Power Sector Emissions under Tightening Carbon Dioxide Quotas*

Travis Roach† and R. Kaj Gittings‡

June 9, 2021

Abstract

The Regional Greenhouse Gas Initiative (RGGI) was passed by an original collection of 10 northeastern states and is the first cap-and-trade policy in the United States to specifically target carbon dioxide emissions from the electricity sector. We exploit the introduction of this policy and subsequent tightening of the carbon cap to assess how carbon dioxide emissions have changed within RGGI states while also evaluating emissions leakages that may have occurred. While RGGI auction prices are relatively low compared to estimates of the social cost of carbon, the policy represents a direct increase in the cost of production and should have behavioral impacts on producers. Using plant-level data and several identification strategies, we find that there are reductions in emissions from coal-fired plants in RGGI states, but there is mixed evidence at natural gas-fired plants. We also show that the emissions cap reductions enacted in 2014 is where the policy began to have more significant impacts on emissions. These conclusions are strongest with careful evaluation of control states since spillover effects of the policy to non-RGGI states are possible within the Eastern Interconnection.

JEL Codes: H23, H71, L94, Q48, Q58
Keywords: Cap and Trade; Carbon Dioxide; Regional Greenhouse Gas Initiative

*We are thankful for comments on early drafts of this work by R. Bruce Williamson of the Maine Public Utilities Commissiom, and discussants and session participants of the AERE@WEAI sessions of the 94th WEAI conference.

†Corresponding author. University of Central Oklahoma, troach2@uco.edu.
‡Texas Tech University, kaj.gittings@ttu.edu

1 Introduction

The warming effect of greenhouse gas accumulation is widely accepted in the scientific community, however the regulation of greenhouse gas emissions remains a hotly contested and debated topic — especially in the United States. Despite global collaboration efforts and agreements to reduce greenhouse gasses, specific policies intended to reduce these emissions have not been widely applied and are mostly regionally-based. Within the United States, market-based instruments like emissions taxes or carbon permit trading systems have been proposed at both state and Federal levels, but they have been met with limited political support.[1] One market-based system that has been put in place is the Regional Greenhouse Gas Initiative (RGGI) which is a coalition of eleven states (originally 10 states) in the northeastern United States that have voluntarily agreed to a gradually more restrictive cap on carbon dioxide emissions from the electric power sector.

The RGGI was the first cap and trade system for carbon dioxide emissions put in place in the United States, and this carbon trading system has been in force since 2009. When the policy began, the RGGI implemented a carbon cap of 188 million allowances for the 10 state region which tightened incrementally on an annual basis. The original coalition of 10 states included Connecticut, Delaware, Maine, Maryland, Massachusetts, New Hampshire, New Jersey, New York, Rhode Island, and Vermont. In 2012 New Jersey left the coalition but rejoined the RGGI in 2020. Virginia also joined the RGGI in 2020 while Pennsylvania passed legislation to join the RGGI in 2022, which will bring total membership to twelve states.

In the time since the policy's original implementation, the consortium of states within the RGGI voted to lower their emissions cap in 2014. This was done because the emissions cap was effectively non-binding due to an unanticipated abundance of cheap natural gas as prices

[1] For example, California has passed an emissions trading scheme, while the state of Washington has proposed a carbon tax that failed to garner enough votes to pass. They have since established a cap-and-trade bill.

collapsed from more than \$13 per million BTU in July 2008 to just over \$3 per million BTU one year later in July 2009. Natural gas prices continued to decline, motivating an organic market driven substitution away from coal and hence lower emissions nationwide. However, the newly enacted emission cap reductions from 2014 were far more strict. For example, the original 2020 cap was 78,175,215 allowances which equated to 41% of the emissions allowances in 2009. The new adjusted RGGI cap for 2020 is 56,283,807 allowances or 29.9% of the 2009 level. After more than 51 quarterly auctions in the 11-year existence of the policy, more than \$3.9 Billion has been collected in revenues which have been proportionally dispersed back to individual RGGI states.[2] This results in an average of more than \$350 million per year that has been extracted from the utilities sector alone in these RGGI states, a sum large enough that should affect production activities of power producing plants.[3]

While it is true that carbon dioxide emissions have been falling across the United States due to coal plant retirements, the cost-competitiveness of natural gas since the price collapse, and increased renewable energy capacity, early analyses of the RGGI indicate that over half of the observed decline in emissions in the RGGI-area are due to the policy (Murray and Maniloff (2015)). However, because the RGGI is a regional initiative that takes place within a larger body of electricity flow between both adopting and non-adopting states, there is great potential for spillover effects from the policy. In fact, early predictions expected carbon-intensive power generation to flow from non-adopting states into RGGI adopting states (Chen 2009). Early empirical research concerning the RGGI has largely borne out this expectation. Kim and Kim (2016) determine that leakage has occurred by using synthetic control methods and yearly state-level data on the share of natural gas used to generate electricity. Lee and Melstrom (2018) also make use of state-level electricity flows, though

[2]Auction results and accumulated proceeds can be found here: https://www.rggi.org/auctions/auction-results

[3]Once dispersed to states, these permit revenues have been used to fund activities like energy efficiency audits and improvements, renewable energy development, low-carbon city services and consumer billing credits to offset higher prices. A report on state by state investments is available in RGGI (2016) which can be found here: https://www.rggi.org/investments/proceeds-investments

at a quarterly interval. Fell and Maniloff (2018) and Chan and Morrow (2019) are the only others besides our work that use plant-level data to show that leakage has occurred. Fell and Maniloff (2018) use the proximity of Ohio and Pennsylvania and designate them 'leaker' states. They find that the capacity factor of natural gas-fired plants have changed in these leaker states following the implementation of the RGGI. Chan and Morrow (2019) show that leakage has occurred for both carbon dioxide emissions and associated emissions, and they are the only other paper to use observed emissions (not simply electricity generation) as we do in this paper. Yan (2021) finds that natural gas consumption has increased by as much as 237% in outside regions. The consensus of this early literature is that carbon leakage is occurring, and that this occurs by increased power production in neighboring (non-adopting) states at natural gas-fired plants. In this research, we show that this conclusion may be biased due to the choice of 'control group' states that are considered. Here, we find similar evidence that spillover has occurred and that it primarily comes through increased natural gas emissions, but that the parallel trends assumption that is critical for identification fails in those specifications as does SUTVA (Stable Unit Treatment Value Assumption) simply by the existence of spillovers into controls. Another unintended spillover effect that has been studied considers how the passage of the RGGI has affected co-pollutants like Sulphur-dioxide (SO_2) and Nitrogen-oxide (NO_x) which are also emitted on combustion. Numerous studies have found that any co-benefits that may accrue are dependent on the level of regulation on the other pollutants (Fullerton and Karney (2018)).[4] An important leakage effect that has been caused by the RGGI's implementation is that SO_2 has shifted from the RGGI region to areas with higher marginal damages from SO_2 (Chan and Morrow (2019)).

Our paper makes several contributions to the existing literature. First, we are able to identify the differential effect of the RGGI on plant-level emissions for both coal and natural gas-fired units. Given the differing carbon content and production costs of each primary

[4]As a practical example, potential interactions with the existing SO_2 permit system.

fossil-fuel, and existing evidence from the literature, it is possible that the RGGI has had a differential impact on plant-level emissions for these alternative electricity generating fuels. This delineation between coal and natural gas plants is important because we ultimately find robust evidence of emissions falling at the former, and mixed evidence of both an increase and a decrease in emissions for the latter. Second, we assess not only the introduction of the RGGI as prior authors have done, but also measure how the tightening of the cap since 2014 has changed emissions. The RGGI is not unique in its need to lower the emissions quota from the original policy design. Haites (2018) notes that all emissions trading systems have accumulated surplus allowances and lowered their carbon caps. By separately estimating reductions during these two treatment periods we are able to determine the incremental effects of the RGGI on emissions - under binding and non-binding cap scenarios. After modelling and addressing the prior two factors within the RGGI (multiple plant types and multiple policy treatment dates), we assess leakage that has occurred as a result of the policy. Prior authors have shown that nearby states and states within the same electricity balancing authority jurisdiction have witnessed an increase in electricity generation, especially from natural gas-fired plants. However, our study is able to determine how plant-level emissions have changed in these non-RGGI regions to make up for reductions in electricity generation within the RGGI region by plant and by treatment period. The results of these analyses cast doubt on prior studies' results because the treatment-free comparison group is not stable. Indeed, spillover into a separate region is evidence in and of itself that the standard group of control plants are tainted. To combat this, we analyze alternative sets of controls that are plausibly free of contamination by excluding plants in 'leaker' states as well as those connected to the same electricity grid as RGGI states. Lastly, we determine how the RGGI has impacted co-pollutants at both the state-level and plant-level cumulatively and by treatment period.

To preview results, we find robust evidence that plant-level emissions have declined at coal-fired plants in the RGGI since the implementation of the policy. Coal plant emissions in RGGI-adopting states have fallen by approximately 20% per plant per year on average (relative to plants in non-adopting states) with larger reductions occurring during the period when the emissions cap was tightened. At natural gas-fired plants, we find mixed evidence of the role that carbon dioxide pricing played in changing the amount of plant-level emissions. Under some control group specifications we find weak evidence that emissions have declined since in most specifications we cannot differentiate the effect from zero, and in one specification there is evidence that emissions actually increased. Further, estimates from our event-study models cast doubt on prior studies' evidence of carbon leakage occurring. In a standard difference-in-differences specification, we too are able to observe potential leakage but our event-study models show that the parallel trends assumption is clearly violated regardless of the control group we adopt. Hence the conclusion in the literature of emissions leakage should be regarded with caution.

2 Empirical Strategy

Our analysis begins with a brief state-level examination of the RGGI effect on emissions relative to non-RGGI states. This analysis does not distinguish the fuel source of electricity generating plants and takes the form of a standard differences-in-differences model. The RGGI is sharply defined along state lines, so electricity generating sources outside of the RGGI are used as a 'treatment-free' comparison group subject to the same global factors that both treated and untreated units face (e.g. commodity price changes).

These models are problematic for three reasons. First, spillover effects and leakage into controls violates SUTVA (Stable Unit Treatment Value Assumption) and the usual assumption that the controls are 'treatment-free'. Second, we show later in our event study analysis that the parallel trends assumption is violated in specifications that allow spillover states to be used as controls. Third, aggregating plant-level emissions to the state-level masks the effects of the policy. Some plants may reduce emissions (i.e., coal powered plants) while other plants may increase emissions (i.e., natural gas plants). In aggregation, these effects may cancel each other out potentially yielding a muted effect of the policy – even if the policy has a strong bite at the plant-level. However, presenting these state-level results also serves two purposes. One is that our work can be compared to previous research investigating the RGGI and emissions reductions at the state-level. The other is that our primary analysis, results and conclusions at the plant-level can be understood in context to a standard state-level analysis.

There is a clear start date to the carbon trading policy which began in 2009 which we use as the first 'treatment' period.[5] Fortunately, we are not concerned with issues of staggered adoption that have been brought to the forefront of modern difference-in-differences methodology because the policy was instituted in all adopting states simultaneously (despite staggered ratification by state governments).[6] We are also able to use variation in the intensity of treatment because the cap on emissions was lowered in 2014. This subsequent policy intervention serves as our second treatment period.

While the state-level analysis is meant to be illustrative, the plant-level analysis is more appropriate since it can distinguish the fuel source of each plant. To integrate the differential effect of the RGGI based on the fuel source each plant uses, we measure the impact of the RGGI on plant-level emissions using a triple difference-in-differences framework. This is an important source of variation because plants that use coal or natural gas will likely be impacted by carbon pricing differently since these fuel sources have different carbon factors.[7] In the plant-level analyses we evaluate the viability of different control groups to combat the violation of SUTVA in the standard model. We also use the plant-level data to conduct event study analyses and evaluate the parallel trends assumption.

[5]The first auction took place on September 25, 2008 but the permits are for future emissions so we begin our treatment policy in 2009.

[6]See, for example, Goodman-Bacon (2018) and Callaway and SantAnna (2020)

[7]Carbon factors are the amount of carbon released per unit of heat (BTU).

2.1 Data Description

The data on plant-level emissions are collected from the Energy Information Administration. These data measure the amount of carbon dioxide (CO_2), sulphur dioxide (SO_2), and nitrogen oxide (NO_x) emissions produced by each power-plant and are at an annual frequency from 2001-2018.[8] We merge the reported emissions for each plant with data on yearly generation by fuel input to determine the thermal input for each plant (e.g. bituminous or lignite coal versus natural gas, etc.). The universe of all electricity generating plants are included in the data - this includes coal powered plants, natural gas powered plants, nuclear powered plants, and plants that use renewable inputs such as wind, solar, geothermal, and biomass. A small number of individual plants use multiple fuel sources - e.g., a plant that uses both natural gas and coal as fuels to generate electricity - but these cases are rare.

Table 1 shows state-level and plant-level summary statistics for each pollutant. It is evident that CO_2 emissions account for by far the largest share of total emissions across all plant types, and coal fired plants emit approximately 5 to 6 times more CO_2 than natural gas plants. Also, there is a tremendous amount of variation in emissions across plants with some plants emitting zero or near zero emissions while other plants have many magnitudes greater emissions than the mean. This highlights the importance of conducting the analysis at the plant-level rather than higher level aggregations since the RGGI policy (and subsequent cap reductions) will not be binding across all plants but will impose meaningful constraints on others.

2.2 Econometric Specification

The basis of our analysis uses the method of differences-in-differences. The state-level version of the benchmark model is given by equation (1).

[8] 2018 is the latest year that plant-level data are available.

Table 1. Summary Statistics

State-level Emissions

	RGGI Region	Non-RGGI
CO_2	15.395	17.240
	(2.205)	(1.377)
SO_2	8.685	10.907
	(2.609)	(1.643)
NO_x	8.960	10.641
	(1.403)	(1.229)

Plant-level Emissions

Coal-powered	RGGI Region	Non-RGGI
CO_2	13.622	13.929
	(1.290)	(1.742)
SO_2	7.699	8.025
	(2.076)	(1.887)
NO_x	6.991	7.369
	(1.272)	(1.587)
Natural Gas-powered		
CO_2	11.076	10.740
	(2.232)	(3.042)
SO_2	1.943	2.099
	(2.577)	(3.028)
NO_x	4.272	4.514
	(1.849)	(2.247)

Notes: Emissions are expressed in natural logarithm. Mean is shown with standard deviations below in parentheses. Unit of observation is a state-year (top) and a plant-year (bottom).

$$CO2_{st} = \beta_1 Begin_t + \beta_2 Lower_t + \beta_3 RGGI_s \tag{1a}$$
$$+\gamma_1 Begin_t * RGGI_s + \gamma_2 Lower_t * RGGI_s + \varepsilon_{st}$$

$$\varepsilon_{st} = \mu_s + \lambda_t + t\mu_s + \nu_{st} \tag{1b}$$

The plant-level benchmark model is shown in equation (2). Here, instead of aggregating all emissions to the state-level we are able to measure the differential impacts that the RGGI policy has had on plant-level emissions depending on their primary fuel source.

$$CO2_{ist} = \beta_1 Begin_t + \beta_2 Lower_t + \beta_3 RGGI_s + \beta_4 C_i + \beta_5 NG_i + \beta_6 Begin_t * RGGI_s \tag{2a}$$
$$+\beta_7 Begin_t * C_i + \beta_8 Begin_t * NG_i + \beta_9 Lower_t * RGGI_s + \beta_{10} Lower_t * C_i$$
$$+\beta_{11} Lower_t * NG_i + \beta_{12} RGGI_s * C_i + \beta_{13} RGGI_s * NG_i$$
$$+\gamma_1 Begin_t * RGGI_s * C_i + \gamma_2 Lower_t * RGGI_s * C_i$$
$$+\gamma_3 Begin_t * RGGI_s * NG_i + \gamma_4 Lower_t * RGGI_s * NG_i + \varepsilon_{ist}$$

$$\varepsilon_{ist} = \mu_i + \lambda_t + t\mu_s + \nu_{ist} \tag{2b}$$

The outcome $CO2_{ist}$ is the amount of carbon dioxide emissions (in natural logarithm) from plant i in state s in year t in both equations. The variable $RGGI_s$ is a dichotomous indicator for RGGI adopting states, $Begin_t$ is a dichotomous indicator that is set to 1 for the early years of the RGGI before the cap reductions took place (2009-2013). The second policy variable ($Lower_t$) indicates when the cap on emissions was updated and lowered in 2014 which is set to a value of 1 from 2014 to the end of the sample period 2018. Addressing the cap reduction in this manner is new to the empirical literature on the RGGI. The variables, C_i and NG_i indicate whether the generating source uses coal or natural gas as a fuel source. The stochastic component ε_{ist} contains plant fixed effects (μ_i), year fixed effects (λ_t), and state-specific linear time trends ($t\mu_s$). All models are estimated by least squares where standard errors are clustered by plant for plant-level estimates, and clustered by state for state-level estimates.[9]

One method of gauging model validity is through the use of an event-study model. The benchmark version of the model is outlined in equation (3) but can be amended to include additional parameters for unobserved heterogeneity. Here, the model produces an estimated effect of the RGGI for each year in the data. The hypothesis is that if the parallel trends assumption is valid, then the evolution of the estimated effects of the policy (β_2^t) should be statistically zero in the pre-period and only deviate from zero in the post-period. Furthermore, the models in equations (1) and (2) produce the estimated effect of the RGGI averaged over two distinct time periods in the post-period, but the event study model in equation (3) allows one to see the evolution of the policy effects year-to-year.

$$CO2_{ist} = \beta_0 + \beta_1 RGGI_s + \sum_{\substack{t=t_0 \\ t \neq 2008}}^{t} \Big[\beta_2^t RGGI * I(Year = t)\Big] + \varepsilon_{ist} \tag{3a}$$

$$\varepsilon_{ist} = \mu_i + \lambda_t + t\mu_s + \nu_{ist} \tag{3a}$$

2.3 Identification

Mechanically, the parameter on $RGGI_s$ is identified because New Jersey left and then reentered the RGGI during the sample period. Similarly, the parameters on C_i and NG_i are identified due to the small number of plants that produce electricity using both coal and natural gas as fuel sources where those cases are coded as both a coal plant and a natural gas plant. The main effects $Begin_t$ and $Lower_t$ are included in equations (1) and (2) but

[9]There are more than 9,000 plants in the data and therefore estimating a model that includes plant-specific effects and plant-specific linear time trends (an additional 9,000 parameters to estimate) is computationally burdensome. Including state-specific trends in the plant-level model imposes the assumption that all plant-level trends within a state are the same.

are absorbed by the year fixed effects. The primary coefficients of interest are γ_1 and γ_2 in the state-level model and those attached to the triple interactions (γ_1 through γ_4) in the plant-level model as these represent the treatment effect of carbon legislation on plant-level emissions for coal or natural gas generating sources relative to plant-level emissions at same-fuel-type plants outside of the RGGI area. In the tables presenting results that follow, we only show parameter estimates from the treatment interactions for the sake of brevity.

One fundamental issue with estimating the effect of the RGGI on emissions or other outcomes is that natural gas prices went through a strong and continual decline during the period in which the RGGI was enacted. As natural gas became less expensive relative to coal, producers had a natural incentive to substitute away from coal to natural gas which would lead to a decline in emissions even without the RGGI policy. Without accounting for this event, the decline in natural gas prices would generate downward bias in our estimates of the independent effect of the RGGI on emissions. The parameter λ_t captures nonlinear period specific shocks that vary over time but are common across states and plants. Therefore, this parameter accounts for the evolution of the market price of natural gas over time. Furthermore, to the extent that the decline in natural gas prices affected plant emissions in states differently (e.g., accessibility through pipelines and distribution), our models also include state-specific time trends ($t\mu_s$) that capture changes in emissions that vary across states. Conditional on plant fixed effects also in the model, any remaining impacts on plant emissions due to the decline in natural gas prices are likely random.

We have particular concerns with interpreting the estimated parameters as causal in the benchmark specifications. The first is that identification requires that the parallel trends assumption is satisfied, and while our inclusion of state-specific time trends may alleviate those issues, it may not be sufficient if agents respond to the policy before it was enacted because the policy was announced in advance and the early response was differential across treated and control units. The second is that we are interested in spillover affects outside

Figure 1: State-level Emissions by Region

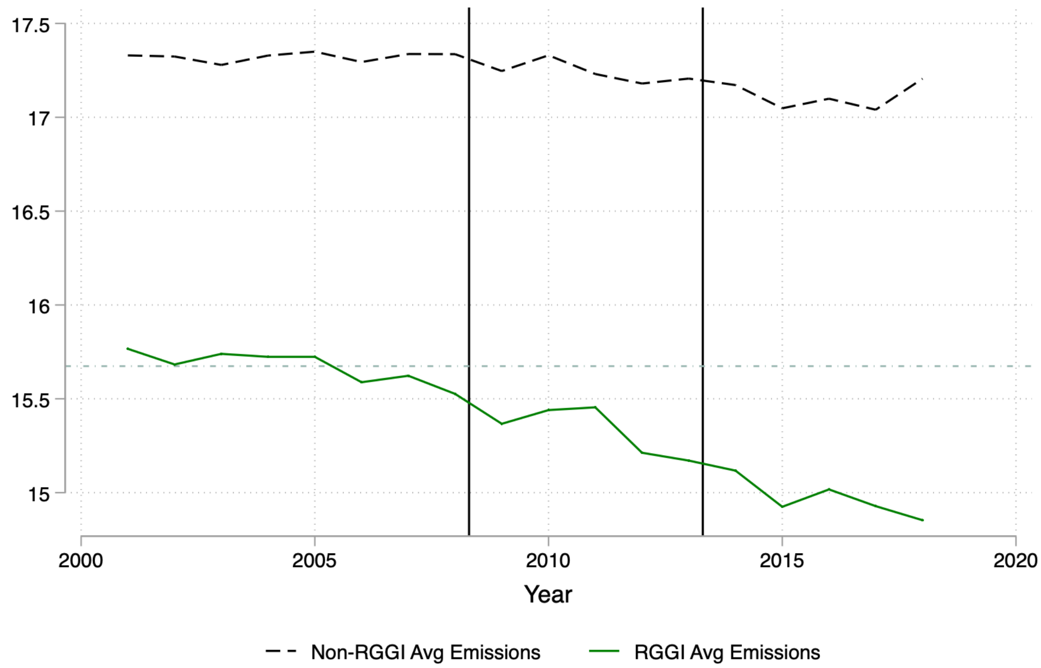

the RGGI region which would contaminate the controls in the benchmark specification and violate SUTVA by definition. Similarly, RGGI states are part of the Eastern Interconnection grid and PJM balancing authority that are composed of states that are not part of the RGGI and are otherwise members of the control group in the benchmark specification. These institutional constructs are also likely contributors to a violation of SUTVA. To address the parallel trends concern, we conduct event study analyses and show under which specifications the parallel trends assumption holds or is violated. To combat the likelihood of contaminated controls and SUTVA violations in the benchmark model, we construct alternative sets of controls that are plausibly unaffected by the RGGI policy.

3 State-level Emissions

Figure 1 offers a first look at our eventual state-level findings. In this figure, we chart average state-level emissions in the RGGI region and average state-level emissions in non-RGGI

Table 2. State-level Estimates

	Dep Var.	(1) CO_2	(2) SO_2	(3) NO_x
RGGI · Begin		-0.070	-0.513**	-0.007
		(0.088)	(0.221)	(0.121)
RGGI · Lower		-0.075	-0.694*	-0.103
		(0.178)	(0.401)	(0.203)
State FE		Y	Y	Y
Year FE		Y	Y	Y
State Trends		Y	Y	Y
R^2		0.994	0.974	0.983
Obs		867	867	867

Notes: Clustered robust standard errors shown in parentheses; *,**,*** denote statistical significance at the 10%, 5%, and 1% levels, respectively

states. We include a dotted line that equals the pre-implementation mean of emissions for the RGGI region. The vertical lines indicate when the RGGI was implemented and when the cap reductions took place. The state-level emissions data shows a decline for both adopting and non-adopting regions – implying a small treatment effect that may not be statistically discernible from similar trends in non-adopting states. Note the difference between this state-level time-series and Figure (2) that presents emissions data at the plant-level across groups. In Figure (2) each fuel source visually decreases relative to non-RGGI states in the post-period. This suggests that analyses at a state-level aggregation may miss important differences across plants.

Nevertheless, in Figure 1 it is more evident that emissions have fallen for both RGGI and non-RGGI states and there does not appear to be much of a difference between the paths of each group. This drop in emissions is is likely driven by many factors, some of which are common to both RGGI and non-RGGI adopting states, including: macroeconomic distress, increases in the amount of installed capacity of renewable energy and its infusion into the generation mix, the 2009 stimulus bill which subsidized wind and solar power, Obama-era coal regulations, and the drop in the price of natural gas which has driven an increase in the quantity demanded of natural gas-powered electricity.

We test whether or not the RGGI has impacted emissions at the state-level in Table 2 by pollutant. All models in columns (1) to (3) are specified according to equation (1). Here, we see that carbon dioxide emissions have fallen in the RGGI, but this effect is not statistically discernible from similar reductions outside of the RGGI. There is evidence that sulphur-dioxide emissions fell in both periods. There is also not a statistically discernible change in nitrogen-oxide emissions due to the RGGI in this state-level analysis.[10]

4 Plant-level Changes in CO_2 Emissions

Figure 2 illustrates emissions by coal-fired (top lines) and natural gas-fired (bottom lines) electricity plants over time for RGGI and non-RGGI regions. For both plant-types we also include a dotted line equal to the pre-implementation RGGI mean emissions.[11] Here, we see that plant-level emissions appear to be stable across non-adopting regions for both fuel types in the pre-period, and that there are declines in mean plant-level emissions following each policy event. For coal-fired facilities, this effect is nearly one log-point lower in post-treatment periods, but for natural gas facilities we notice a much smaller change over time.

Table 3 displays our initial results from estimating variations of equation (2) and reflect how the RGGI policy has effected plant-level emissions. Columns (4) and (6) are specified as equation (2) whereas columns (3) and (5) account for fewer unobservables. The $Begin_t$ variable in columns (1) and (2) is coded differently than our benchmark specification. In columns (1) and (2), $Begin_t$ is set equal to one in 2009 and does not turn off through the sample period. Therefore, in columns (1) and (2) the estimates of the RGGI reflect the average effect over the entire policy period in the same way previous literature has done. Columns (3) and (4) differentiate the treatment effect for the initial policy implementation and the subsequent cap reduction in 2014 as originally described in Section 2. Columns (5)

[10] Table A1 in the appendix shows state-level estimates for all pollutants along with estimates for any state-level leakage effects.

[11] Vertical lines indicate the beginning of the policy in 2009 and the cap reduction in 2014.

Figure 2: Plant-level Emissions by Region; Coal (top) NG (bottom)

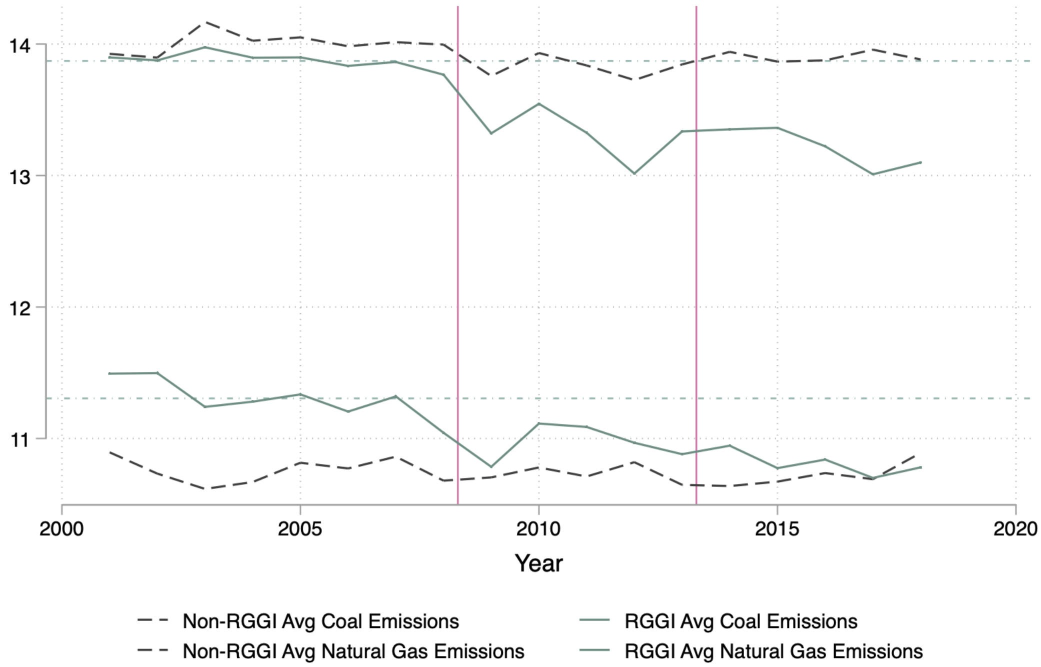

and (6) are a robustness check to determine if the new cap-and-trade policy was preempted by firms in the RGGI region who may have reduced emissions in advance given that the policy was pre-announced. These results in columns (5) and (6) are discussed in the next section but are presented here for ease of comparison to the main results.

The fully specified model in column (2) indicates that over the entire period of the RGGI policy, emissions have fallen by approximately 0.22 log-units (19.7% from column (2))[12] on average per year at coal-fired plants in the RGGI relative to plants outside the RGGI, but that natural gas-fired plants have not experienced statistically different emissions reductions as a result of the overall policy. When the policy is separated to identify the effect of the early period (2009-2013) and the cap reduction period (post-2013), it becomes clear that the emissions reductions primarily occurred when the emissions cap was reduced. The fully specified model in column (4) suggests emissions fell by about 0.46 log points (36.9%) in

[12]These percentages are calculated as $e^{\hat{\beta}} - 1$.

Table 3. Plant-level Estimates

		(1)	(2)	(3)	(4)	(5)	(6)
	Dep. Var	CO_2	CO_2	CO_2	CO_2	CO_2	CO_2
RGGI · Coal · Begin		-0.286***	-0.219**	-0.187*	-0.140	-0.043	-0.008
		(0.103)	(0.111)	(0.096)	(0.101)	(0.116)	(0.121)
RGGI · Coal · Lower		-	-	-0.588***	-0.462***	-0.454***	-0.360**
				(0.155)	(0.174)	(0.164)	(0.181)
RGGI · NatGas · Begin		-0.103	-0.134	-0.079	-0.119	-0.049	-0.095
		(0.099)	(0.097)	(0.096)	(0.093)	(0.119)	(0.116)
RGGI · NatGas · Lower		-	-	-0.133	-0.181	-0.107	-0.152
				(0.125)	(0.128)	(0.142)	(0.146)
RGGI · Coal · Early						0.265***	0.267***
						(0.088)	(0.088)
RGGI · NatGas · Early						-0.013	-0.042
						(0.101)	(0.100)
Plant FE		Y	Y	Y	Y	Y	Y
Year FE		N	Y	N	Y	N	Y
State Trends		N	Y	N	Y	N	Y
R^2		0.944	0.946	0.944	0.946	0.945	0.947
Obs		55035	55035	55035	55035	55035	55035

Notes: Clustered robust standard errors shown in parentheses; *,**,*** denote statistical significance at the 10%, 5%, and 1% levels, respectively

RGGI coal plants during the cap reduction period relative to plants outside the RGGI, but there was a marginally significant and much smaller reduction in the early period of the RGGI.

For natural gas plants, we do not find any evidence that emissions have decreased as a result of carbon pricing in Table 3. On the one hand, this is surprising since while natural gas is cleaner fuel source, carbon pricing ultimately still increases production costs of electricity produced by natural gas facilities. However, the dramatic decrease in natural gas prices likely offset much of the costs imposed by the RGGI for natural gas plants. The lower carbon factor of natural gas allows for the price-differential to still decrease enough to cause inter-fuel substitution to natural gas, even within the RGGI and during the tighter carbon cap period. For example, natural gas produces 117 pounds of carbon dioxide per million BTU while coal produces between 205.7 - 251.6 pounds per million BTU depending on the type of coal.[13] So, all else equal, about twice as many carbon permits must be purchased for the same amount of heat output in coal plants.

4.1 Alternative Specifications

In this section, we explore alternatives to the set of control group observations we use to measure an average treatment effect. This analysis evaluates concerns of contaminated control units and violations of SUTVA. We also perform additional robustness checks that account for common concerns with the use of difference-in-differences modeling as a method to determine treatment effects. We also determine whether or not electricity generating plants within the RGGI region pre-empted the policy by making production changes in advance of RGGI states signing on to the agreement. We then provide an event study analysis for both coal and natural gas plant emissions to evaluate in which models the common pre-period trends assumption holds and to provide a clearer picture of how the RGGI effects evolve in the post-period.

We begin by exploring alternative control groups as a comparison to RGGI states. In our benchmark estimates above, the estimated treatment effects are with reference to changes in plant-level emissions in all states that are not members of the RGGI. The mere fact that previous research (and ours) finds there may be spillover emissions to nearby regions highlights that SUTVA may not hold and the estimated treatment effects might be biased. Column (1) of Table 4 replicates column (4) of Table 3 for ease of comparison, and columns (2)-(5) make incremental changes to the type of unobserved trends or control groups.

First we investigate how our estimates change when we include trends based on the fuel-type a plant uses. This inclusion stands to capture unobserved heterogeneity in trends in emissions by plant type, since coal use specifically has declined independent of the RGGI due to local and federal regulation of coal ash and other bi-products that are monitored

[13] Bituminous and Coke, respectively.

in the toxic release inventory, while natural gas prices have also declined significantly over the sample period. Here, we see that our estimated treatment effects for coal plants are unchanged, but now there exists a statistically discernible impact at natural gas-fired facilities during the RGGI period.

Next in column (3), we exclude all plants located in California from the set of assumed treatment-free comparison plants as a comparison to our benchmark model in column (1) of Table 4. This is because California implemented their own carbon dioxide pricing scheme beginning in 2013. When we exclude California, we again find consistent treatment effect estimates for coal plants, and now find that each treatment period is associated with lower emissions at natural gas plants.

The final two columns make two drastic changes relative to one another. Although the RGGI coalition is defined sharply along state lines, electricity generation and balancing is dependent on supply and demand conditions across geopolitical barriers like state lines or balancing authority regions which vary in their state-to-state makeup. For instance, the New York ISO balances supply and demand within New York only, but the PJM balancing authority balances generation and demand across West Virginia, Pennsylvania, Ohio (and parts of other states) *in addition to* many of the RGGI states. Moreover, these balancing authorities coordinate with one another to help electricity flow unimpeded within the entire Eastern Interconnection that covers states east of Colorado and north of Texas.[14] There is also the Western Interconnection and the Texas 'ERCOT' interconnection. Electricity within each interconnection flows at a different phase, so electricity physically cannot (and generally is not) transported across interconnections without the phase of the current being transformed.

In column (4), we only include plants that are part of the Western Interconnection or ERCOT as comparison plants to calculate the average treatment effect while still excluding

[14]Including parts of northern Texas.

California. This should be a very clean control group because electricity does not flow between the separate interconnections and so there literally cannot be spillover from the RGGI policy area into treatment-free comparison plants.[15] In this specification we see larger reductions in coal plant emissions and both policy periods are now statistically significant. For natural gas facilities the point estimate is now *positive* for both periods, but not statistically different from zero.

Column (5) only uses plants located within the Eastern Interconnection as the control group. Because balancing authorities in this region are connected, we may expect exogenous weather and wind profiles to impact electricity dispatch and load demand profiles such that plants in this area follow a similar trend. This set of controls is most suspect to violations of SUTVA due to spillover effects and electricity flows between RGGI and non-RGGI states. We note that this model is similar to the primary robustness check presented in Yan (2021) who finds large spillover effects. In this model, we again find a reduction in coal plant emissions, and now find each policy period is associated with lower emissions at natural gas-fired facilities.

To summarize, across all model specifications we see that coal plant emissions fall relative to non-RGGI plants as a direct result of the RGGI policy and the reductions are strongest and most precisely estimated in the cap reduction period. The largest coal emissions reductions are estimated in the model with the cleanest control group (column (4)) and the smallest coal emissions reductions are estimated using the most contaminated control group (column (5)). For natural gas emissions, point estimates are consistently negative in all models that incorporate potentially contaminated controls but turn positive with the cleanest control group (column (4)). This pattern of results is consistent with leakage from the policy occurring in neighboring states outside the RGGI but within the Eastern Interconnection and PJM balancing authority. While the finding of leakage is not new and we will demon-

[15] In the Western Interconnection and ERCOT (Texas) electricity actually flows at a completely different phase pattern.

Table 4. Alternate Controls

	(1)	(2)	(3)	(4)	(5)
Dep Var.	CO_2	CO_2	CO_2	CO_2	CO_2
RGGI · Coal · Begin	-0.140	-0.156	-0.154	-0.469***	-0.097
	(0.101)	(0.104)	(0.102)	(0.127)	(0.103)
RGGI · Coal · Lower	-0.462***	-0.541***	-0.472***	-0.850***	-0.379**
	(0.174)	(0.176)	(0.174)	(0.212)	(0.176)
RGGI · NatGas · Begin	-0.119	-0.158*	-0.163*	0.011	-0.187*
	(0.093)	(0.093)	(0.094)	(0.130)	(0.095)
RGGI · NatGas · Lower	-0.181	-0.234*	-0.250*	0.149	-0.321**
	(0.128)	(0.127)	(0.130)	(0.189)	(0.132)
Plant FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
State Trends	Y	Y	Y	Y	Y
Fuel-type Trends	N	Y	N	N	N
Excludes CA	N	N	Y	Y	Y
Eastern Interconnection Only	N	N	N	N	Y
Non-Eastern Interconnection Control Group	N	N	N	Y	N
R^2	0.946	0.947	0.949	0.933	0.950
Obs	55035	55035	49066	14929	41020

Notes: Clustered robust standard errors shown in parentheses; *,**,*** denote statistical significance at the 10%, 5%, and 1% levels, respectively

strate the occurrence of leakage in our framework below, the results in Table 5 illustrate sources of bias (contaminated controls and differential trends) in the traditional framework used to evaluate the RGGI in the literature.

The RGGI was operational at the beginning of 2009, but the policy was discussed among eventual member states as early as 2003. In late December 2005, seven of the RGGI states announced an agreement to implement the RGGI, as outlined in a Memorandum of Understanding (MOU) signed by the Governors of Connecticut, Delaware, Maine, New Hampshire, New Jersey, New York, and Vermont. This MOU, as amended, still provides the major policy design elements and governance rules of the RGGI to this day. Two years later, Massachussets and Rhode Island both signed the MOU, and Maryland soon followed. Given the foreknowledge of joining the RGGI and the delay before the policy officially began in 2009, it is possible that plants within RGGI adopting states began to implement or change their generation *prior* to the first compliance period. If this occurred, then the estimated effect for emissions reductions in our benchmark models above would be biased toward zero.

To address this issue, we augment the model in equation (2) with an additional pre-2009 indicator ($Early_t$) and its interactions to determine whether or not plants within the RGGI began to change behavior prior to the actual implementation of the carbon permit trading system.[16] The 'pre-treatment' effects (parameter estimates for $RGGI * Coal * Early$ and $RGGI * NatGas * Early$) are found in columns (5) and (6) of Table 3. We find no evidence that natural gas plants pre-empted the policy, but the estimated treatment effect at coal units is positive and statistically significant.

To better understand the evolution of the impact of the RGGI across the pre- and post-adoption periods, we return to our event study analyses. Figure (3) displays coefficients of the event-study model for coal plants and Figure (4) shows the same for natural gas plants. In each figure: the top panel plots period specific coefficients analogous to our benchmark model (Table 3, column 4); the middle panel plots period specific coefficients when we include fuel-type trends (analogous to Table 4, column 2); and the bottom panel plots period specific coefficients when we only use plants outside of the Eastern Interconnection as the treatment-free comparison group (analogous to Table 4, column 4).

Several findings emerge from Figures (3) and (4). First is that the top panel of Figure (3) illustrates the preemptive relative increase in coal emissions using the benchmark model (and all non-RGGI states as controls) and the inclusion of fuel-type trends does not alleviate this result. However, the bottom panel that only includes plants in states outside the Eastern Interconnection as controls does not indicate that RGGI states preempted the policy. This suggests that plants in non-RGGI states within the Eastern Interconnection behaved differently than plants within RGGI states before the policy took effect and that the common trends assumption within that sample is likely violated and that the controls are contami-

[16] $Early_t = 1$ when a state signs the MOU.

Figure 3: Coal Event Study Estimates

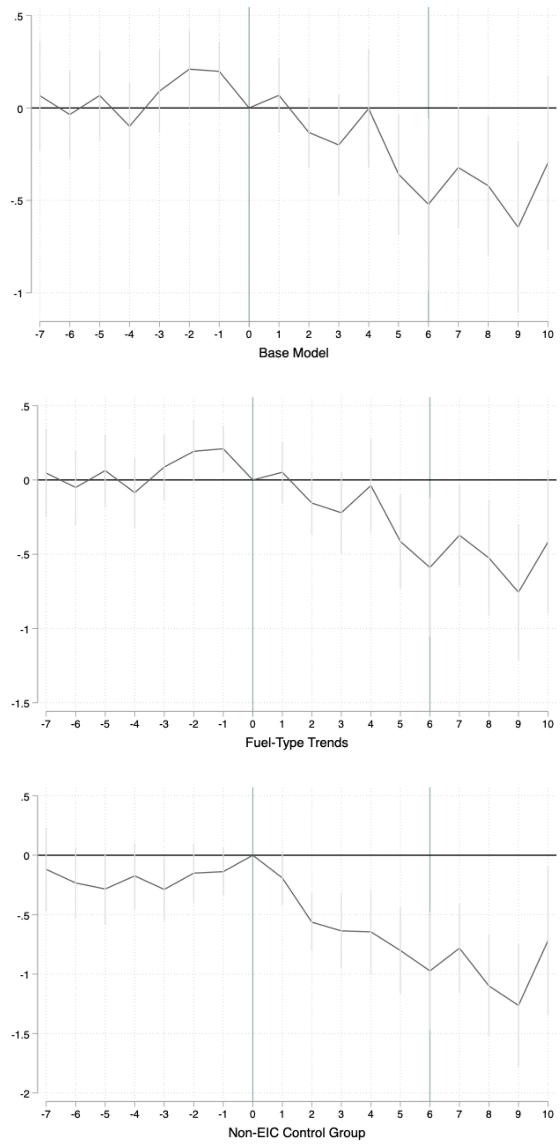

Figure 4: Natural Gas Event Study Estimates

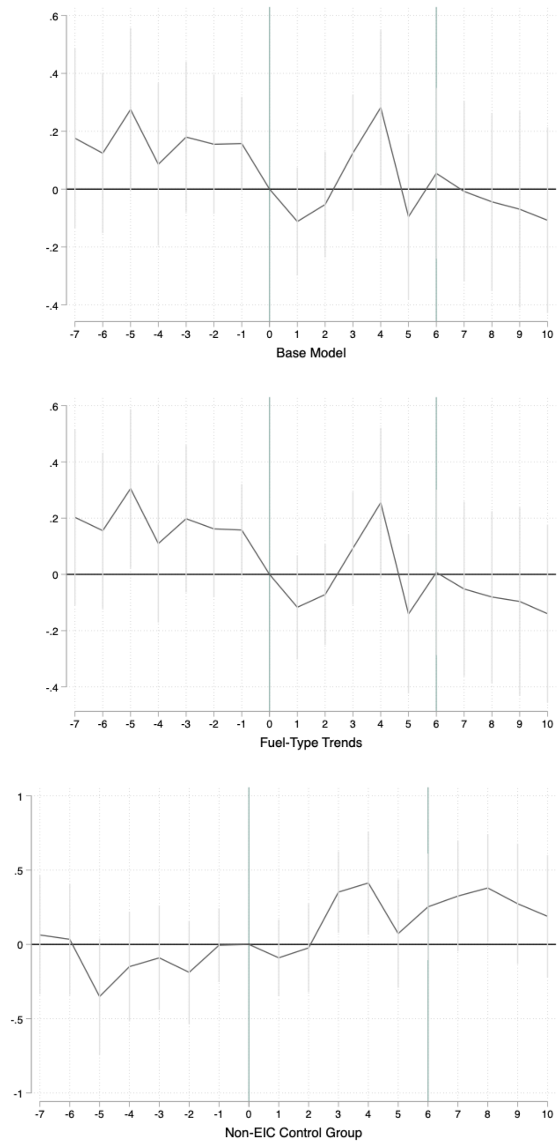

nated. However, the point estimates are stable and statistically zero when only considering plants in states outside the Easter Interconnection as controls. Second, coal emissions reduced more strongly as the cap became more strict. With contaminated controls (top two panels), the early period of RGGI shows to not have been consequential but with a cleaner control group (the bottom panel) the RGGI induced meaningful coal emissions reductions from the beginning. Again, this is to be expected regardless of auction prices relative to the social cost of carbon since requiring the purchase of emissions permits increase the cost of production directly and should influence behavior on the margin – even with relatively low permit prices.

Some similar patterns emerge for natural gas emissions (Figure 4). With contaminated controls (the top two panels), there appears to be a marginal difference in trend in the pre-period but with the cleaner control group (bottom panel) the pre-period trend is more stable and statistically indistinguishable. Furthermore, there are periods of positive and statistically significant increase in natural gas emissions in the bottom panel but not the top two panels. Taken together with Figure (3), this suggests that the RGGI induced not only a reduction in coal emissions but potentially a meaningful substitution in fuel source toward natural gas.

5 Policy Spillovers

It is evident that carbon emissions have fallen at coal-fired power generation plants within the RGGI region. There is some evidence that natural gas-fired plants within the RGGI may have increased emissions in response, and it remains unclear whether or not these emissions from the RGGI region have simply been replaced by emissions produced outside the region as prior work implies might have happened. Here, we use the two treatment period model from before to examine how emissions have changed at the plant and state levels outside of the RGGI region. We also test how emissions of SO_2 and NO_x have changed.

To identify whether or not leakage has occurred we rely on the fact that electricity supply and demand is subject to the physical limitations of transmission. Moreover, we exploit variation in the overlapping network of 'balancing authorities' whose job it is to dispatch and curtail electricity transmission depending on the load demand profile of an area which does not necessarily fall evenly across RGGI-adopting state lines. However, the overlapping grid network and interdependent systems immediately cast doubt on the validity of the identification strategy to measure spillover and we will evaluate this potential further.

We estimate the change in emissions from coal and natural gas-fired plants in non-RGGI states that are part of the PJM balancing region as in Chan and Morrow (2019), as well as in the 'leaker' states of Ohio and Pennsylvania as in Fell and Maniloff (2018) and Yan (2021). These estimates are shown in Table 5 with columns (1) and (2) measuring leakage into 'leaker' states not in the PJM, and columns (3) and (4) measuring leakage into plants that reside in states that are in the PJM region but are not signatories to the RGGI agreement.

Focusing first on the estimated effect within the RGGI region in Table 5, we find the change in plant-level emissions is quantitatively similar to the estimates found before. We see that each phase of the RGGI policy has decreased coal plant emissions but there is no discernible impact on natural gas emission. These estimates are inline with the prior estimates of Tables 3 and 4.

Looking now to carbon emissions leakage, we find evidence that some carbon emissions were pushed out to the surrounding 'leaker' plants and plants that are non-RGGI but in the PJM region. Columns (1) and (2) both consider how emissions have changed at plants in Ohio or Pennsylvania ('leaker' states), but differ in their inclusion of state fixed effects and state trends. Columns (3) and (4) investigate plants that are part of the PJM balancing authority, but that are outside of the RGGI region. In both cases, the point estimates are slightly attenuated when including state trends and fixed effects. We find evidence across all specifications that some emissions have been pushed outside of the RGGI region, and that

Table 5. Plant-level Leakage

	(1)	(2)	(3)	(4)
Dep. Var.	CO_2	CO_2	CO_2	CO_2
RGGI · Coal · Begin	-0.177*	-0.126	-0.187*	-0.141
	(0.097)	(0.102)	(0.097)	(0.103)
RGGI · Coal · Lower	-0.603***	-0.474***	-0.643***	-0.524***
	(0.156)	(0.175)	(0.156)	(0.176)
RGGI · NatGas · Begin	-0.058	-0.092	-0.022	-0.056
	(0.096)	(0.093)	(0.097)	(0.094)
RGGI · NatGas · Lower	-0.082	-0.116	-0.037	-0.075
	(0.125)	(0.128)	(0.125)	(0.128)
Leaker · Coal · Begin	0.073	0.140		
	(0.118)	(0.121)		
Leaker · Coal · Lower	-0.172	-0.073		
	(0.192)	(0.194)		
Leaker · NatGas · Begin	0.331**	0.328**		
	(0.148)	(0.149)		
Leaker · NatGas · Lower	0.663***	0.660***		
	(0.237)	(0.238)		
PJM · Coal · Begin			0.012	0.024
			(0.102)	(0.105)
PJM · Coal · Lower			-0.269*	-0.290*
			(0.150)	(0.157)
PJM · NatGas · Begin			0.511***	0.473***
			(0.117)	(0.117)
PJM · NatGas · Lower			0.730***	0.692***
			(0.183)	(0.184)
Plant FE	Y	Y	Y	Y
Year FE	N	Y	N	Y
State Trends	N	Y	N	Y
R^2	0.945	0.947	0.945	0.947
Obs.	55035	55035	55035	55035

Notes: Clustered robust standard errors shown in parentheses; *,**,*** denote statistical significance at the 10%, 5%, and 1% levels, respectively

this occurs at natural gas-fired facilities.

While uncovering leakage and spillover effects of the RGGI are an important components of evaluating the overall effect of the policy, their are also sources of bias when trying to estimate the causal effect of the policy on plants within the PJM and leaker states similar to our concerns regarding the RGGI. That is, are the control plants/states for PJM and leaker plants appropriate? The models in Table 5 utilize our benchmark specification in equation (2) and all other plants in the data as controls. We have already provided evidence that including plants within the Eastern Interconnection as controls is problematic when estimating the effect of the RGGI. The evidence of spillover effects in Table 5 suggests that these controls may indeed be tainted by the RGGI policy treatment. We also demonstrated earlier that it is the group of plants outside the Eastern Interconnection that exhibits stable and common pre-treatment trends with RGGI plants. To evaluate whether the leaker and PJM estimates from the benchmark model in Table 5 are appropriate, we estimate the modified version of our event study model in equation (3) below:

$$\begin{aligned} CO2_{ist} = \beta_0 + \beta_1 RGGI_s + \sum_{\substack{t=t_0 \\ t\neq 2008}}^{t} \left[\beta_2^t RGGI * I(Year = t)\right] + \\ \beta_3 PJM_s + \sum_{\substack{t=t_0 \\ t\neq 2008}}^{t} \left[\beta_4^t PJM * I(Year = t)\right] + \\ or \\ \beta_5 Leaker_s + \sum_{\substack{t=t_0 \\ t\neq 2008}}^{t} \left[\beta_6^t Leaker * I(Year = t)\right] + \varepsilon_{ist} \end{aligned} \tag{4a}$$

$$\varepsilon_{ist} = \mu_i + \lambda_t + t\mu_s + \nu_{ist} \tag{4b}$$

Figure (5) plots parameter estimates of β_4^t and Figure (6) parameter estimates of β_6^t. In both figures, the top rows are the estimates specific to coal plants and the bottom rows

Figure 5: PJM Event Study Estimates, Coal (top row) Natural Gas (bottom row)

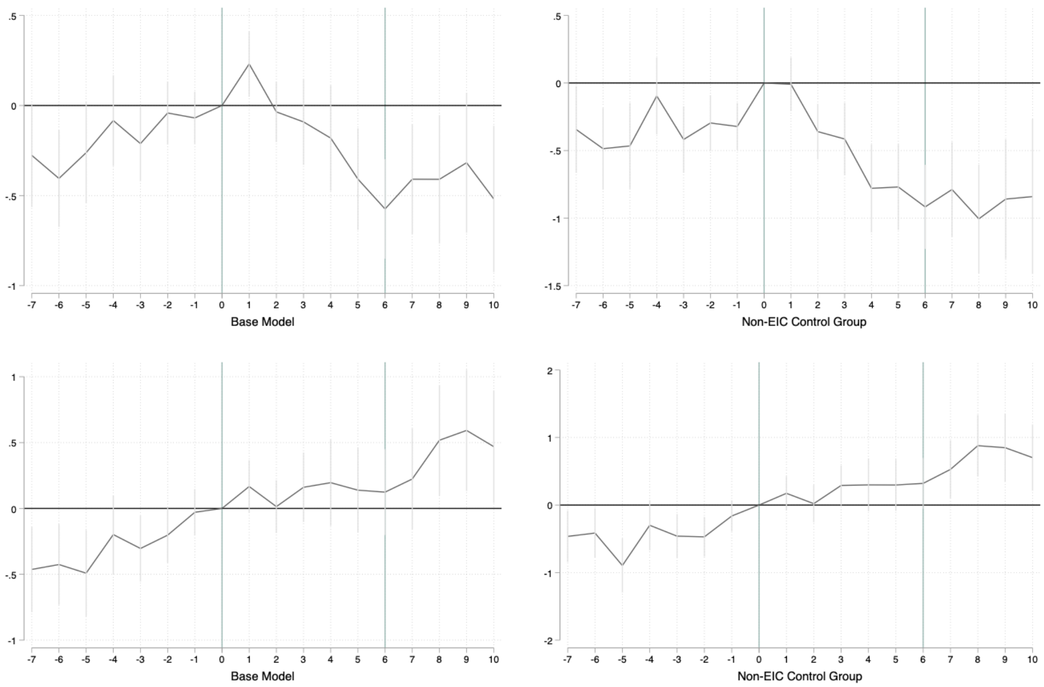

Figure 6: Leaker Event Study Estimates, Coal (top row) Natural Gas (bottom row)

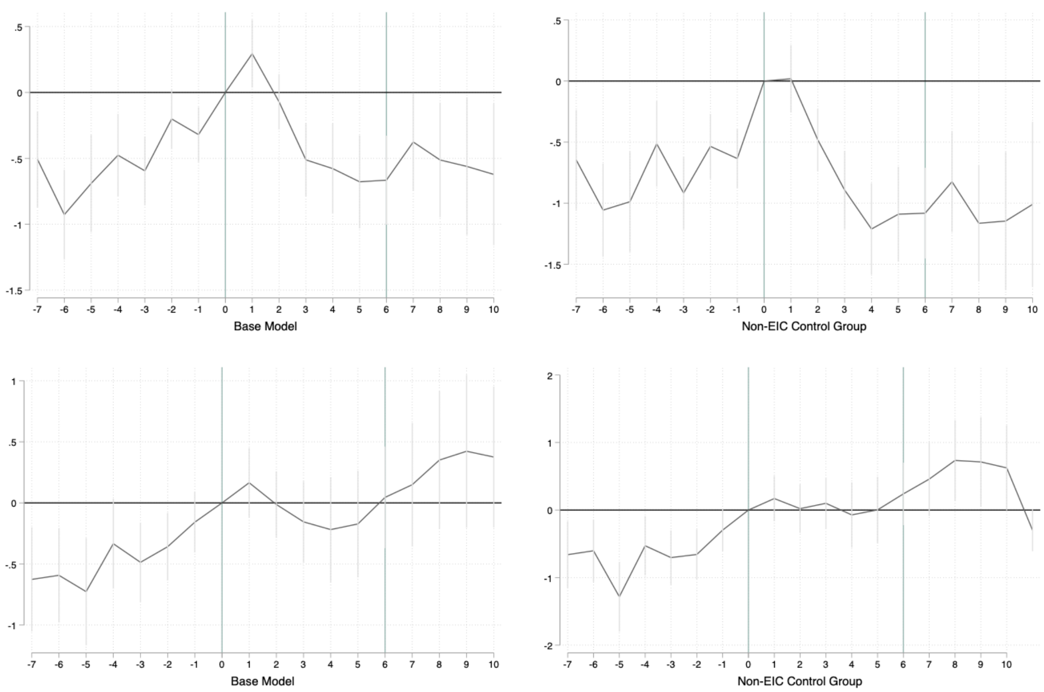

are the estimates specific to natural gas plants. The left two panels in both figures are from models that incorporate all other plants in the control group and the right two panels incorporate only plants *outside* the Eastern Interconnection in the control group.

Figures (5) and (6) are problematic for conclusions regarding leakage in our Table 5 and other findings in the literature (Chan and Morrow (2019); Fell and Maniloff (2018); Yan (2021)). The assumption of common trends in the pre-period fails for both coal and natural gas plants regardless of the control group we utilize. While our earlier results suggest that plants outside the Eastern Interconnection are appropriate controls for RGGI plants, this does not appear to be the case for plants in Leaker states or the PJM. There is something fundamentally different about the evolution of emissions in Leaker and PJM plants even after accounting for unobservables in the model.

5.1 RGGI Impact on Non-regulated Emissions

Finally, to measure the full impact of the RGGI on electric power sector emissions we also investigate whether or not associated emissions have changed as a result of the policy (those that are not regulated under the RGGI carbon permit market). In this section we re-estimate both the plant-level and state-level models with Sulphur Dioxide and Nitrous Oxide emissions as the dependent variable to see if plant-level (Table 6) or state-level (Table A1) emissions have changed for these pollutants.

Beginning at the plant-level, we do find evidence that the RGGI has led to decreases in in SO_2 and NO_x for both coal and natural gas plants but that the relative declines in these pollutants is smaller for natural gas plants once the emissions caps were reduced. This matches the narrative from before that more natural gas is used within the RGGI during the cap-tightening period. As we saw with carbon-leakage, here we again find some evidence of emissions leakage of SO_2 in the non-RGGI PJM territory, though not in Pennsylvania and Ohio. This finding is only statistically significant for natural gas plants in the non-RGGI

Table 6. Plant-level Associated Emissions

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.	SO_2	SO_2	SO_2	NO_x	NO_x	NO_x
RGGI · Coal · Begin	-0.382**	-0.358*	-0.370*	-0.305**	-0.301**	-0.326**
	(0.195)	(0.194)	(0.196)	(0.135)	(0.135)	(0.136)
RGGI · Coal · Lower	-0.985***	-0.967***	-0.975***	-0.703***	-0.716***	-0.746***
	(0.303)	(0.302)	(0.305)	(0.210)	(0.211)	(0.212)
RGGI · NatGas · Begin	-0.628***	-0.629***	-0.606***	-0.281***	-0.264***	-0.233***
	(0.110)	(0.110)	(0.111)	(0.086)	(0.086)	(0.087)
RGGI · NatGas · Lower	-0.720***	-0.697***	-0.659***	-0.218*	-0.180	-0.141
	(0.142)	(0.142)	(0.142)	(0.114)	(0.114)	(0.115)
Leaker · Coal · Begin		0.164			0.051	
		(0.203)			(0.106)	
Leaker · Coal · Lower		0.192			-0.100	
		(0.267)			(0.158)	
Leaker · NatGas · Begin		0.122			0.208*	
		(0.176)			(0.115)	
Leaker · NatGas · Lower		0.284			0.426**	
		(0.222)			(0.172)	
PJM · Coal · Begin			0.068			-0.082
			(0.170)			(0.096)
PJM · Coal · Lower			0.130			-0.197
			(0.216)			(0.135)
PJM · NatGas · Begin			0.192			0.310***
			(0.150)			(0.097)
PJM · NatGas · Lower			0.395**			0.508***
			(0.181)			(0.139)
Plant FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
State Trends	Y	Y	Y	Y	Y	Y
R^2	0.946	0.946	0.946	0.916	0.917	0.917
Obs.	58619	58619	58619	58607	58607	58607

Notes: Clustered robust standard errors shown in parentheses; *,**,*** denote statistical significance at the 10%, 5%, and 1% levels, respectively

PJM territory. For NO_x emissions leakage, we find evidence of an increase in emissions at natural gas plants in both outside regions. Regarding state-level impacts (columns (4) to (9) of Table A1), there is mixed evidence that the RGGI has impacted non-regulated emissions in outside regions. There is some evidence of a decrease in SO_2 emissions in Pennsylvania and Ohio due the introduction of the cap-and-trade system, but these state-level estimates aggregate emissions across plant type and are difficult to interpret as discussed before.

6 Conclusions and Policy Implications

Despite global collaboration efforts and agreements to reduce greenhouse gasses, specific policies intended to reduce greenhouse gas emissions have not been widely applied and are mostly regionally-based. In this paper, we analyze the efficacy of the first cap-and-trade program that was passed in the United States, the Regional Greenhouse Gas Initiative (RGGI), and also discuss broader impacts of the policy such as leakage into neighboring areas and effects on unregulated co-pollutants.

Our paper makes several contributions to the existing literature that has investigated the RGGI and its effects. First, we are able to identify the differential effect of the RGGI on plant-level emissions for both coal and natural gas-fired units. Second, we assess not only the introduction of the RGGI as prior authors have done, but also measure how the tightening of the cap since 2014 has changed emissions. Next, we assess leakage that may have occurred as a result of the policy and carefully evaluate how this may bias our estimates due to contamination of the control group. Lastly, we determine how the RGGI has impacted co-pollutants (such as SO_2 and NO_x) at both the state-level and plant-level cumulatively and by treatment period. We also replicate our main results for emissions aggregated to the state-level for comparison to our plant-level models and previous literature that has focued on state-level emissions.

We find that plant-level emissions have declined in the RGGI region at coal-fired plants, relative to treatment-free comparison plants. Averaged over the entire RGGI policy period, emissions in coal fired plants fell by about 20% per year relative to plants outside the RGGI. This effect is stronger in the period when the carbon cap was lowered (about 40% lower), and is even larger in our models that incorporate a smaller but more plausible control group of plants. For natural gas plants, we find mixed evidence. In some model specifications we see a statistically discernible decrease in emissions, and in models with the most plausible and clean control group we find increases in emission from natural gas plants. Taken together, this mixed evidence is consistent with the lower carbon content and large wholesale price declines for natural gas – both of which stand to actually increase generation at natural gas-fired units even within the RGGI. Concerning the RGGI's total impact on emissions, some of our estimates support prior findings that carbon leakage is occurring, though we also argue and provide support that these spillover effects contaminate the control group and generate bias.

Our results offer insight into how future policies may better address carbon emissions and climate change as well as highlight some of the pitfalls and roadblocks associated with formulating an effective policy. For example, a market-based policy intervention such as the cap and trade RGGI policy (or a carbon tax) appears to be successful at reducing emissions with two caveats: (1) External factors such as the recent decline in natural gas and subsequent substitution away from coal or technological innovation can quickly make existing caps ineffective or inappropriate. Although emissions in coal plants declined almost immediately, the RGGI cap was clearly more binding after the 2014 cap reduction as electricity generation was already moving away from coal to cheaper natural gas. However, in the case that natural gas prices would have increased rather than decreased, the caps could have been too restrictive too quickly. It is important for future policies to be both flexible to change and responsive to external factors in the marketplace. Similar evidence is uncovered in the transportation sector in Andersson (2019) who shows that drivers respond more strongly to changes in carbon tax rates than equivalent (market determined) gasoline price changes. In our setting, it is possible that price movements due to changes in market forces are perceived differently by electricity generating firms than equivalent carbon price changes. (2) Addressing climate change with policies that are restrictive to geopolitical boundaries will be less effective due to spillovers and the exportation of emissions. Domestically proposed climate change policies receive significant political push-back in part for this reason as well as the differential effect on competitiveness for domestic firms. A broad reaching policy crossing political boundaries is clearly the more ideal approach, particularly if it should have meaningful and identifiable enforcement mechanisms.

7 Appendix

Table A1. State-level Full Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dep Var.	CO_2	CO_2	CO_2	SO_2	SO_2	SO_2	NO_x	NO_x	NO_x
RGGI · Begin	-0.070	-0.067	-0.066	-0.513**	-0.519**	-0.511**	-0.007	-0.009	-0.012
	(0.088)	(0.088)	(0.088)	(0.221)	(0.222)	(0.219)	(0.121)	(0.122)	(0.121)
RGGI · Lower	-0.075	-0.077	-0.074	-0.694*	-0.710*	-0.716*	-0.103	-0.110	-0.082
	(0.178)	(0.179)	(0.179)	(0.401)	(0.402)	(0.400)	(0.203)	(0.203)	(0.204)
Leaker · Begin		0.058**			-0.123			-0.050	
		(0.025)			(0.122)			(0.182)	
Leaker · Lower		-0.054**			-0.335**			-0.150	
		(0.026)			(0.127)			(0.099)	
PJM · Begin			0.034			0.085			-0.105
			(0.101)			(0.335)			(0.148)
PJM · Lower			0.018			-0.089			0.066
			(0.105)			(0.269)			(0.130)
State FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
State Trends	Y	Y	Y	Y	Y	Y	Y	Y	Y
R^2	0.994	0.994	0.994	0.974	0.974	0.974	0.983	0.983	0.983
Obs	867	867	867	867	867	867	867	867	867

Notes: Clustered robust standard errors shown in parentheses; *,**,*** denote statistical significance at the 10%, 5%, and 1% levels, respectively

References

Andersson, J. J. (2019). Carbon Taxes and CO_2 Emissions: Sweden as a Case Study, *American Economic Journal: Economic Policy* **11**(4): 1–30.
URL: *https://www.aeaweb.org/articles?id=10.1257/pol.20170144*

Callaway, B. and SantAnna, P. H. C. (2020). Difference-in-Differences with multiple time periods, *Journal of Econometrics* .
URL: *https://www.sciencedirect.com/science/article/pii/S0304407620303948*

Chan, N. W. and Morrow, J. W. (2019). Unintended consequences of cap-and-trade? Evidence from the Regional Greenhouse Gas Initiative, *Energy Economics* **80**: 411–422.
URL: *http://www.sciencedirect.com/science/article/pii/S0140988319300222*

Chen, Y. (2009). Does a regional greenhouse gas policy make sense? A case study of carbon leakage and emissions spillover, *Energy Economics* **31**(5): 667–675.
URL: *http://www.sciencedirect.com/science/article/pii/S0140988309000309*

Fell, H. and Maniloff, P. (2018). Leakage in regional environmental policy: The case of the regional greenhouse gas initiative, *Journal of Environmental Economics and Management* **87**: 1–23.
URL: *http://www.sciencedirect.com/science/article/pii/S0095069616302984*

Fullerton, D. and Karney, D. H. (2018). Multiple pollutants, co-benefits, and suboptimal environmental policies, *Journal of Environmental Economics and Management* **87**: 52–71.
URL: *http://www.sciencedirect.com/science/article/pii/S0095069617305764*

Goodman-Bacon, A. (2018). Difference-in-Differences with Variation in Treatment Timing, *Technical Report w25018*, National Bureau of Economic Research.
URL: *https://www.nber.org/papers/w25018*

Haites, E. (2018). Carbon taxes and greenhouse gas emissions trading systems: what have we learned?, *Climate Policy* **18**(8): 955–966.
URL: *https://doi.org/10.1080/14693062.2018.1492897*

Kim, M.-K. and Kim, T. (2016). Estimating impact of regional greenhouse gas initiative on coal to gas switching using synthetic control methods, *Energy Economics* **59**: 328–335.
URL: *http://www.sciencedirect.com/science/article/pii/S0140988316302316*

Lee, K. and Melstrom, R. T. (2018). Evidence of increased electricity influx following the regional greenhouse gas initiative, *Energy Economics* **76**: 127–135.
URL: *http://www.sciencedirect.com/science/article/pii/S0140988318303992*

Murray, B. C. and Maniloff, P. T. (2015). Why have greenhouse emissions in RGGI states declined? An econometric attribution to economic, energy market, and policy factors, *Energy Economics* **51**: 581–589.
URL: *http://www.sciencedirect.com/science/article/pii/S0140988315002273*

RGGI (2016). RGGI_proceeds_report_2014.pdf, *Technical report*, The Regional Greenhouse Gas Initiative.
URL: *https://www.rggi.org/sites/default/files/Uploads/Proceeds/*$RGGI_Proceeds_Report_2014$*.pdf*

Yan, J. (2021). The impact of climate policy on fossil fuel consumption: Evidence from the Regional Greenhouse Gas Initiative (RGGI), *Energy Economics* p. 105333.
URL: *https://www.sciencedirect.com/science/article/pii/S0140988321002395*