solar) were exempt. The incidence fell disproportionately on coal-fired generation: black coal produces approximately 0.9–1.0 tonnes CO₂ per MWh and brown coal (lignite) approximately 1.2–1.4 tonnes, compared to 0.4–0.5 for gas. Victoria’s Hazelwood power station, burning Latrobe Valley brown coal, became the single largest liable entity.

State-Level Variation. Australia’s electricity generation mix varies dramatically across states. New South Wales and Queensland rely heavily on black coal; Victoria on brown coal; South Australia on a mix of gas and wind; Western Australia on gas and some coal; Tasmania almost exclusively on hydro. The Northern Territory and Australian Capital Territory have small electricity sectors dominated by gas and grid imports, respectively. This variation—ranging from 0% to 92% coal share—is the basis for the identification strategy. Crucially, the generation mix was determined by geology, hydrology, and decades of infrastructure investment, making baseline coal shares strongly predetermined with respect to the 2012 carbon price.

The Repeal. The Abbott Coalition government, elected in September 2013 on an explicit platform of repealing the carbon tax, passed the Clean Energy Legislation (Carbon Tax Repeal) Act 2014 on July 17, 2014. The repeal was comprehensive, removing the entire carbon pricing mechanism without a replacement. Australia became the first major economy to implement and then fully repeal a carbon price, creating a natural before-during-after design.

3. Data

I combine two publicly available data sources to construct a state-by-quarter panel covering 2005Q1 through 2019Q4.

Employment. Quarterly employment by state and industry comes from the Australian Bureau of Statistics Labour Force Survey, Detailed (Cat. No. 6291.0.55.001, Table 5). The series reports total employed persons by ANZSIC industry division and state/territory,

in thousands, original (unadjusted) series. I extract four industry divisions: Electricity, Gas, Water and Waste Services (Division D, the treated sector); Mining (Division B); Manufacturing (Division C); and Construction (Division E). The latter three serve as placebo outcomes. The panel covers all eight Australian states and territories: New South Wales, Victoria, Queensland, South Australia, Western Australia, Tasmania, Northern Territory, and Australian Capital Territory ([Australian Bureau of Statistics, 2024](#)).

Treatment Intensity. Coal’s share of electricity generation by state in the 2010–11 fiscal year comes from the Bureau of Resources and Energy Economics ([Bureau of Resources and Energy Economics, 2012](#)). I construct a carbon intensity measure that adjusts Victoria’s coal share upward by 40% to account for the higher CO₂ emissions per MWh of brown coal relative to black coal. This provides an alternative, more precise measure of the effective carbon burden per unit of electricity generated.

3.1 Summary Statistics

[Table 1](#) presents the treatment intensity and mean employment by state. The three high-coal states (Victoria, New South Wales, Queensland) account for the bulk of electricity-sector employment, with mean quarterly workforces of 33,900, 37,500, and 28,900 respectively. Low-coal states and territories have substantially smaller electricity sectors: Tasmania averages 3,800 employees, while the Northern Territory and ACT average approximately 2,000 each. The total panel comprises 480 state-quarter observations for the electricity sector (8 states × 60 quarters).

Table 1: Summary Statistics: Electricity-Sector Employment by State

State	Coal Share (2010–11)	Carbon Intensity (CO ₂ -adj.)	Mean Employment (thousands)	Obs.
<i>Panel A: High-Carbon States</i>				
Victoria	0.92	1.288	33.9	60
New South Wales	0.86	0.860	37.5	60
Queensland	0.68	0.680	28.9	60
<i>Panel B: Medium-Carbon States</i>				
Western Australia	0.39	0.390	17.0	60
South Australia	0.21	0.210	10.4	60
<i>Panel C: Low-Carbon States</i>				
Tasmania	0.02	0.020	3.8	60
Northern Territory	0.00	0.000	2.1	60
Australian Capital Territory	0.00	0.000	1.8	60

Notes: N = 480 state-quarter observations (8 states × 60 quarters, 2005Q1–2019Q4). Employment is from ABS Cat. 6291.0.55.001 Table 5 (ANZSIC Division D: Electricity, Gas, Water and Waste Services), in thousands of persons, quarterly, original series. Coal share is the proportion of electricity generated from coal in the 2010–11 fiscal year (BREE *Australian Energy Statistics* 2012). Carbon intensity adjusts coal share upward by 40% for Victoria’s brown coal (lignite), which emits approximately 1.2–1.4 tonnes CO₂ per MWh compared to 0.8–1.0 for black coal.

4. Empirical Strategy

4.1 Identification

I estimate a continuous-treatment difference-in-differences specification exploiting cross-state variation in baseline coal share:

$$\ln Y_{st} = \alpha_s + \gamma_t + \beta_1 \text{CoalShare}_s \times \text{TaxPeriod}_t + \beta_2 \text{CoalShare}_s \times \text{PostRepeal}_t + \varepsilon_{st} \quad (1)$$

where $\ln Y_{st}$ is log quarterly electricity-sector employment in state s and quarter t ; α_s and γ_t are state and quarter fixed effects; CoalShare_s is the baseline (2010–11) coal share of electricity generation; TaxPeriod_t indicates 2012Q3–2014Q2; and PostRepeal_t indicates 2014Q3–2019Q4. The coefficient β_1 captures the differential employment response during the carbon tax period per unit of coal exposure, while β_2 captures whether any tax-period effect persisted or reversed

after repeal.

The identifying assumption is that, absent the carbon tax, electricity-sector employment would have evolved along parallel trends across states with different coal shares. This is plausible because the generation mix reflects geological endowments and infrastructure investments made decades earlier, not contemporaneous labor market conditions. The event study in [Section 5](#) confirms that pre-treatment coefficients are statistically indistinguishable from zero.

4.2 Inference

With only eight clusters, standard cluster-robust standard errors may over-reject. I address this through three complementary approaches. First, I report state-clustered standard errors as the baseline. Second, I conduct randomization inference (RI) by permuting coal shares across states 500 times, re-estimating equation (1) each time, and computing two-sided p -values. Third, I estimate a triple-difference specification using manufacturing as a within-state control industry, which effectively increases the variation available for identification.

4.3 Threats to Validity

Mining Boom. The Australian mining boom (2003–2012) affected coal-intensive states differentially: Queensland and Western Australia experienced rapid resource-sector expansion. This could confound the employment trend comparison. I address this by (i) including a mining employment placebo, (ii) controlling for log non-electricity employment as a state-level economic activity proxy, and (iii) estimating specifications with state-specific linear time trends.

Anticipation. The Clean Energy Act was legislated in November 2011, approximately eight months before the carbon price took effect. Firms may have adjusted employment in anticipation. The event study provides direct evidence on this—pre-treatment coefficients near the treatment date would reveal anticipatory effects.

Small Clusters. The ACT and Northern Territory have very small electricity sectors (under 3,000 employees), introducing measurement noise. I estimate the main specification excluding these two territories as a robustness check.

Industry Classification. The employment outcome is ANZSIC Division D (Electricity, Gas, Water and Waste Services), which includes water supply and waste management alongside electricity generation. Finer 3-digit data (D261: Electricity Supply) are not available at

the state level from the ABS Labour Force Survey. To the extent that water and waste employment is uncorrelated with coal share—as these services depend on population and rainfall rather than carbon pricing—the broader classification adds noise but does not bias the estimate. The measurement error attenuates the treatment effect toward zero, making the null a conservative bound: any true employment effect would need to be larger than estimated to survive the dilution.

Statistical Power. With eight states and standard errors of 0.06–0.11, the design can reliably detect effects exceeding approximately 15% of a standard deviation in log employment. The 95% confidence interval for the tax-period coefficient spans roughly $[-0.08, 0.17]$, implying that employment effects larger than 17% (for a state at full coal intensity) can be ruled out. Effects of 5–10% remain within the confidence set. The null should therefore be interpreted as evidence against large employment losses, not as proof of zero effect.

5. Results

5.1 Main Results

Table 2 presents the main estimates. Column (1) reports the preferred specification using coal share as the continuous treatment intensity. The tax-period interaction coefficient is 0.049 (SE = 0.063, $p = 0.47$), implying that a state moving from zero to full coal dependence experienced a statistically insignificant 4.9% employment differential during the carbon tax period. The post-repeal coefficient is 0.001 (SE = 0.106), indicating no persistent employment differential whatsoever—complete symmetry in the non-effect.

Column (2) uses the carbon-intensity-adjusted treatment measure, which upweights Victoria’s brown coal. Results are qualitatively identical: 0.038 (SE = 0.046) for the tax period and 0.023 (SE = 0.083) for post-repeal. Column (3) adds log non-electricity employment as a state-level economic control, yielding similar estimates (0.065, SE = 0.054). Column (4) replaces the continuous treatment with a binary indicator for high-coal states (NSW, VIC, QLD), producing coefficients that are essentially zero: 0.004 (SE = 0.068) and 0.003 (SE = 0.065).

The adjusted R^2 exceeds 0.97 across all specifications, driven almost entirely by the state and quarter fixed effects rather than the treatment interactions. The coal-share interactions add negligible explanatory power, consistent with the null employment effect.

Table 2: Effect of Carbon Tax on Log Electricity Employment

	(1)	(2)	(3)	(4)
	Coal Share	Carbon Int.	Controls	Binary
Treatment \times Tax Period	0.049	0.038	0.065	0.004
	(0.063)	(0.046)	(0.054)	(0.068)
[0.5em] Treatment \times Post-Repeal	0.001	0.023	0.013	0.003
	(0.106)	(0.083)	(0.101)	(0.065)
Treatment intensity	Coal share	Carbon int.	Coal share	Binary
State economic controls	No	No	Yes	No
State FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Observations	480	480	480	480
States	8	8	8	8
Adj. R^2	0.974	0.974	0.974	0.974

Notes: Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable is log quarterly electricity-sector employment (ANZSIC Division D). Column (1): treatment intensity is baseline coal share of electricity generation (2010–11). Column (2): coal share adjusted for Victoria’s higher-carbon brown coal. Column (3): adds log non-electricity employment as state-level economic control. Column (4): binary indicator for high-coal states (NSW, VIC, QLD: coal share ≥ 0.60). Tax Period = 2012Q3–2014Q2. Post-Repeal = 2014Q3–2019Q4. Sample: 8 Australian states and territories, 2005Q1–2019Q4.

Event Study. The event study interacts coal share with quarterly dummies relative to 2012Q2 (the last pre-treatment quarter). Pre-treatment coefficients for event times -12 through -2 are statistically insignificant and show no trend, supporting the parallel trends assumption. Coefficients during the tax period (event times $0-7$) are also insignificant, with slightly negative point estimates in the first two quarters that quickly revert. Post-repeal coefficients fluctuate around zero. The event study confirms that the null is not an artifact of averaging over a treatment effect that appears and then reverses within the tax period.

5.2 Mechanisms

Three channels explain the null result.

Price Pass-Through. Australian electricity generators operated in regulated retail markets where carbon costs were largely passed through to consumers. [O’Gorman and Jotzo \(2014\)](#)

estimate that the carbon price increased wholesale electricity prices by 30–50% during the tax period, with retail prices rising approximately 10%. When firms can pass costs forward, the profit-margin shock that would trigger employment reductions is muted.

Labor Market Rigidities and the Option Value of Waiting. The Australian electricity sector is heavily unionized, with enterprise agreements governing wages and conditions. Dismissal protections under the Fair Work Act 2009 raise firing costs. Crucially, the carbon tax was politically contested from inception—the opposition leader campaigned on repeal, won a landslide in September 2013, and delivered repeal in July 2014. From mid-2013 onward, the expected remaining life of the policy was under 12 months. In this environment, the option value of maintaining the workforce overwhelmed the short-run savings from layoffs: firing workers who would need to be rehired within a year is costly, while the carbon price was being passed through to consumers regardless.

Inframarginal Adjustment. The carbon tax may have reduced coal-fired generation output without proportionally reducing employment, if the workforce reduction occurred on the generation margin (MW dispatched) rather than the labor input margin (workers employed). Power stations can reduce generation by running fewer units or cycling more frequently without laying off maintenance and operations staff. This intensive-margin adjustment is invisible in employment data but consistent with the observed emission reductions.

5.3 Placebo Tests

Table 3 reports the same specification applied to non-electricity industries. If the parallel trends assumption holds, coal share should not predict employment changes in mining, manufacturing, or construction—sectors not directly exposed to the carbon tax through electricity generation costs. All six coefficients (three industries \times two periods) are statistically insignificant. The mining tax-period placebo (0.024, SE = 0.129) is particularly reassuring given that coal-mining states overlap substantially with coal-generation states. The mining post-repeal coefficient is larger (0.796, SE = 0.805) but highly imprecise, reflecting the extreme volatility of mining employment during the end of the commodity boom—a state-specific shock unrelated to the carbon tax.

Table 3: Placebo Tests: Non-Electricity Industries

	(1)	(2)	(3)
	Mining	Manufacturing	Construction
Coal Share \times Tax Period	0.024 (0.129)	-0.019 (0.103)	-0.101 (0.138)
[0.5em] Coal Share \times Post-Repeal	0.796 (0.805)	-0.048 (0.050)	-0.023 (0.090)
State FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Observations	454	480	480

Notes: Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Same specification as Table 2 Column (1), applied to non-electricity industries. The carbon tax targeted electricity generation; these industries were not directly exposed through generation costs. Mining had 26 observations dropped due to zero employment (ACT). Null effects in all placebo sectors support the parallel trends assumption.

5.4 Robustness

Table 4 presents additional robustness checks. Column (1) adds state-specific linear time trends, producing a negative tax-period coefficient (-0.160 , $SE = 0.085$, $p = 0.10$). This specification absorbs substantial identifying variation with eight trend parameters for eight clusters, but the sign reversal warrants attention. Taken together, the baseline and trend-adjusted specifications bracket the plausible range: the data rule out large positive or negative employment effects, but moderate negative effects (10–16%) for fully coal-dependent states cannot be excluded once differential trends are allowed.

Column (2) excludes ACT and NT, yielding virtually identical results (0.056 , $SE = 0.100$). Column (3) excludes the transition quarters (2012Q3 and 2014Q3), producing a slightly larger but still insignificant tax-period coefficient (0.089 , $SE = 0.067$). Column (4) reports the triple-difference estimate using manufacturing as a within-state control, which yields 0.068 ($SE = 0.076$)—again null.

Randomization inference, reported at the bottom of the table, produces p -values of 0.638 (tax period) and 1.000 (post-repeal) from 500 permutations of coal shares across states. The observed coefficients are well within the permutation distribution.

Table 4: Robustness Checks

	(1)	(2)	(3)	(4)
	State Trends	Excl. Small	Excl. Trans.	Triple Diff.
Treatment \times Tax Period	-0.160*	0.056	0.089	0.068
	(0.085)	(0.100)	(0.067)	(0.076)
[0.5em] Treatment \times Post-Repeal	-0.373*	0.102	0.001	0.049
	(0.208)	(0.083)	(0.110)	(0.103)
State-specific trends	Yes	No	No	No
Excl. ACT & NT	No	Yes	No	No
Excl. transition Q	No	No	Yes	No
Manufacturing control	No	No	No	Yes
Observations	480	360	464	960
RI p -value (tax)		0.638		
RI p -value (post)		1.000		

Notes: SEs clustered at state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Col. (1): state-specific linear trends. Col. (2): drops ACT and NT (<3K employees). Col. (3): excludes transition quarters (2012Q3, 2014Q3). Col. (4): triple-difference with manufacturing as within-state control. RI p -values from 500 permutations of coal shares.

6. Discussion

The central finding—that Australia’s carbon tax produced no detectable employment effect in the most exposed sector—joins a growing body of evidence suggesting that the employment costs of carbon pricing are substantially overstated in political discourse. The result echoes Yamazaki (2017) on British Columbia, Martin et al. (2016) on the UK Climate Change Levy, and the cross-country evidence surveyed by Köppl and Schratzenstaller (2023), all finding negligible aggregate employment effects from carbon pricing. What this paper adds is the reversal: neither the imposition nor the removal of the carbon price moved employment, providing a symmetric falsification test that permanent policies cannot offer.

The null does not mean the carbon tax was costless. Electricity prices rose, transferring surplus from consumers to generators and the government. Emission reductions occurred through fuel switching and reduced generation intensity, not through employment-mediated output contractions. The welfare analysis of carbon pricing should focus on these real price and emission channels rather than on phantom employment losses that the data do not support.

Two caveats are important. First, the design exploits cross-state variation in coal share, which limits identification to the extensive margin of coal exposure. Within-state reallocation—from coal to gas or renewables employment—would be absorbed by the state fixed effects and would not appear as a treatment effect. The null is therefore a statement about net state-level employment, not about individual worker displacement, which may well have occurred. Second, the two-year duration of the tax may have been too short for full adjustment. A permanent carbon price might eventually produce employment effects that a temporary one did not. The null applies to the policy as implemented, not to a hypothetical perpetual carbon tax.

7. Conclusion

Australia’s carbon pricing mechanism—imposed on the developed world’s most coal-dependent grid and then uniquely repealed—provides a rare controlled experiment in climate policy. The “job-killing” narrative that drove the repeal finds no empirical support for large employment losses in the labor market data, though moderate effects cannot be fully excluded. Electricity-sector employment in high-coal states evolved indistinguishably from low-coal states during the tax period, and the repeal triggered no employment recovery. The missing cliff suggests that carbon pricing’s labor market costs, at least at the sectoral level and over policy-relevant horizons, may be far smaller than the political debate assumes. The real costs and benefits of carbon pricing lie elsewhere: in emission reductions, energy prices, and the long-run direction of technological change.

Acknowledgements

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

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

Contributors: @olafdrw

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

References

- Australian Bureau of Statistics**, “Labour Force, Australia, Detailed,” Cat. No. 6291.0.55.001 2024. Accessed April 2026.
- Azevedo, Deven, Hendrik Wolff, and Akio Yamazaki**, “Do Carbon Taxes Kill Jobs? Firm-Level Evidence from British Columbia,” *Climate Change Economics*, 2023, 14 (2), 2350010.
- Bureau of Resources and Energy Economics**, “Australian Energy Statistics 2012,” Technical Report, Australian Government, Department of Industry, Canberra 2012.
- Greenstone, Michael**, “The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufactures,” *Journal of Political Economy*, 2002, 110 (6), 1175–1219.
- Köppl, Angela and Margit Schratzenstaller**, “Carbon Taxation: A Review of the Empirical Literature,” *Journal of Economic Surveys*, 2023, 37 (4), 1353–1388.
- Martin, Ralf, Laure B. de Preux, and Ulrich J. Wagner**, “The Anatomy of a Climate Policy: Lessons from the UK Climate Change Levy,” *Journal of Public Economics*, 2016, 140, 1–16.
- Meng, Sam, Mahinda Siriwardana, and Judith McNeill**, “The Environmental and Economic Impact of the Carbon Tax in Australia,” *Environmental and Resource Economics*, 2013, 54 (3), 313–332.
- , – , and – , “The Impact of the Australian Carbon Tax on Industries and Households,” *Margin: The Journal of Applied Economic Research*, 2014, 8 (1), 15–37.
- Metcalf, Gilbert E. and James H. Stock**, “The Unequal Economic Consequences of Carbon Pricing,” *NBER Working Paper*, 2023, (31221).
- O’Gorman, Marianna and Frank Jotzo**, “Impact of the Carbon Price on Australia’s Electricity Demand, Supply and Emissions,” Technical Report, Centre for Climate Economics and Policy, Australian National University 2014.
- Walker, W. Reed**, “The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce,” *Quarterly Journal of Economics*, 2013, 128 (4), 1787–1835.

Yamazaki, Akio, “Jobs and Climate Policy: Evidence from British Columbia’s Revenue-Neutral Carbon Tax,” *Journal of Environmental Economics and Management*, 2017, 83, 197–216.

A. Data Appendix

Employment Data. The ABS Labour Force Survey, Detailed (Cat. No. 6291.0.55.001) provides quarterly employment estimates by ANZSIC industry division and state/territory. Table 5 reports total employed persons, employed full-time, and employed part-time. I use total employed persons (original series). The survey is based on a rotating sample of approximately 26,000 households each month, with quarterly estimates derived from the middle month of each quarter. Data are available from November 1984 through February 2026. I restrict the analysis to 2005Q1–2019Q4 to avoid the disruptions of the COVID-19 pandemic and to provide a balanced panel with sufficient pre- and post-treatment observations.

Coal Share. Baseline coal shares of electricity generation by state are from the Bureau of Resources and Energy Economics (BREE) *Australian Energy Statistics 2012*, Table O, which reports generation by fuel type and state for the 2010–11 fiscal year. These shares reflect the permanent infrastructure of each state’s generation fleet: NSW and Queensland have extensive black coal reserves; Victoria has the Latrobe Valley brown coal deposits; Tasmania has major hydro capacity; South Australia invested early in wind; and the Northern Territory relies on local gas.

Sample Restrictions. I include all eight states and territories. The ACT has very small and volatile electricity employment (mean 1,800) and zero coal share; the Northern Territory is similar (mean 2,100, zero coal). Excluding these two territories does not affect the main results (Table 4, Column 2).

B. Identification Appendix

Pre-Trends. Table 5 reports selected event study coefficients from the specification interacting coal share with quarterly indicators. Pre-treatment coefficients for event times -12 through -2 are jointly insignificant (F -test $p > 0.50$), and no individual coefficient exceeds 0.46 in absolute value, all with wide confidence intervals spanning zero. Tax-period and post-repeal coefficients are similarly insignificant.

Table 5: Event Study Coefficients: Coal Share \times Quarterly Indicators

Event Time	Period	Estimate	SE	95% CI
-12	Pre	0.455	(0.518)	[-0.560, 1.469]
-8	Pre	0.447	(0.391)	[-0.319, 1.213]
-4	Pre	0.003	(0.124)	[-0.241, 0.246]
-2	Pre	0.098	(0.202)	[-0.297, 0.493]
0	Tax	-0.101	(0.109)	[-0.314, 0.113]
+2	Tax	-0.232	(0.278)	[-0.776, 0.313]
+4	Tax	0.505	(0.545)	[-0.562, 1.573]
+8	Post-repeal	0.131	(0.253)	[-0.365, 0.627]
+12	Post-repeal	0.354	(0.249)	[-0.135, 0.842]
+16	Post-repeal	-0.049	(0.394)	[-0.821, 0.724]
+20	Post-repeal	0.034	(0.278)	[-0.511, 0.580]

Notes: Selected event-time coefficients from the specification $\ln Y_{st} = \alpha_s + \gamma_t + \sum_{k \neq -1} \delta_k \cdot \text{CoalShare}_s \times \mathbf{1}[t = k] + \varepsilon_{st}$. Reference period is event time -1 (2012Q2, the last pre-treatment quarter). Standard errors clustered at the state level in parentheses. Pre-treatment coefficients (event times -12 through -2) are individually and jointly insignificant. Tax period: event times $0-7$ (2012Q3–2014Q2). Post-repeal: event times $8+$ (2014Q3 onward). Full set of 33 event-time coefficients available upon request.

Randomization Inference. I conduct 500 permutations of the vector of coal shares across the eight states, re-estimating the main specification each time. The distribution of permuted tax-period coefficients has mean -0.007 and standard deviation 0.11 . The observed coefficient of 0.049 falls at the 36th percentile of this distribution (two-sided RI $p = 0.638$). The permuted post-repeal distribution has mean -0.012 and standard deviation 0.18 ; the observed coefficient of 0.001 falls near the median (RI $p = 1.000$).

C. Robustness Appendix

Additional specifications tested but not tabulated include: (i) dropping 2008–09 (Global Financial Crisis), with results nearly identical to baseline; (ii) using quarterly log first-differences instead of log levels, which produces noisier but qualitatively similar estimates; (iii) weighting by baseline state employment, which slightly increases precision without changing

the sign or magnitude of coefficients.

D. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SE	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>								
Log employment	Tax period	0.0487	0.0630	0.187	1.211	0.0075	0.0097	Small positive
Log employment	Post-repeal	0.0005	0.1059	0.288	1.211	0.0001	0.0252	Null
<i>Panel B: Heterogeneous</i>								
High-coal states	Tax period	-0.1266	0.1868	0.282	0.228	-0.1564	0.2306	Large negative
Medium-coal states	Tax period	0.6598	0.4588	0.090	0.658	0.0898	0.0625	Moderate positive

- Notes:** **Country:** Australia. **Research question:** Did Australia’s carbon pricing mechanism (AUD 23–24/tonne CO₂, July 2012–July 2014) reduce electricity-sector employment in coal-intensive states? **Policy mechanism:** The Clean Energy Act 2011 imposed a carbon price on approximately 500 entities covering 60% of Australia’s emissions, directly increasing the cost of coal-fired electricity generation relative to gas and renewables, with the policy repealed after two years. **Outcome definition:** Log quarterly employment in ANZSIC Division D (Electricity, Gas, Water and Waste Services) from ABS Cat. 6291.0.55.001, measuring the total number of employed persons in the electricity-sector workforce. **Treatment:** Continuous — baseline coal share of state electricity generation (2010–11), ranging from 0.00 (ACT, NT) to 0.92 (Victoria). **Data:** ABS Labour Force Detailed (Cat. 6291.0.55.001) Table 5, quarterly, 2005Q1–2019Q4, 8 states and territories, 480 state-quarter observations. **Method:** Continuous-treatment difference-in-differences with state and quarter fixed effects, standard errors clustered at the state level. **Sample:** All eight Australian states and territories; electricity sector only (ANZSIC Division D). SDE = $\hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$ where $\text{SD}(X)$ is the standard deviation of the treatment interaction and $\text{SD}(Y)$ is the unconditional standard deviation of log employment. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).