The Effect of California's Carbon Cap and Trade Program on Co-pollutants and Environmental Justice: Evidence from the Electricity Sector

Ryan Walch[1]
Department of Economics
University of Oregon
rwalch@uoregon.edu

November 1st, 2018
Click Here For Most Recent Version

Abstract: Carbon dioxide emissions are globally uniformly mixing pollutants, but the same industrial processes that produce carbon emissions typically also produce co-pollutants, such as NO_x and SO_x. Environmental Justice advocates have expressed concern that California's cap-and-trade program may cause low-income and minority communities to experience greater exposures to these co-pollutants. I use emissions data from almost all power plants in the United States, and a variety of strategies for constructing counterfactual emissions for plants covered by the California program (including a semi-parametric matching estimator and a synthetic control design). None of these methods suggests that California's carbon cap-and-trade program has increased average co-pollutant emissions. If anything, average co-pollutant emissions may have decreased. From the EJ perspective, average co-pollutant emissions at plants located in low-income or minority communities covered by the program have not gone up relative to co-pollutant emissions at plants in similar communities outside of California.

Keywords: Cap-and-Trade, Carbon Pricing, Environmental Justice, Co-pollutants
JEL Codes: Q48, Q52, Q53, Q54

[1] *Acknowledgments*: I would like to thank Trudy Cameron, Grant McDermott and Ben Hansen for helpful guidance and comments throughout the writing of this paper. All remaining errors are my own.

1 Introduction

Cap-and-trade programs have become a popular tool for policy makers who wish to take steps to mitigate climate change by reducing carbon emissions. This popularity reflects the fact that cap-and-trade achieves a fixed level of abatement at minimum cost. Carbon cap-and-trade programs exist in the European Union, California, and Quebec, and are scheduled to begin in China.

Environmental justice groups in both California and other jurisdictions have expressed concerns that cap-and-trade programs may increase pollution levels in disadvantaged communities. In such a case, the aggregate net welfare effects from the policy may be positive, but some of the distributional consequences may be regressive.[2] Carbon dioxide, itself, is a uniformly mixing pollutant and therefore the spatial distribution of abatement actions, ultimately, has no relation to the distribution of benefits from the climate change mitigation objective of the program. However, pollutants that often co-occur with carbon dioxide, such as NO_x and SO_x, do not mix uniformly in the atmosphere. Any redistribution of carbon emissions as a result of a cap-and-trade program may, in fact, have concurrent effects on NO_x and SO_x emissions, and thus have local effects on population exposures. It is possible for carbon pricing to alter the spatial distribution of co-pollutant damages by changing the spatial distribution of economic activity.

There is a growing literature on the interaction between carbon pricing and co-polluants (e.g. Muller (2012), Agee et al. (2014), Fullerton and Muehlegger (2017) Novan (2017)). The theoretical relationship between co-pollutants and carbon dioxide is ambiguous. NO_x and SO_x may be either complements or substitutes, relative to carbon dioxide, which means that carbon pricing may reduce or increase co-pollutant levels. Additionally, carbon pricing is frequently enacted in a setting where regulations on co-pollutants do not properly reflect the full damages of their emissions. This incomplete regulation presents an additional challenge in predicting the effects of carbon-pricing on co-pollutants and creates the possibility of adverse interactions between carbon pricing and co-pollutant emissions.

In the debate over the renewal of California's greenhouse-gas permit-trading program, environmental justice (EJ) concerns have created opposition to the program from several environmental groups such as the Sierra Club and the California Environmental Justice Alliance. EJ concerns about the program are not limited solely to non-profit advocacy groups. Such concerns have also been expressed during the California Air Resource Board's (CARB) formal rulemaking process by the board's official Environmental Jus-

[2]See: Barboza and Megerian (2017), Guerin (2017), Geuss (2017), Kahn (2016), Mason et al. (2017), Climate Hawks Vote (2017)

tice Advisory Committee. In a meeting on February 15th, 2017, the advisory committee issued a statement criticizing cap-and-trade, stating that the program “does not reflect best practices in research or serve the interests of poor communities, communities of color, and indigenous communities in California and around the world.” (California Air Resource Board (2017))

EJ groups have argued that the ability of firms to reallocate carbon emissions from plant to plant (by transferring permits) will result in higher levels of pollution in minority and low-income communities. Such concerns are based on the perception among the EJ community that dirtier plants, which are disproportionately located in low-income and minority communities, will increase their emissions when allowed to buy permits, whereas cleaner plants will respond by lowering their emissions. These concerns tend to be voiced by non-economists, and thus often do not contain references to formal economic logic. Such a focus on (absolutely) dirtier plants ignores the important role of *marginal* abatement costs in determining changes in pollution levels after the introduction of an emissions trading program.[3]

Although the arguments by environmental justice groups are not grounded upon formal economic theory, regressive distributional effects could occur if high-abatement-cost firms are located in disadvantaged communities. According to economic theory, it is the spatial distribution of marginal abatement costs among plants that will determine the changes in pollutants as a result of the program. Low-marginal-abatement-cost firms will have an incentive to lower their emissions to sell permits to high-marginal-abatement-cost firms. Thus the core distributional concerns of the EJ groups could be valid if firms with relatively high marginal abatement costs are more likely to be located in disadvantaged communities.[4]

Instead of market-based methods of regulation, EJ advocacy groups argue for command-and-control regulation, including per-facility carbon dioxide emission limits in addition to technology standards. Similar EJ concerns resulted in strong opposition from environmental groups when the state of Washington included a carbon tax referendum bill in the 2016 election and, the concerns, contributed to the failure of the referendum.

California’s policy makers are required to consider the EJ impacts of both the current cap-and-trade program and any future programs. By law, all state government organizations must ensure that environmental regulation does not systematically harm individuals on the basis or race of income. U.S. federal regulations have also long required policy makers to consider the distributional impacts of environmental regulation. Executive Or-

[3]See Farber (2012) for a detailed discussion of the EJ arguments made against cap-and-trade.

[4]Section 2 discusses these arguments and the corresponding economic reasoning in more detail.

der 12898, issued in 1994, requires that the EJ impacts of all environmental regulations be considered when evaluating policy. An understanding of the distributional impacts of cap-and-trade programs is therefore important to policy-makers both in California, and for any future state or federal carbon pricing programs. Even though the economic basis for these advocacy group arguments is not always clear, and the implicit assumptions may not demonstratively hold, the question of the EJ impacts of cap-and-trade programs has become central to both the political economy of carbon pricing, and the legal obligations of policy makers.

There is extensive evidence that existing levels of pollution are often higher in low-income and minority communities than in other types of communities. However, previous work on emission markets has found little evidence of systematic differences in policy-induced abatement levels with respect to spatial variations in race or income. Most notably, Fowlie et al. (2012) study the effects of the Southern California NO_x emissions trading program (RECLAIM). They use a matched difference-in-differences estimator and fail to reject the null hypothesis of no systematic variation in program benefits according to either the racial composition or the income levels of the affected communities. However, the task of quantifying the change in the spatial distribution of co-pollutants that occur after the introduction of a *carbon* emission-permit market has received relatively little attention in the literature. While portions of my empirical approach will follow Fowlie et al. (2012), the setting for their paper differs in important ways from the setting for my paper. Fowlie et al. (2012) consider a permit market program which directly regulates a harmful local pollutant whereas my paper addresses the effects of a permit market on pollutants that are not directly covered by the market in question but for which emission levels are none the less correlated with emissions of the regulated pollutant. There is no theoretical justification to believe that the pattern of abatement across areas with different demographics or income levels will the same in the carbon cap-and-trade program as it was in RECLAIM.

I utilize a dataset of hourly plant-level emissions for all power plants with more than 25 megawatts of capacity across the United States. I observe emissions of CO_2, NO_x, and SO_x. I combine these emissions data with demographic data from the American Community Survey conducted by the U.S. Census, as well as EPA data on the characteristics of individual power plants. My data set allows me to use non-California entities to help control for unobserved region-wide shocks that differentially affect communities based on their income levels or racial composition. Controlling for these region-wide shocks is not possible in studies that rely solely on administrative data from California's cap-and-trade program.

My empirical strategy first utilizes a semi-parametric matched difference-in-difference estimator to construct a control group for each regulated plant in California. There are many strategies for matching. My initial approach is analogous to that of Fowlie et al. (2012). I match each treated unit to the closest M controls based on their distance in covariate space as defined by the Mahalanobis norm. This and similar estimators are discussed in Heckman et al. (1997), Heckman et al. (1998), Abadie and Imbens (2006), Abadie and Imbens (2011) and Haninger et al. (2017). This matching method allows for more flexibility in constructing counterfactual values than parametric methods, and limits the influence of non-similar control plants. This method also allows me to construct heterogeneous treatment effects which vary systematically with the demographics of nearby communities, permitting a direct test for any adverse environmental justice outcomes of California's cap-and-trade program.

As a second approach, I make use of the synthetic control method developed by Abadie et al. (2010). This method constructs the counterfactual outcomes for California emissions from the linear combination of control-state emissions that best tracks California pre-treatment emissions. While I cannot directly compute heterogeneous treatment effects for each individual plant, I can compare estimates where the sample is restricted to low-income or high-minority-share communities to the results estimated on the full sample.

My results suggest that, on average, California electricity plants saw a reduction in co-pollutant emissions due to the carbon cap-and-trade program. The sign and magnitude of the key coefficient is negative regardless of the specification of the control group and for both methods, but it is not statistically significant for all possible choices of matched control groups or for the synthetic control results. Importantly, I find no robust evidence that this effect varies with either the income or the racial composition of the communities surrounding the plant. Thus, there is no compelling evidence for adverse environmental justice impacts for co-pollutants in low-income or high-minority-share communities in California.

Two previous papers have made attempts to characterize the distribution of gains across demographics for California's carbon cap-and-trade program. Cushing et al. (2016) present statistics from Californian administrative data showing that regulated facilities are more likely to be located in low-income or minority communities, and that several industries have experienced increases in both their carbon emissions and their co-pollutant emissions. However, the Cushing et al. paper consists mainly of summary statistics and includes no formal statistical analysis. No attempt is made to control for unobserved heterogenity, or to estimate a causal effect for the program.

Meng (2017) uses CARB administrative data to estimate a difference-in-differences model to assess the levels of carbon abatement across advantaged and disadvantaged communities. These administrative data consist of carbon emissions reported to the state of California to document compliance with the cap-and-trade program. Meng finds no evidence that carbon abatement varies systemically with race or income. Meng also finds suggestive evidence of perhaps *more* abatement in low income and minority communities, but the estimates are not statistically significant at conventional levels.

My approach differs from Meng (2017) in several ways. Meng's data allow him to see the full universe of entities regulated under California's cap-and-trade program, whereas my data are limited to electricity-generating firms. However, the data used in my paper have two distinct advantages relative to Meng's. First, I can directly observe co-pollutant levels, whereas Meng can observe only carbon emissions and has no data on co-pollutant emissions. Second, all the firms in Meng's dataset are located only in California. Thus, his identification strategy cannot control for national or regional trends unrelated to cap-and-trade that differentially affected emissions in advantaged and disadvantaged communities. The data used in my paper are for the entire U.S. so I can use patterns of emissions for various sets of matched-non-California firms as controls.

The paper is organized as follows. Section 2 outlines the institutional background of the California cap-and-trade program. Section 3 explains the data and methodology. Section 4 presents the results. Section 5 discusses some limitations of the data and the analysis, and proposes some additional research that may be appropriate, as more data accumulates in the coming years and as firms reoptimize their capital stocks over the longer run. Section 6 concludes.

2 Program Background

In 2006 the California legislature passed Assembly Bill 32 (AB32). The law mandated a reduction in carbon dioxide emissions to 1990 levels by 2020. To meet these goals, California chose to establish a cap-and-trade program. California also adopted a low-carbon fuel standard, implemented energy efficiency regulations, and required electrical utilities to obtain more of their electricity from renewable sources.

In a cap-and-trade program, each firm must surrender a permit for each ton of carbon dioxide that it emits. The total quantity of permits is capped and firms are allowed to buy and sell permits for cash payments. A cap-and-trade program achieves a particular level of abatement at least cost because it allows firms with higher marginal abatement costs to "bargain" with other firms with lower marginal abatement costs to reduce their emissions instead. Thus the equilibrium pattern of emissions is determined by the distribution of marginal abatement costs across firms. Firms with the lowest marginal abatement costs will typically do the most abatement, freeing up permits for sale to other firms that have marginal abatement costs higher than the market price of a permit.

California's cap-and-trade program began in 2013. The cap was initially set at two percent below 2012 emissions, declined two percent in 2014 and was scheduled to decline by three percent in each subsequent year (until 2020). The state estimates that the required decline in emissions represents a 15 percent reduction from the counterfactual "no program" trend in emissions. The AB32 program covers carbon dioxide as well as several other greenhouse gases.[5]

All electricity producers in California are covered by the program as well as all large industrial sources emitting more than 25,000 megatons of CO_2-equivalent emissions per year. Fuel suppliers were brought under the cap in 2015. Around 450 entities, in total, were covered by the program as of 2015. CARB estimates that eighty percent of all Californian carbon emissions are subject to the cap.[6] Permit allocations to large industrial emitters were initially distributed at no cost to firms, based on the firm's historical emissions and the firm's energy efficiency, but an increasing proportion of permits will be auctioned as time goes on.[7] Electricity generators received free permits on the condition that all profits from the permits must benefit utility rate-payers. Permits may be banked, and firms may meet part of their compliance obligation by purchasing "offsets" that support other types of approved carbon-dioxide-reducing projects.[8] Permits were initially traded at $22 per ton of CO_2-equivalent. However, after some volatility, equilibrium prices fell and eventually settled around a price of $12 to $13 a ton by 2014.

As noted in the introduction, the spatial distribution of carbon abatement activity does not affect the distribution of global benefits from carbon emissions reductions. Carbon dioxide is considered to be "globally uniformly mixing," meaning that it spreads evenly throughout the earth's atmosphere. It is the global concentration of carbon dioxide that determines the pace of global warming. Consequently, it does not matter which firm(s) choose to abate their carbon emissions; all that matters for climate change is the aggregate abatement. This feature of carbon-emissions means that there can be opportunities to decrease the overall cost of regulation by facilitating the allocation of abatement responsi-

[5]The full list includes carbon dioxide (CO_2), methane (CH_4), nitrous oxide (NO_2), hydrofluorocarbons (HFCs), perfluorocarbons (PFCs), sulfur hexafluoride (SF_6), nitrogen trifloride (NF_3).

[6]Sources of greenhouse emissions that remain outside the cap include agriculture and emissions from residential and commercial sources.

[7]Increasing auction shares reflect the transition of property rights (to carbon emissions) from firms to the general population.

[8]Common offset projects including preserving or planting forests and disposing of certain types of ozone-depleting gases which also have greenhouse effects. These projects may be located outside of the state.

bility towards cost-minimizing locations.

The expected spatial pattern of abatement under a cap-and-trade program is determined by the spatial pattern in marginal abatement costs. Environmental justice concerns would be warranted if plants with higher marginal abatement costs are located in disadvantaged communities, and if changes in carbon dioxide emissions are strongly correlated with changes in local co-pollutants. However, if marginal abatement costs are uncorrelated with characteristics of the surrounding community, then it is unlikely that changes in the location of carbon emissions, attributable to cap-and-trade, will differentially change co-pollutants for low-income and minority communities.

Theory is ambiguous on which case will hold. There are three possibilities: (1) dirtier plants located in disadvantaged communities may not yet have taken full advantage of all available abatement technologies, implying that marginal abatement costs for these firms could be lower than for cleaner state-of-the-art plants located in wealthier non-minority communities; (2) dirtier plants located in low-income communities have higher emissions because they tend to have higher marginal abatement costs (often implicitly assumed by those who oppose cap-and-trade programs on EJ grounds); (3) there is simply no correlation between marginal abatement costs and the low-income and minority-shares of surrounding communities. Unfortunately, it is not possible to observe marginal abatement costs directly in the available data and possible proxies are insufficiently informative, so it is not possible to simply observe which of (1) through (3) hold.

These environmental justice concerns have taken a central role in the debate about whether to renew California's cap-and-trade program after the expiration of AB32 in 2020. These concerns can also be seen explicitly in the legislation passed. The process to extend California's carbon cap-and-trade beyond 2020 began with the passage of a separate California Senate bill, SB32, in 2016, which mandated a forty percent reduction in carbon emissions below 1990 emission levels by 2030. In response to widespread EJ concerns, SB32 explicitly mandates that the emission reductions must be achieved in a "manner that benefits the state's most disadvantaged communities." AB 398, which was passed in 2017, established a more aggressive cap-and-trade system to achive these reductions.

In addition to the cap-and-trade "extension" in AB 398, a separate bill was passed (in conjunction with the extension) that mandates stricter regulation of local pollutants. As a condition for their support of the bill, industry groups demanded that no new GHG regulations could subsequently target entities already participating in the cap-and-trade program. This final condition was viewed as an attempt to forestall any traditional command-and-control regulations, such as plant-specific abatement targets or emissions limits, which had become a popular policy proposal in environmental justice circles to address their distributional concerns.

Thus the distributional effects of changes in co-pollutants has become an important part of both the political economy and legal obligations of California policy makers. A better understanding of the pattern of abatement due to California's cap-and-trade program is therefore of first-order importance to policy-makers.

Another factor that could affect the proper function of California's carbon market is the potential for out-of-state or out-of-country "leakages" of carbon emissions. Abatement of carbon emissions in California could be at least partially undone by increases in emissions outside the state, because carbon pricing would make production outside of California relatively more profitable. Specifically to deter leakage, California freely allocates a portion of the total number of permits based on a firm's output and its efficiency relative to the industry. These criteria act as an output subsidy and encourage firms (and production) to stay in California instead of moving to an unregulated state.[9] Furthermore, California directly taxes imports of electricity from other states.

3 Data and Empirical Strategy

My dataset consists of plant-level emissions from 2010 to 2016 for almost all power plants in the continental United States. Emissions data can be retrieved from the EPA's Clean Air Market Data (CAMD) which includes all generators with a capacity greater than 25 MWh. These data are collected from continuous emission monitoring systems (CEMS) which record emissions data at an hourly frequency. All units must report CO_2, NO_x and SO_x emissions data. These emissions are flows measured at the point of emission from the plant and do not represent readings or imputations of concentrations from ambient pollution monitors. Wind speed and direction, temperature, precipitation and atmospheric chemistry, for example, will all affect the "fate and transport" of these emissions and eventual exposure of the population to the resulting ambient levels of pollution.

Data on the characteristics of electricity generators have been retrieved from the EPA's Emission and Generation Resource Integrated Database (eGRID). eGRID is an extensive database on both the environmental and the technical characteristics of U.S. electricity generators. eGRID is published every other year. In years without eGRID data, I assign plants the characteristics contained in the eGRID release from the previous year. Plant characteristics that I use in this analysis include the primary fuel type of the plant, the nameplate capacity, annual net generation, and the heat rate. Nameplate capacity refers

[9]For a model of the role of output subsidies in preventing leakages see Fischer and Fox (2012)

to the amount of electricity an electrical plant can generate in a given period of time (usually in megawatts per hour) and is a proxy for plant size. Heat rate is a measure of the plant's efficiency. It reflects how much energy is needed to generate one unit of electricity. Plants with *lower* heat rates are more efficient. In addition, latitude and longitude coordinates for the plant's location are drawn from eGRID.

My demographic data are drawn from the five-year moving average of the American Community Survey (ACS) data at the census-tract level.[10] To study the environmental justice impacts of the program, I focus on two variables: per-capita income and the proportion of the population which belongs to a minority group.[11] Demographic data for the neighborhood surrounding each plant are based on all census tracts which intersect a one-mile buffer centered on the plant's latitude/longitude location.

There are two natural control groups to consider for the California plants: the North American Electric Reliability Corporation's (NERC) Western Electricity Coordinating Council (WECC)[12] and the entire set of non-California power plants in the United States. The WECC enforces many federal regulations and writes rules to ensure power plant compliance across its region. The regulations are designed to ensure equal access to transmission infrastructure and to minimize the chance of a wide-scale power failure. In addition, the WECC overlaps with the Western Interconnection. The Western Interconnection is one of three grid interconnections in the US and covers the portion of the United States that lies west of the Rocky Mountains.[13] Technological constraints make it difficult for plants inside the Western Interconnection to transmit power to consumers outside of the interconnection and for firms outside the interconnection to transmit power in.[14] Thus plant inside the interconnection form a market, and experience similar market conditions and regulations.[15] By limiting the sample to the WECC it is less likely that unobserved demand shocks across regions will bias the results.

[10] The values I use reflect the terminal year of the five-year window, not the midpoint. This is a necessary compromise because midpoint data for 2016 will not be available until the ACS data for 2018 are available. I also cannot use earlier years because the five-year ACS only became available in 2010.

[11] Following Fowlie et al. (2012) I use the proportion of residents who identify as African-American or Hispanic for the minority variable.

[12] The NERC is a non-profit collective of electricity-generating firms charged by the federal government with (a) ensuring reliability throughout the grid by ensuring compliance with federal regulations, and (b) collecting data.

[13] Non-California states included in my sample are Arizona, Colorado, Idaho, Montana, New Mexico, Oregon, Utah, Washington, and Wyoming.

[14] In recent years there has been a push to better integrate the different interconnections. However, these projects are either in their early phases or relatively small.

[15] Note that I use a more-general definition of "market" here, where a market is the set of firms and buyers whose actions influence the price. Often when people speak of electricity markets, they are referring to a wholesale market administered by an Independent System Operator(ISO).

The WECC, however, is an imperfect control group due to the limited number of plants that are available for use with the matching estimator as potential controls. Matches may therefore be of lower quality, in the sense that the control plants may differ in covariates. The estimator may incorrectly estimate the counterfactual for the treated plant. Expanding the control group to include all U.S. plants outside of California could allow a greater greater pool of potential controls, thus increasing the chance that there are good matches with similar covariates for each treated unit. Additionally, a greater number of plants will provide more statistical power and decrease finite sample bias.

Given that there are plausible arguments for the choice of either of these control groups, I will show estimates using both of these groups. Qualitatively, the signs and magnitudes of the key parameter estimates do not change with the choice of control groups. The difference is in the sizes of the standard errors and therefore in the statistical significance of the results.

I aggregate the hourly raw emissions data for each plant to the cumulative yearly level for that plant because demographic data are only available on the annual level. The unit of observation is thus the plant-year. I merge the eGRID data on plant characteristics with the CAMD emissions data using the Office of Regulatory Information System PLant (ORISPL) codes. Plants without eGRID information must be dropped from the sample. Given that California power plants primarily use natural gas instead of coal or oil, I limit the sample to plants for which primary fuel type is natural gas. Additionally, I balance the panel, by dropping all plants that do not report emissions for the full seven years of the sample.[16] This leaves me with a total of 69 eligible control plants from the WECC and 662 control plants from the entire United States. There are 86 treated plants in the California sample.[17]

Table 1 shows summary statistics for California and the two candidate control groups. California plants as a group are different than the control plants. They tend to be cleaner and smaller. The communities surrounding the California plants are more diverse and have more income than the communities surrounding the rest of the WECC plants. However, average income levels in communities surrounding plants across the entire U.S. outside of California do not differ from income levels in the communities surrounding California

[16]Generally this is because a plant has shutdown. When the plant is in California there is a potential that this is a result of the cap-and-trade program. If that is the case, this fact will bias my results towards zero, suggesting smaller abatement then actually occurred.

[17]In total I exclude 1,768 observations from 368 plants from the overall sample. The plants I exclude tend to be smaller than the plants included in the sample. Among the controls, the excluded plants tend to be larger emitters although the plants excluded from the California sample emit at levels similar to the included California plants. Efficiency measures between dropped and included plants are similar across both the California and control sample.

plants.

Figure 1 plots trends in the emissions for the treatment and control groups for both NO_x and SO_x and for both the regional and national-level control groups. WECC control plants for NO_x and the national level controls for SO_x seem to provide somewhat plausible control groups as it is arguable that the assumption of parallel prior trends seems to hold. For the other two sets of trends, it seems like the parallel trends assumption is violated.

The most simplistic approach to answering the research question poised in this paper would be to compare the mean emissions before and after the program. Many of the EJ groups concerned about cap-and-trade in California implicitly make such an argument and cite research such as Cushing et al. (2016) that follows this method. If I were to replicate this approach with my data, I would find a statistically insignificant decrease of 10.2 tons a year in NO_X and a statistically significant increase of .062 tons per year for SO_x. However, there are major concerns about the validity of this approach. It is impossible to separate the effect of the program from changes in co-pollutant levels that would have occurred anyway. To get proper estimates of the program's *causal* impact, we need to find a proper control group that would allow us to estimate what would have happened at the California plant's under the no-program counterfactual.

The most natural way to construct a counterfactual would be to estimate a simple difference-in-differences model using plants outside of California as a control group. However, the differing pre-trends should lead us to approach these simple difference-in-differences results with caution. Table 2 shows estimates from such a difference-in-differences model with the WECC plants acting as the control group. Rudimentary difference-in-difference specifications imply that the California carbon program caused co-pollutants to increase. The estimates from these specifications are implausibly large, implying an increase of around 3 standard deviations for NO_x and close to 100 standard deviations for SO_x. The inclusion of state-by-year time trends causes the results to lose statistical significance. This suggests that the simple difference-in-differences results (showing an increase in co-pollutants) are *not* properly controlling for economic shocks that differentially impact the treatment and control groups, and that the model does a poor job of adjusting for differences in covariates across treatment and control groups. Even in the most defensible specification, including state-specific time trends, the standard errors are large and the thus the parameters are very imprecisely estimated. A more sophisticated method of selecting the control group could both reduce bias and shrink standard errors by controlling better for differing pre-trends.

To construct a more-valid control group, I turn to a nearest-neighbor matched difference-and-difference estimator. There is a large literature concerning the properties and implementation of this matching estimator (Heckman et al. (1997), Heckman et al. (1998), Abadie (2005), Abadie and Imbens (2006)m Abadie and Imbens (2011), Fowlie et al. (2012)). The estimator matches treated units with their nearest neighbors, where "distance" is decided by similarity in covariates. By limiting the control group to the nearest neighbors of the treated plants, dissimilar control plants that may bias the results are removed from the estimation. The advantage of a semi-parametric approach, compared to the standard difference-in-difference model, is that it allows more flexibility in terms of the functional form of the relationship between the treated and control groups when estimating treatment effects.

Following the potential outcome framework from Rubin (1973), suppose that each plant has two potential levels of emissions based on its inclusion in a carbon pricing program like the one in California. Let $Y_{it}(1)$ be the emissions from plant i in time period t under carbon pricing and $Y_{it}(0)$ represent the emissions under the no-carbon-pricing counterfactual. Let $D_i = 1$ if plant i was actually subject to the California cap-and-trade and $D_i = 0$ if it was not. I wish to estimate the average treatment effect on the treated (ATT)

$$ATT = E[Y_{i,t}(1) - Y_{i,t}(0)|D_i = 1] \tag{1}$$

For the California plants, the econometrican observes only $Y_{it}(1)$ The challenge is to find a consistent estimator of $Y_{it}(0)$. The semi-parameteric matching estimator uses the M nearest neighbors in covariate space to estimate $Y_{i,t}(0)$. In other words, the estimator constructs a control group for each Californian plant. Distance between plants is measured by the Mahalanobis norm, which scales the difference in each plant attribute by the standard deviation of that attribute. This means a higher penalty is assigned to plants which differ in attributes that do not have much variation than to plants that differ in attributes that vary widely across the sample.

It can be shown that the ATT can be consistently estimated with

$$\frac{1}{N_{\text{treat}}}\left[\sum_{i\in\mathcal{I}_{\text{treat}}}(Y_{i,t_1}(1) - Y_{i,t_o}(0)) - \frac{1}{M}\sum_{j\in\mathcal{J}(i)}(Y_{j,t_1}(0) - Y_{j,t_0}(0))\right] \tag{2}$$

Where $\mathcal{I}_{\text{treat}}$ is the set of treated plants and $\mathcal{J}_i$ is the set of the M closest matches to plant i from the group of control plants. Y_{i,t_0} is average emissions of plant i in the pretreatment period ($t = t_0$) and Y_{i,t_1} indicates the average emissions of plant i after the introduction of the cap-and-trade program ($t = t_1$).

I also implement the finite bias adjustment suggested in Abadie and Imbens (2011).

This approach uses OLS to adjust for remaining differences in covariates between the treated entity and the matched control group.

To test for adverse environmental justice effects of the program, I run the following regression.

$$\begin{aligned}\Delta\text{Emissions}_i = \beta_0 + \beta_1 Treat + \beta_2 Treat \times PropMinority \\ + \beta_3 Treat \times PerCapitaIncome + \gamma X_i + \eta_{\mathcal{J}(i)} + \epsilon_i \end{aligned} \tag{3}$$

Where the X_i are a set of plant-level controls and $\eta_{\mathcal{J}(i)}$ is a match-group fixed effect. The regression thus compares each treated unit to the within-group variation from the controls chosen by the semi-parametric matching estimator. The inclusion of interaction terms captures how changes in co-pollutants vary systematically with demographic characteristics of the surrounding community.

For robustness, I also estimate a synthetic control model. The synthetic control approach provides a data-driven method for choosing a control group. Instead of finding controls based on covariate similarity, as in matching, a synthetic control model finds the linear combination of controls that best approximates (by minimizing the mean-square prediction error) the path of California emissions in the before-treatment periods. This linear combination of controls can then be used as a counterfactual, or "synthetic control", for the treated entity. The effect of the policy is the difference between the actual post-treatment outcome for the treated entity and the outcome predicted by the synthetic control.

Given that the synthetic control method is designed to analyze only one treated entity, I aggregate the emissions data to the state level and use average per-plant emissions as my outcome variable. The synthetic control approach makes no provisions for heterogeneous treatment effects. However, to test for environmental justice concerns, I can find the program's effect on low-income and high-minority share communities by limiting the sample to only include these groups.

4 Results

Table 3 shows the estimate of the ATT computed by the semi-parametric matching estimator. The first column of results shows the estimates for the program's effect on NO_x and the second column of results shows the estimates for the program's effect on SO_x. Rows 1 and 2 show estimates using the western United States as a control group. Rows 3 and 4 show estimates allowing all (non-California) plants in the United States to be used as potential controls. Rows 2 and 4 use the Abadie and Imbens (2011) finite-bias adjustment, whereas Rows 1 and 3 do not. Estimates with and without the bias adjustment are qualitatively similar.

Regardless of the control group used, the signs of the key parameter estimates in Table 3 suggest that, on average, California's cap-and-trade program has decreased co-pollutants. However, the effect is statistically significant only when the control plants outside California are drawn from the entire rest of the country. This difference could reflect the fact that, with a larger sample size, there is greater statistical power and thus smaller standard errors. Even if the difference in standard errors is not due to sample size, there is no evidence that the program, on average, had an adverse effect on co-pollutant emissions.

Figure 2 plots emissions of the treated plants both in California and for those non-California plants that have been selected at least once as a control. The comparability of the pre-trends seems to have improved although the NO_x pre-trends for the rest of the U.S. are still an imperfect match to the California trends.

Table 4 shows estimates, for equation 3, for heterogeneous treatment effects by race and income. Column 1 shows estimates for NO_x using the WECC as the control group. Column 2 shows estimates for SO_x, also using the WECC as the control group. Columns 3 and 4 shows estimates for NO_x and SO_x, respectively, allowing every natural-gas plant in the US, outside of California, to be used in matching.

The coefficients of interest for the environmental justice implications of the policy are those on the interaction terms ($Treat\times$ *Proportion Minority*) and ($Treat\times$ *Per-Capita Income*). These coefficients show how changes in emissions have differed with the racial composition and income of the surrounding communities. Any statistically significant results for the estimates of these coefficients would therefore suggest that the gains or losses do not accrue evenly across income and race.

To test for adverse environmental justice impacts on minority communities, the relevant coefficient to consider is the one on ($Treat\times$ *Proportion Minority*). A positive estimate of this coefficient would imply that minority communities saw additional increases in co-pollutants as a result of the policy and that the EJ concerns about the program are supported by the data. A negative coefficient estimate would imply that minority communities saw additional *decreases* in emissions as a result of the policy. The signs of the estimates for this coefficient vary by pollutant and control group. All except one are negative, which would imply disproportionate gains for minority populations from the policy. However, none of the four estimates are statistically significant and thus I cannot reject the null-hypothesis that gains from this policy are *unrelated* to the racial composition of the surrounding communities.

To test the environmental justice concerns along the income dimension, we need to examine the coefficient on ($Treat\times$ Per-Capita Income). Concerns about adverse environmental justice would be valid if the estimates of this coefficient were negative. This would mean that lower-income communities saw additional increases, or smaller decreases, in emissions than higher-income communities. Three of the four coefficient estimates are positive, while one is negative. As before, none of these estimates are statistically significant and I cannot reject the null hypothesis that emission reductions are unrelated to income.

To summarize: a statistically significant result for either of these coefficients would imply that the pattern of abatement gains differs systematically with the racial and income characteristics of the community. However, across all pollutants and all control groups, I find no evidence of adverse environmental justice impacts in either race or income as a result of California's cap-and-trade program.

4.1 Robustness Checks

Table 5 displays several robustness checks designed to address possible threats to identification.[18] For comparison, the first two rows of the table reproduce the (bias adjusted) estimates of the effects of the program from Table 3.

Leakages. One concern is that the program may have affected electricity generators outside of California, thereby contaminating the controls. This could occur if, for instance, the increased cost of carbon made unregulated electricity outside of California more attractive to buyers inside California, causing "leakages." Given that the structure of the grid makes it difficult to transfer power outside of the Western Interconnection, I can test for spillovers by eliminating all WECC plants as potential national-level controls. The key estimates remain qualitatively similar, with overlapping confidence intervals.

Anticipatory Effects. There may also be some concern that the results may be biased by anticipatory effects in the lead-up to the introduction of the regulation. Initial permit allocations were determined in part by a firm's historical record of emissions, so firms may have had an incentive to increase their emissions right before the start of the program. There does seem to be an increase in emissions around the program start date, but this increase also occurs for both control groups. To test whether these anticipation effects are significantly biasing my estimates, I drop the year immediately preceding the program (2012) from the sample ,since emissions in that year where used to determine permit allocations, and rerun the estimator. The results for this robustness test are displayed in rows

[18]For brevity I only display the ATT estimates. The heterogeneous treatment effects estimates are also robust and can be found in the appendix.

4 and 5 of Table 5. The resulting estimates are qualitatively similar and, if anything, suggest that the original estimates may understate the NO_x emission reductions due to the program.

San Onofre Closure. In January 2012, the San Onofre nuclear power plant permanently closed. San Onofre was a large source of electricity to California that provided eight percent of in-state electricity generation. Its closure increased the demand for electricity from natural gas plants. As documented in Davis and Hausman (2016), transmission constraints made it difficult for plants to easily replace the lost generation and therefore San Onofre's closure resulted in an increase in emissions and a change in the spatial distribution of electricity generation across the state. Given that matches are made, in part, based off of a plant's emission history between 2010 and 2012, this change could lower the quality of the matches if a plant's pre-closure emissions history no longer predicts the plant's post-closure behavior. To assess the effect of this closure on the matches made by the estimator, I explore a specification that uses only the post-San Onofre closure, but pre-program, emissions history to construct matches. The results, in the last two rows of Table 5, remain qualitatively similar.

Placebo Tests. One concern is that my results could be due to an overfitting of the model by my choice of covariates in a way that produces a spurious statistically significant result. Figure 3 shows results from 50,000 placebo tests, where "treatment" status is randomized across all units in the sample. The estimates for the actual set of treated plants from Table 3 are marked by the thin vertical line. If the estimator is not overfitting, or otherwise downward biasing the results, the distribution of estimated placebo treatment e effect sizes should be centered around zero. This is what occurs.

4.2 Synthetic Control

Figure 4 plots the results of the synthetic control model. The data have been disaggregated to the monthly frequency. Actual California emissions are shown by the solid line and the synthetic control is shown by the dotted line.[19] Treatment effects can be computed by examining the distance between the two lines. Table 6 shows yearly averages of these estimated treatment effects. The estimated effect is much lower than the effect estimated from the matching estimator and is frequently close to zero. There is, however, no evidence of a damaging increase in co-pollutants.

Inference in synthetic control models is done through a permutation test as suggested in Abadie et al. (2010). The basic idea of this test is, if the policy had an effect on the

[19]Weights for the synthetic control can be found in the appendix.

outcome variable of interest, that there should be an increase in the distance between the synthetic control and the actual outcome for the treated entity after the start of the program. Recall that, the synthetic control is constructed to minimize the distance between the synthetic control and the pre-treatment outcome. If the policy has an effect on the outcome variable of interest, the policy will change the relationship between the outcome for the treated entity and the outcome for the synthetic control. In this case, the goodness-of-fit between the synthetic control and the actual data will deteriorate after the start of the program. This "forecast deterioration" can be quantified by calculating the ratio of the MSPE of the synthetic control after the treatment period to the MSPE before the treatment period. A higher MSPE ratio signifies that the program caused a greater deterioration in the goodness of fit, indicating that the estimate is statistically significant.

How large the MSPE ratio of a given synthetic control needs to be, to achieve statistical significance, is determined by comparing it to the distribution of placebo MSPE ratios. In this permutation test, a placebo synthetic control model is estimated for each control unit wherein that unit designated as the treated entity. If the MSPE ratio for the actual treated entity is above the $(1 - \alpha)$ percentile of all the placebo MSPE ratios we reject the null hypothesis that the policy had no effect at the α significance level.

Figure 5 shows the result of this permutation test. The pre-post MSPE ratio for California is in the middle of the distribution of the placebo estimates for both pollutants, suggesting that the policy had no statistically significant effect on co-pollutant emissions.

Table 6 also shows estimates of the cap-and-trade program on low-income and high-minority-share communities. Figure 6 shows synthetic control results where the sample is limited to plants for which the surrounding community has an average income below the California median. Similarly, Figure 7 shows results for plants for which the surrounding community has a minority share above the California median. The estimates of the program's effect for both restricted samples are small, close to zero and statistically insignificant.[20] Although these comparisons of synthetic control estimates between various samples lack the formal statistical testing of the matching estimator, these results suggests that abatement patterns across plants have not differed for plants in communities with lower-than-average incomes or higher-than-average minority shares, when compared to the state as a whole.

Lastly, Figure 8 and Figure 9 plot the difference between the actual data and the synthetic control for the California and placebo estimates, and therefore show the treatment effect. This difference is illustrated for the entire sample as well as the low-income and high-minority-share restricted samples. The California effect lies in the center of the placebo

[20]Permutation tests for the restricted samples can be found in the Appendix.

estimates. These plots can be more intuitive than the permutation-test figures, and they also show that the synthetic control approach implies a statistically insignificant effect of the cap-and-trade program on co-pollutants. There is no evidence that this small and statistically insignificant effect differs for low-income and/or high-minority-share communities.

5 Caveats and Directions for Future Research

My data sources and statistical approaches have several limitations and qualifications. First, my analysis focuses on just one sector, electricity generation, and does not take into account the effect of carbon cap-and-trade on the distribution of pollutants from other sectors. Within-state electricity generation represents only eleven percent of California's greenhouse gas emissions. The other major contributing sectors are transportation (39 percent) and other industrial sources (23 percent). The external validity of my results depends on the degree of similarity in the pattern of abatement in the electricity generation sector compared to other sectors that also fall under the carbon cap-and-trade program. It is worth noting, though, that the concerns of environmental justice activists often rely on general arguments that are not specific to any one sector. Future work could expand this analysis to address changes in the distribution of co-pollutant abatement in other industries.

Given the relatively short period of time that has passed since the implementation of this carbon trading program, my results reflect only the short-run responses of firms and cannot describe possible long-run effects. For instance, the cap-and-trade program could affect plant entry and exit decisions differently for clean and dirty plants. Cap-and-trade could encourage cleaner plants to open, or dirtier plants to close. Thus, long-term co-pollutant abatement patterns may be different from short-term abatement patterns.

The California carbon cap-and-trade program was announced many years before it became law, so there may be reason to believe that there was enough lead time for firms to invest in abatement capital and other anticipation measures. Unfortunately, no data exists on this behavior, so I am unable to characterize fully the extent of relevant capital investment that occurred before or during my sample period. If there was a wave of anticipatory investment before the start of the program, the direction and degree of bias in my estimates would depend on how the pattern of that investment is related to local income and demographic characteristics.

Finally, my results may not generalize in a straightforward fashion to other carbon-dioxide trading programs. California's inventory of power plants was already much cleaner relative to the average grid. This property reflects both a lack of coal generation and long-standing tighter air quality regulations in general. Carbon cap-and-trade may lead to more co-pollutant abatement in a jurisdiction that relies more heavily on dirty fuels like coal, when coal could be easily replaced by generally cleaner fuels like natural gas.

6 Conclusion

Environmental justice advocates have expressed grave concerns about the potential negative effects of California's carbon cap and trade program on co-pollutants. Oversimplified analytical approaches can indeed lead to the impression that co-pollutants have increased as a result of the introduction of the program. In this paper, however, I address these EJ concerns using modern econometric program-evaluation methods and detailed data from the electricity sector. I use both a semi-parametric matching estimator and synthetic control approach to construct relevant control groups for California's treated plants in a systematic data-driven way.

Co-pollutant abatement will be determined by the spatial pattern of carbon abatement which is in turn determined by spatial pattern of marginal abatement costs. I find no evidence of increase in co-pollutants emissions, or that the changes in these emissions have varied systematically with the race or income of the communities surrounding these electricity generating plants. This suggests, that marginal carbon abatement costs do not differ systematically with the sociodemographic characteristics of the communities that surround these plants. Additionally, I find some evidence (although this inference is not robust to all plausible specifications) that the program may actually have caused a decrease in average co-pollutant emissions.

Further work should examine the impact of carbon cap-and-trade programs in other industries, and in other settings, to ensure that these results hold more generally. Additionally, as more years of data accumulate, researchers will be able to explore the longer-term impacts of the program, and growing sample sizes will permit greater precision in the key parameter estimates.

For the electricity sector, at least, my results suggest that policymakers have no apparent basis for worrying about environmental justice concerns relating to the California carbon cap-and-trade program.

References

Abadie, Alberto (2005), "Semiparametric Difference-in-Difference Estimators." *Review of Economic Studies*, 72, 1–19.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller (2010), "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." *Journal of the American Statistical Association*, 105, 493–505, URL `http://www.tandfonline.com/doi/abs/10.1198/jasa.2009.ap08746`.

Abadie, Alberto and Guido W. Imbens (2006), "Large sample properties of matching estimators for average treatment effects." *Econometrica*, 74, 235–267.

Abadie, Alberto and Guido W. Imbens (2011), "Bias-corrected matching estimators for average treatment effects." *Journal of Business and Economic Statistics*, 29, 1–11.

Agee, Mark D., Scott E. Atkinson, Thomas D. Crocker, and Jonathan W. Williams (2014), "Non-separable pollution control: Implications for a CO2 Emissions Cap-and-Trade system." *Resource and Energy Economics*, 36, 64–82, URL `http://dx.doi.org/10.1016/j.reseneeco.2013.11.002`.

Barboza, Tony and Chris Megerian (2017), "Questions remain as Gov. Brown signs legislation to address neighborhood-level air pollution." URL `http://www.latimes.com/local/lanow/la-me-ln-air-pollution-law-20170725-story.html`.

California Air Resource Board (2017), "The California Environmental Justice Advisory Committee's Declaration in Support of Carbon Pricing Reform in California."

Climate Hawks Vote (2017), "California: Broad Coalition Of Environmental Justice, Climate Groups Oppose Cap And Trade Bill." URL `http://climatehawksvote.com/news/press-releases/california-broad-coalition-environmental-justice-climate-groups-oppose-cap-trade-bill`

Cushing, Lara J, Madeline Wander, Rachel Morello-Frosch, Manuel Pastor, Allen Zhu, and James Sadd (2016), "A Prelimnary Environmental Equity Assesment of California's Cap and Trade Program."

Davis, Lucas and Catherine Hausman (2016), "Market impacts of a nuclear power plant closure." *American Economic Journal: Applied Economics*, 8, 92–122.

Farber, Daniel A. (2012), "Pollution Markets and Social Equity: Analyzing the Fairness of Cap and Trade." *Ecology Law Quarterly*, 39, 2–53.

Fischer, Carolyn and Alan K. Fox (2012), "Comparing policies to combat emissions leakage: Border carbon adjustments versus rebates." *Journal of Environmental Economics and Management*, 64, 199–216, URL `http://dx.doi.org/10.1016/j.jeem.2012.01.005`.

Fowlie, Meredith, Stephen Holland, and Erin Mansur (2012), "What Do Emissions Markets Deliver and to Whom ? Evidence from Southern California ' s NOx Trading Program." *American Economic Review*, 102, 965–993.

Fullerton, Don and Erich Muehlegger (2017), "Who Bears the Economic Costs of Environmental Regulations." *NBER Working Paper Series*.

Geuss, Megan (2017), "California approves cap-and-trade scheme until 2030." URL `https://arstechnica.com/tech-policy/2017/07/california-approves-cap-and-trade-scheme-until-2030/`.

Guerin, Emily (2017), "Environmental Groups Say California's Climate Program Has Not Helped Them." URL `http://www.npr.org/2017/02/24/515379885/environmental-groups-say-californias-climate-program-has-not-helped-them`.

Haninger, Kevin, Lala Ma, and Christopher Timmins (2017), "The Value of Brown fi eld Remediation." *Journal of the Association of Environmental and Resource Economists*, 4, 197–241.

Heckman, J, H Ichimura, J Smith, and P Todd (1998), "Characterizing selection bias using experimental data." *Econometrica*, 66, 1017–1098, URL `http://www.nber.org/papers/w6699`.

Heckman, J. J., H. Ichimura, and P. E. Todd (1997), "Matching As An Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme." *The Review of Economic Studies*, 64, 605–654, URL `https://academic.oup.com/restud/article-lookup/doi/10.2307/2971733`.

Kahn, Debra (2016), "California to Extend Cap-and-Trade System to 2050." URL `https://www.scientificamerican.com/article/california-to-extend-cap-and-trade-system-to-2050/`.

Mason, Melanie, Chris Megerian, and Tony Barboza (2017), "Here's why progressives aren't thrilled with Gov. Brown's cap-and-trade plan." URL `http://www.latimes.com/politics/la-pol-ca-cap-and-trade-reactions-20170712-story.html`.

Meng, Kyle C (2017), "Is Cap-and-Trade Causing More Greenhouse Gas Emissions in Disadvantaged Communities ?"

Muller, Nicholas Z. (2012), "The Design of Optimal Climate Policy with Air Pollution Co-Benefits." *Resource and Energy Economics*, 34, 696–722, URL `http://dx.doi.org/10.1016/j.reseneeco.2012.07.002`.

Novan, Kevin (2017), "Overlapping Environmental Policies and the Impact on Pollution." *Journal of Environmental Economics and Management*, 4, S153—-S199.

Rubin, Donald B. (1973), "Matching to Remove Bias in Observational Studies." *Biometrics*, 29, 159–183.

Tables and Figures

Figure 1: Emissions for Treatment and Control Plants From 2009-2016

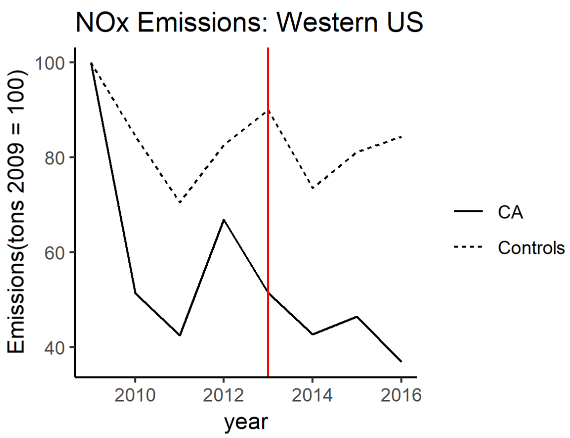

(a) NOx Western U.S.

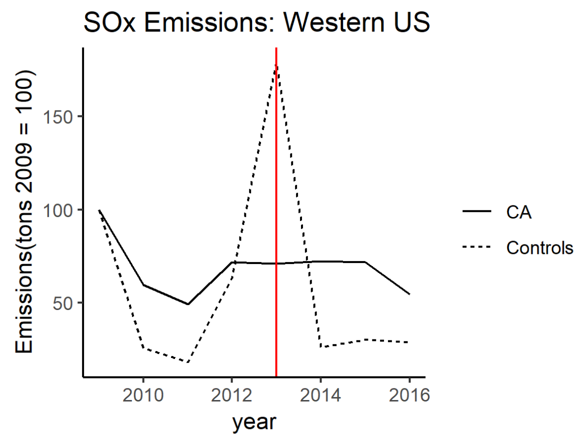

(b) SOx Western U.S.

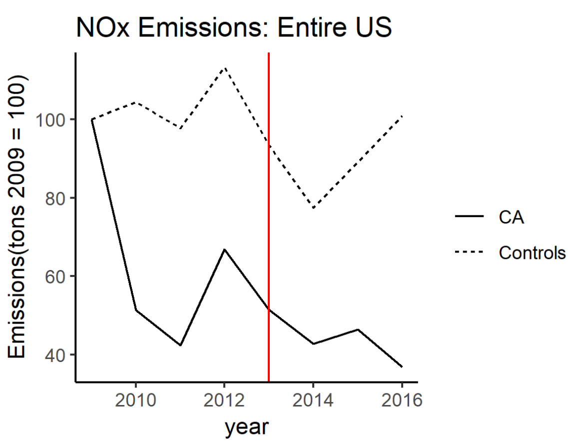

(c) NOx Entire U.S.

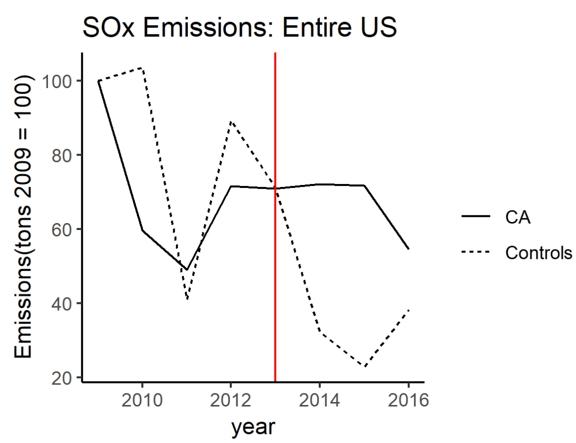

(d) SOx Entire U.S.

Notes: This figure shows historical emissions for the treated plants as well as the historical emissions for all control plants.

Figure 2: Emissions for Treated Plants and Plants Used For Matching

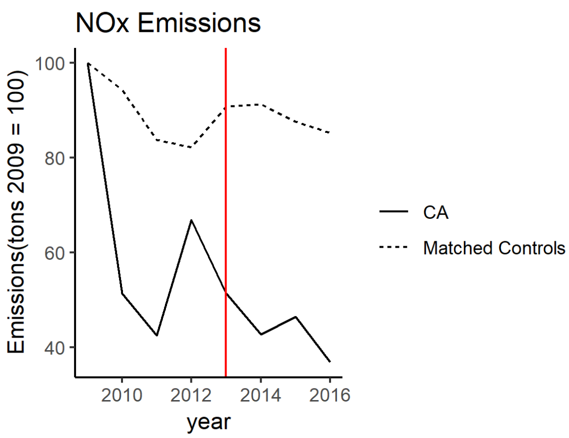

(a) NOx Western U.S.

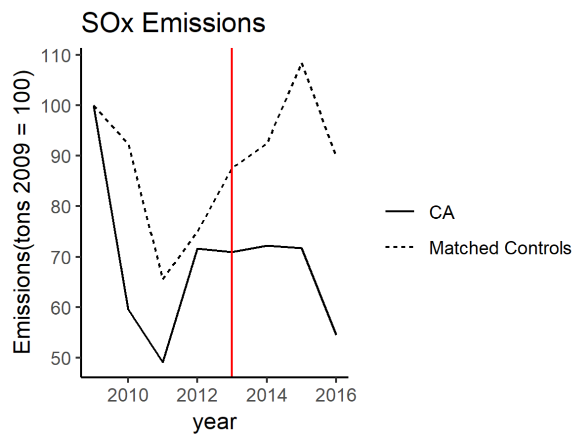

(b) SOx Western U.S.

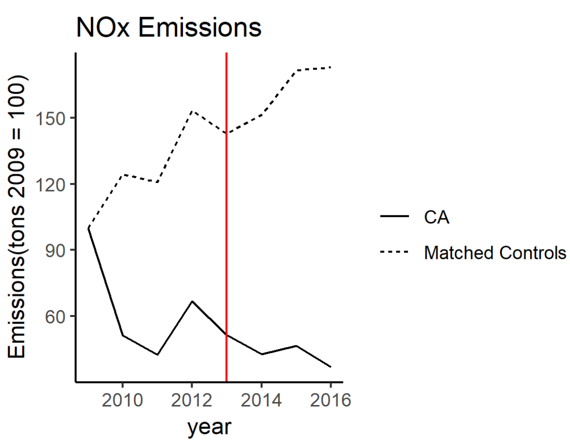

(c) NOx Entire U.S.

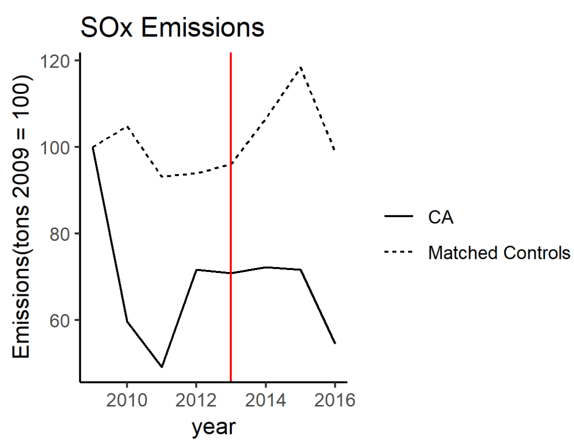

(d) SOx Entire U.S.

Notes: This figure shows historical emissions for the treated plants as well as the historical emissions for all control plants that were matched at least once to a treated plant.

Figure 3: Placebo Test for Semi-Parametric Matching Estimator

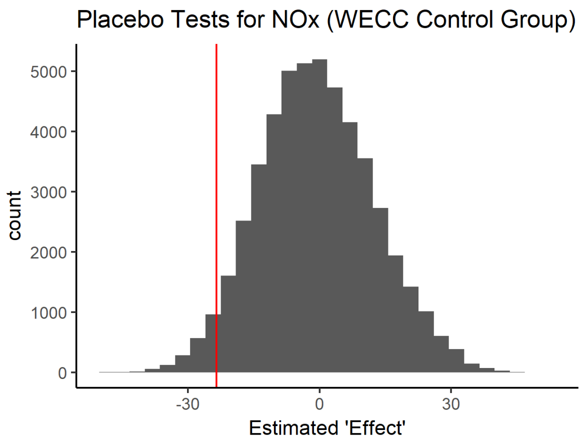

(a) NOx Western U.S.

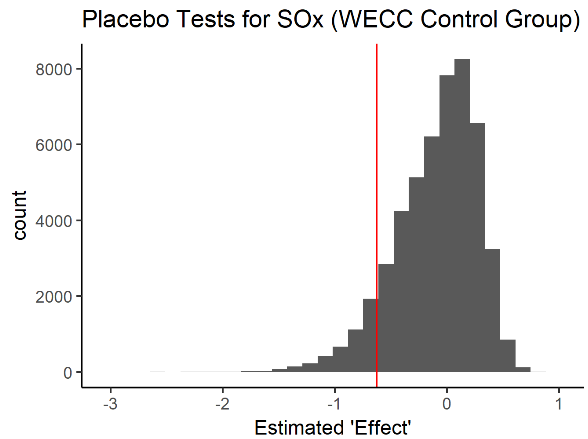

(b) SOx Western U.S.

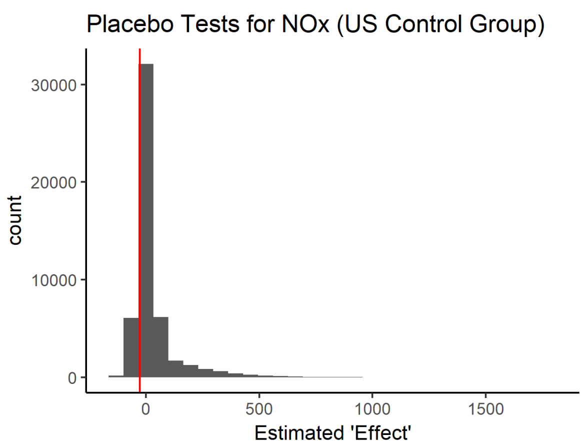

(c) NOx Entire U.S.

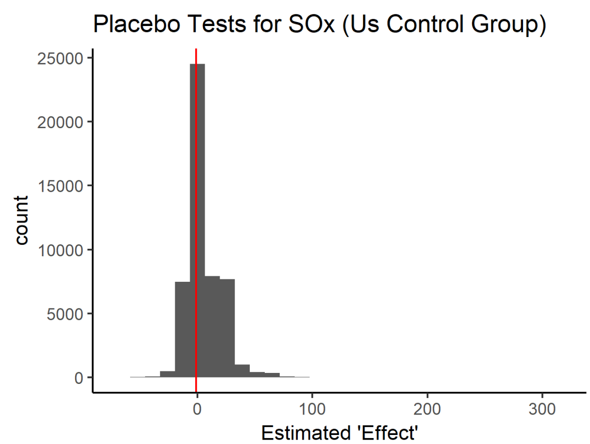

(d) SOx Entire U.S.

Notes: This figure shows ATT estimates from a placebo test of the semi-parametric nearest-neighbor estimator. “Treatment” status was randomly assigned to a subset of plants from the entire sample. The vertical red line shows the estimate calculated from designating the actual treated plants as treated.

Figure 4: Synthetic Control Estimates

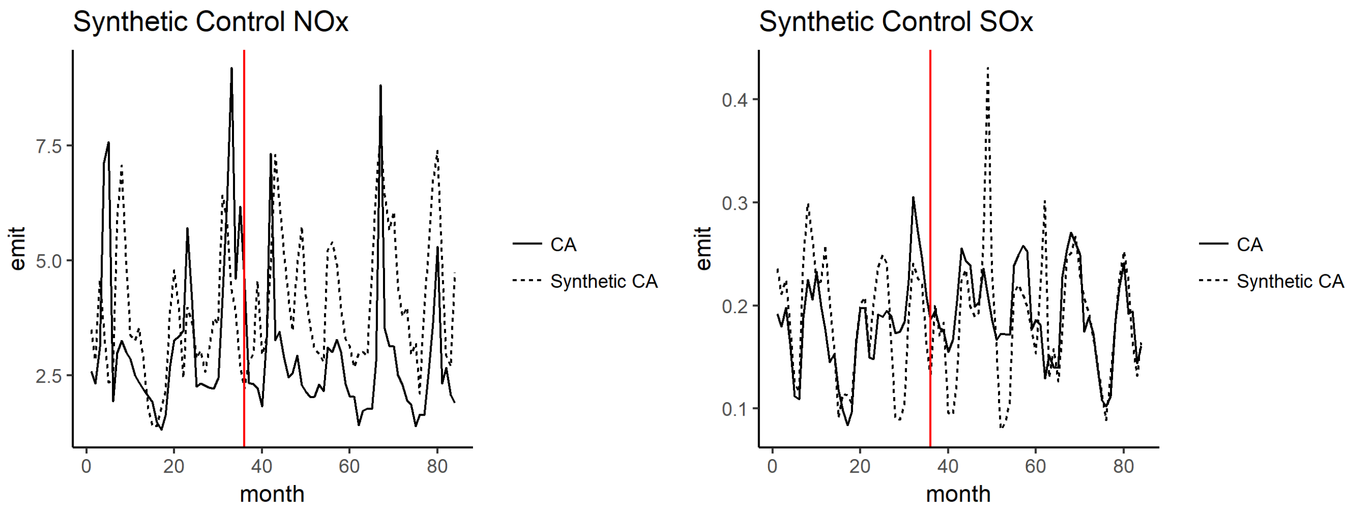

(a) NOx (b) SO_x

Notes: This figure shows historical emissions from California plants (solid line) and the imputed counterfactual emissions from the "synthetic control" (dotted line). The distance between the actual and synthetic control emissions shows the estimated effect of the policy. The vertical red line marks the start of California's carbon cap-and-trade policy.

Figure 5: Permutation Test for Synthetic Control

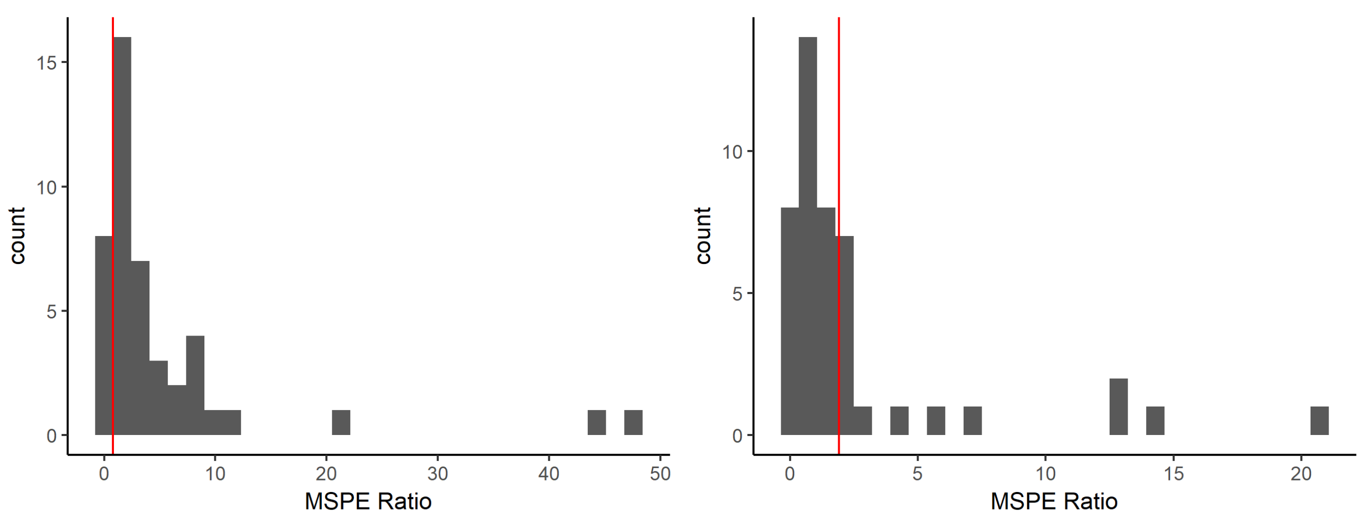

(a) NOx (b) SO_x

Notes: This figure shows the distribution of the post-treatment to pre-treatment MSPE ratios for California and the placebo treated states. The red line shows the MSPE ratio for California.

Figure 6: Synthetic Control Estimates: Only Low-Income Communities Included

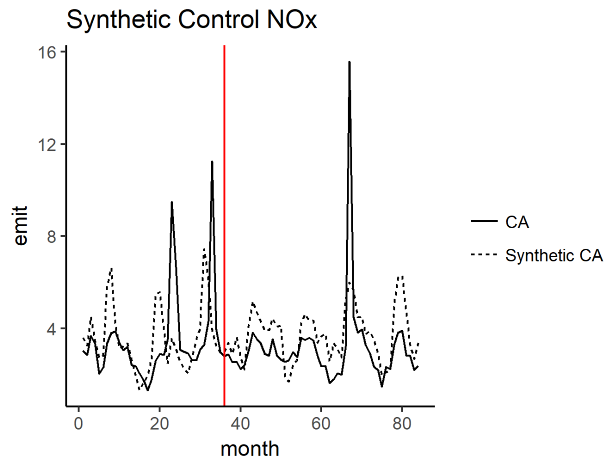

(a) NOx

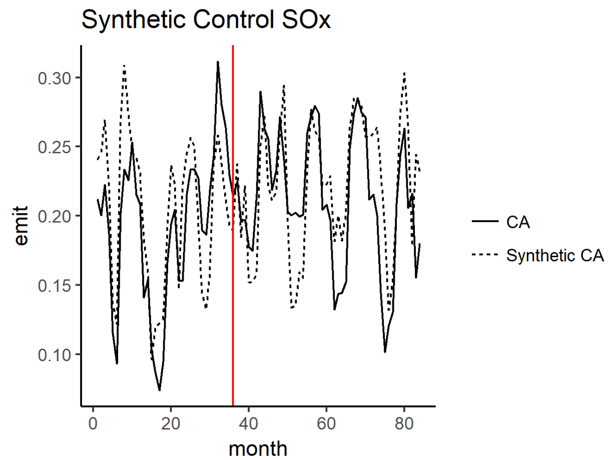

(b) SO_x

Notes: This figure shows the results from a synthetic control estimator when the sample is limited to those plants where the surrounding communities per-capita income is below the California median. The solid line shows historical California emissions and the dotted line shows the calculated counterfactual emissions of the "synthetic control". The difference between both lines is the implied treatment effect. The start of California's cap-and-trade program is marked by the red vertical line.

Figure 7: Synthetic Control Estimates: Only High-Minority-Share Communities

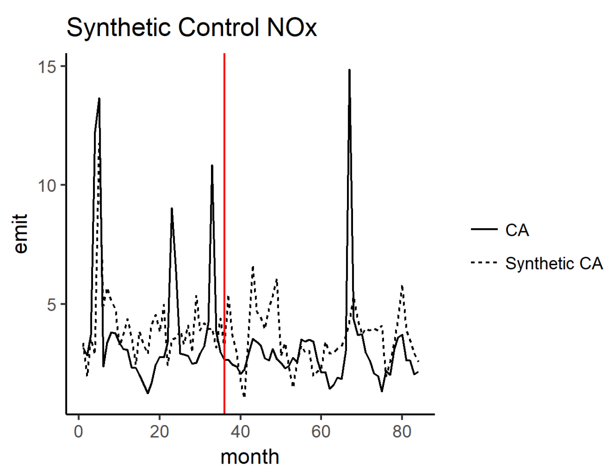

(a) NOx

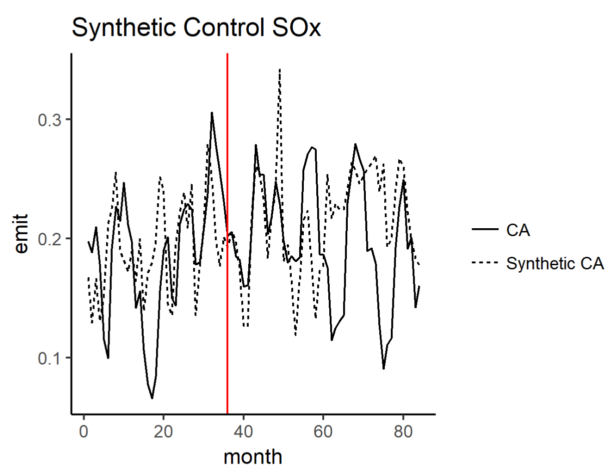

(b) SO_x

Notes: This figure shows the results from a synthetic control estimator when the sample is limited to those plants whose communities have proportion of minorities above the Californian median. The solid line shows historical California emissions and the dotted line shows the calculated counterfactual emissions of the "synthetic control". The difference between both lines is the implied treatment effect. The start of California's cap-and-trade program is marked by the red vertical line.

Figure 8: NO_x Permutation Test: Difference Between Actual Data and Synthetic Control in Treated vs Control States

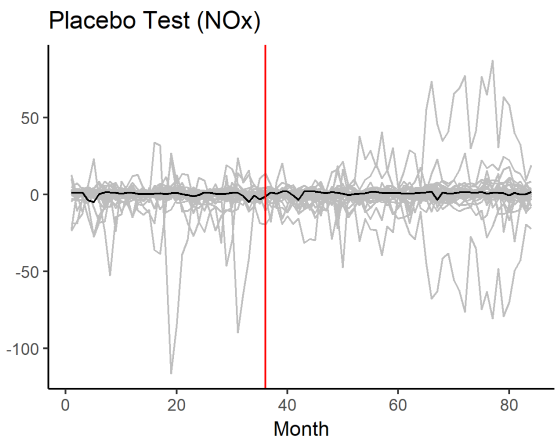

(a) Full Sample

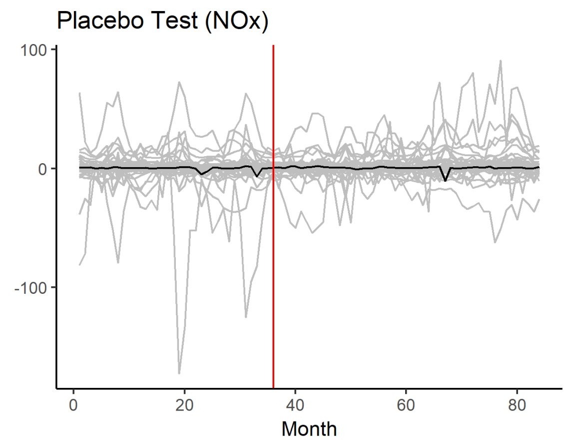

(b) Low-Income Communities Only

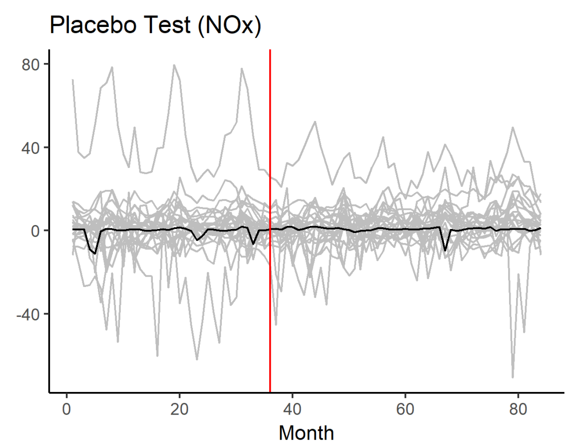

(c) High-Minority-Share Communities Only

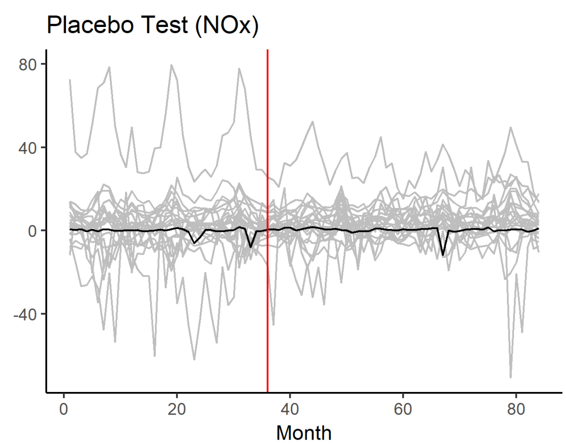

(d) Low-Income and High-Minority-Share Communities

Notes: This figure shows the difference between the observed emission data and the estimated synthetic control counterfactual with each state is designated as the "treated unit". California is shown as the dark line.

Figure 9: SO_x Permutation Test: Difference Between Actual Data and Synthetic Control in Treated vs Control States

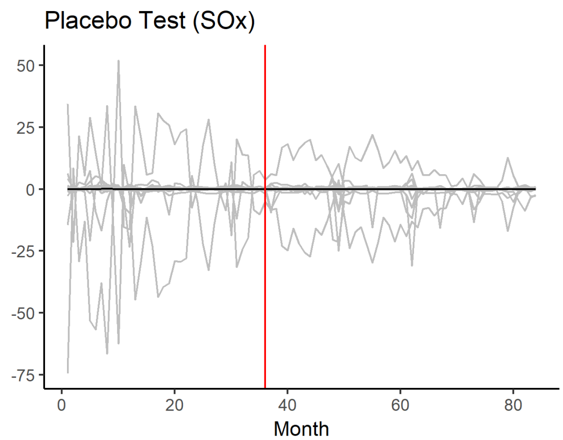

(a) Full Sample

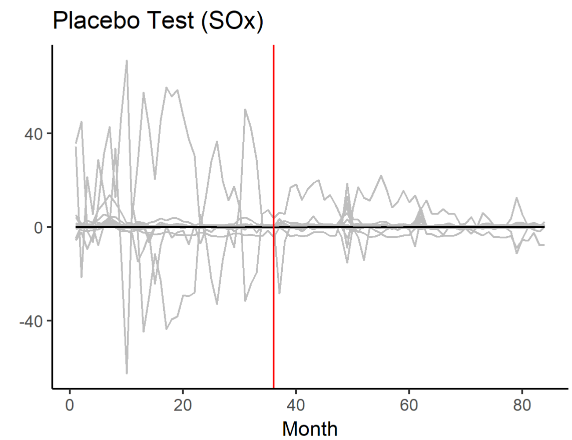

(b) Low-Income Communities Only

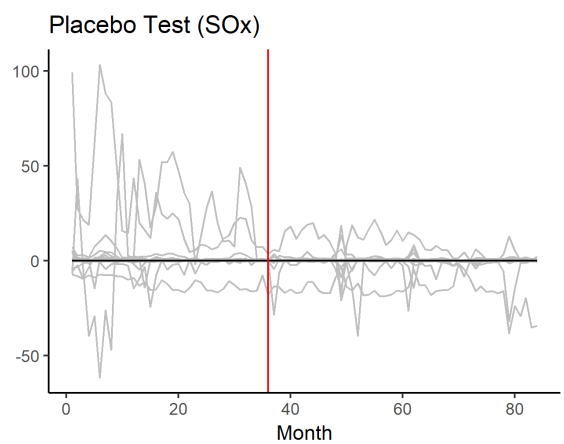

(c) High-Minority-Share Communities Only

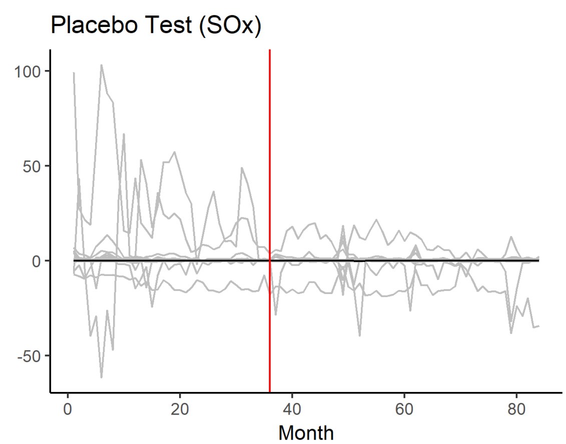

(d) Low-Income and High-Minority-Share Communities

Notes: This figure shows the difference between the observed emission data and the estimated synthetic control counterfactual with each state is designated as the "treated unit". California is shown as the dark line.

Table 1: Summary Statistics

	CA	WECC	EntireUS	$CA - WECC$	$CA - EntireUS$
Plant Emissions					
NO_x Emissions (Tons/Year)	39.3	105.8	198.1	-66.5^{***}	-158.8^{***}
	(124.1)	(156.8)	(506.2)		
SO_x Emissions (Tons/Year)	2.33	3.13	36.0	-0.802^{***}	-33.7^{***}
	(4.55)	(4.02)	(407.1)		
CO_2 Emissions(Thousands of Tons/Year)	423.7	513.9	609.3	-90.2^{***}	-185.6^{***}
	(623.8)	(600.3)	(949.2)		
Plant Attributes					
Nameplate Capacity (MWh)	424.8	476.1	564.6	-51.3^{*}	-139.7^{***}
	(548.8)	(408.9)	(551.9)		
Heat Rate (Thousands of Btu/kWh)	9.67	9.35	13.2	0.322^{*}	-3.48^{*}
	(3.09)	(2.72)	(144.4)		
Annual Net Generation(tWh)	0.96	1.09	1.25	-0.128	-0.294^{***}
	(1.50)	(1.33)	(2.04)		
Community Characteristics					
Proportion Minority	0.491	0.318	0.267	0.174^{***}	0.224^{***}
	(0.248)	(0.215)	(0.233)		
Per-Capita-Income (Thousands of Dollars)	25.2	21.4	24.8	3.77^{***}	0.33
	(11.9)	(5.35)	(8.19)		
n	687	552	5,241		
Average Within-Plant Standard Deviations Over Time					
Proportion Minority	0.008	0.007	0.012	0.001	-0.004
Per-Capita-Income	1.36	1.29	1.40	0.072	-0.039
Number of Plants	86	69	662		

Notes: Summary statistics for natural gas fired power plants. The treatment group is California. Two possible control groups shown are the WECC (which includes Arizona, California, Colorado, Idaho, Montana, New Mexico, Oregon, Utah, Washington, and Wyoming) and all non-California gas-fired plants in the U.S. Standard errors in parentheses.

Table 2: Effects on Co-pollutants of California's
Carbon Cap and Trade: Non-Matched Difference-in-Difference (2010 - 2016)

	Dependent variable:							
	NOx Emissions (tons)				SOx Emissions(tons)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CA	−223.0	802.7*			−251.2**	−138.6		
	(205.7)	(411.2)			(114.2)	(172.1)		
CA × Per_Cap_Inc		145.2*	18.8	27.1		44.4**	142.7	125.1*
		(75.3)	(60.5)	(56.6)		(21.7)	(97.8)	(69.9)
CA × Prop_Minority		19,802.4**	1,275.5	2,127.8		5,422.2*	−48.6	−827.0
		(9,040.3)	(1,834.8)	(2,558.0)		(2,989.9)	(2,827.9)	(3,208.4)
Post	−609.6***	−530.2***	−475.9***	−1,129.3	−542.9***	−508.8***	−423.7***	−469.3
	(185.8)	(172.0)	(182.7)	(705.3)	(163.3)	(159.7)	(140.4)	(373.9)
Post × Per_Cap_Inc		107.5	−114.1*	−25.2		12.3	−31.1	−14.6
		(131.8)	(69.3)	(71.3)		(62.5)	(123.0)	(115.0)
Post × Prop_Minority		−5,865.0	350.3	3,942.6		−6,085.9*	−1,341.0	503.9
		(6,493.6)	(2,526.9)	(3,254.0)		(3,308.7)	(3,892.8)	(4,500.5)
CA × Post	485.6***	363.8**	474.6***	1,109.9	505.5***	457.8***	438.0***	482.7
	(169.6)	(155.7)	(181.2)	(701.3)	(163.7)	(160.4)	(142.9)	(379.3)
CA × Post × Per_Cap_Inc		−224.7	117.0*	27.2		−76.9	32.8	14.8
		(184.2)	(69.3)	(71.4)		(71.2)	(122.5)	(115.0)
CA × Post × Prop_Minority		−25,883.2	−557.3	−4,113.7		−4,071.1	1,399.2	−503.2
		(17,915.9)	(2,534.7)	(3,282.6)		(6,396.9)	(3,897.4)	(4,503.2)
Constant	−766.6***	2,004.8*			−85.8	200.0		
	(282.7)	(1,085.6)			(164.4)	(476.9)		
Per_Cap_Inc		−90.7**	−20.1	−27.2		−12.6	−144.0	−125.8*
		(38.6)	(60.7)	(56.6)		(14.4)	(98.1)	(70.0)
Prop_Minority		−2,945.8**	−1,095.0	−1,965.4		−101.0	27.8	811.1
		(1,303.7)	(1,822.0)	(2,529.4)		(499.0)	(2,831.4)	(3,211.5)
Fixed Effects	No	No	Yes	Yes	No	No	Yes	Yes
State Specific Time Trends	No	No	No	Yes	No	No	No	Yes
Observations	1,329	1,328	1,328	1,328	1,329	1,328	1,328	1,328

Notes: This table shows parameter estimates from a panel diff-in-diff for the effect of California's carbon cap-and-trade program conditional on the plants surrounding demographics. Controls include Primary fuel type and nameplate capacity. All interactions are demeaned. Standard errors, in parenthesis, are clustered at the plant level.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Table 3: Average Treatment Effect on the Treated Estimates From Matched Difference-in-Difference For California's Carbon Cap-and-Trade on Co-pollutants

	NO_x	SO_x	N-Treated	N-Control
Western U.S.	-21.4	-0.587	86	69
	(16.0)	(0.444)		
Western U.S.(With Bias Adjustment)	-23.5	-0.626	86	69
	(15.7)	(0.46)		
Entire U.S.	-22.1^{**}	-0.611	86	662
	(9.15)	(0.680)		
Entire US (With Bias Adjustment)	-25.7^{***}	-1.22^{*}	86	662
	(9.35)	(0.680)		

Notes: This table shows parameter estimates from a nearest-neighbor matched difference-in-difference estimator. Distance is computed using the Mahalanobis norm based off of efficiency (heat rate), nameplate capacity and past emission histories.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Table 4: Heterogeneous Treatment Effects

	Emissions (Tons/Year)			
	Western US NOx	Western US SOx	Entire US NOx	Entire US SOx
	(1)	(2)	(3)	(4)
Treat	−4,118.36	−33.37	2,868.50	−2,158.89
	(3,207.33)	(70.11)	(2,686.87)	(2,670.03)
Proportion Minority	106.13	1.41	−15.29	142.13*
	(124.45)	(2.72)	(55.19)	(84.88)
Per-Capita Income	−3.28	−0.09	1.34	3.77*
	(3.62)	(0.09)	(1.88)	(2.15)
Treat× Proportion Minority	−178.54	−1.43	117.41	−86.49
(Adverse EJ $\implies coef > 0$)	(138.92)	(3.06)	(109.15)	(108.21)
Treat× Per-Capita Income	3.86	0.06	1.63	−2.01
(Adverse EJ $\implies coef < 0$)	(3.66)	(0.08)	(3.43)	(2.92)
Constant	33.76	2.09	−11.60	−61.05
	(76.06)	(3.29)	(70.69)	(79.46)
Observations	155	155	748	748

Notes: This table shows estimates of heterogeneous treatment effects. Estimates are computed from a regression of changes in co-pollutant emissions on demographic variables and a match-group fixed effect.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Table 5: Robustness Checks

Control Group	NOx	SOx	N-Treated	N-Controls
Western US	-23.5	-0.626	86	69
	(15.7)	(0.456)		
Entire US	-25.7^{***}	-1.22^{*}	86	662
	(9.35)	(0.68)		
No Western States	-26.5^{***}	-1.23^{*}	86	593
	(9.40)	(0.690)		
No 2012 West	-24.1	-0.662	86	69
	(18.1)	(0.525)		
No 2012 Entire US	-33.0^{***}	-1.23^{*}	86	649
	(10.8)	(0.677)		
After San Onofre Close (2012) Only: West	-23.6	-1.17^{**}	86	69
	(19.1)	(0.474)		
After San Onofre Close (2012) Only: Entire U.S.	-21.4^{***}	-1.47	86	661
	(8.19)	(3.29)		

Notes: This table shows various robustness checks. The first two rows repeat the results shown in Table 3. The third row drops all WECC plants from the control group to test for spillovers. The fourth and fifth row show estimates when the year 2012 is dropped from the sample to test for anticipation effects.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Table 6: Synthetic Control Estimates

	2013	2014	2015	2016
		Full Sample		
NO_x	-1.32	-0.79	-0.69	-0.77
SO_x	-0.038	-0.003	-0.015	0.0002
		Only Low-Income Communities		
NO_x	-0.819	-0.480	-.167	-1.04
SO_x	0.015	.017	-0.028	-0.042
		Only High-Minority-Share Communities		
NO_x	-1.25	-0.044	-0.052	-1.13
SO_x	0.01	0.029	-0.048	-0.061

Notes: This table shows estimates, using the synthetic control method, of the effect of California's cap-and-trade program for each year. Only low-income communities and only high minority share communities only include communities with income (minority share) below(above) the California median. Inference is done using a permutation test. None of the results are statistically significant. Weights for the synthetic control are shown in the appendix.

Appendix

Figures A1 and A2 show the results of the permutation test for the synthetic control estimates for low-income and high-minority-share communities respectively. The MSPE ratio for California is shown by the vertical line. The fact that the Californian MSPE ratio is smaller than the MSPE ratio of many placebo estimates implies that the synthetic control estimates are statistically insignificant.

Table A1 shows the non-zero weights on emissions in other states used to construct the synthetic control for California. The synthetic control is calculated by taking a weighted average of these control-state emissions.

Table A2 shows nearest-neighbor matched difference-in-difference estimates for the California carbon cap-and-trade program where "distance" is computed using the simple propensity score instead of the Mahalanobis norm. The magnitudes of the estimates are qualitatively similar and the p-values are somewhat lower, resulting in estimates with a higher degree of statistical significance. Table A3 shows the heterogeneous treatment effect results obtained from these alternative matches. Like my preferred specification, using the Mahalanobis norm, I cannot reject the null hypothesis that average emission changes are invariant to race and income of the surrounding communities associated with the plants when I instead use the propensity score.

Table A4 shows the estimates for the heterogeneous treatment effect associated with the robustness checks in Table 5. The results are again qualitatively similar to those in Table 4.

Lastly Tables A5 and A6 show how the key results vary with the choice of the number of nearest neighbors. The results are likewise qualitatively similar to those presented in the body of the paper.

Appendix Figures and Tables

Figure A1: Permutation Test for Synthetic Control: Low-Income Communities Only

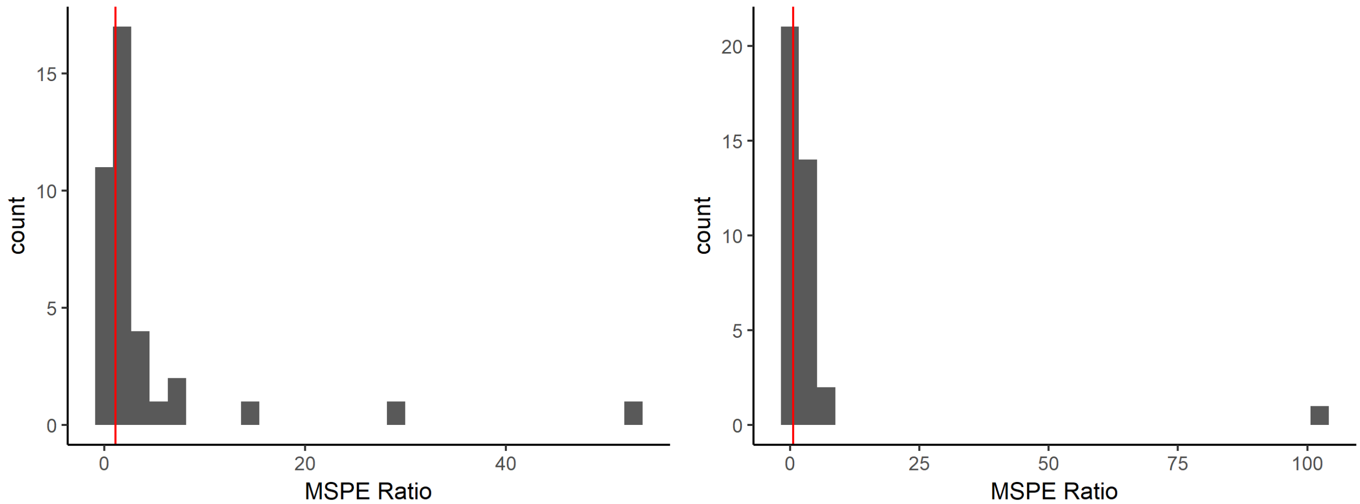

(a) NOx (b) SO_x

Notes: This figure shows the distribution of the post-treatment to pre-treatment MSPE ratios for California and the placebo treated states. The red line shows the MSPE ratio for California. The sample of plants is limited to those whose surrounding communities have a per-capita income below the median of Californian plant-adjacent communities.

Figure A2: Permutation Test for Synthetic Control: Only High-Minority Share Communities

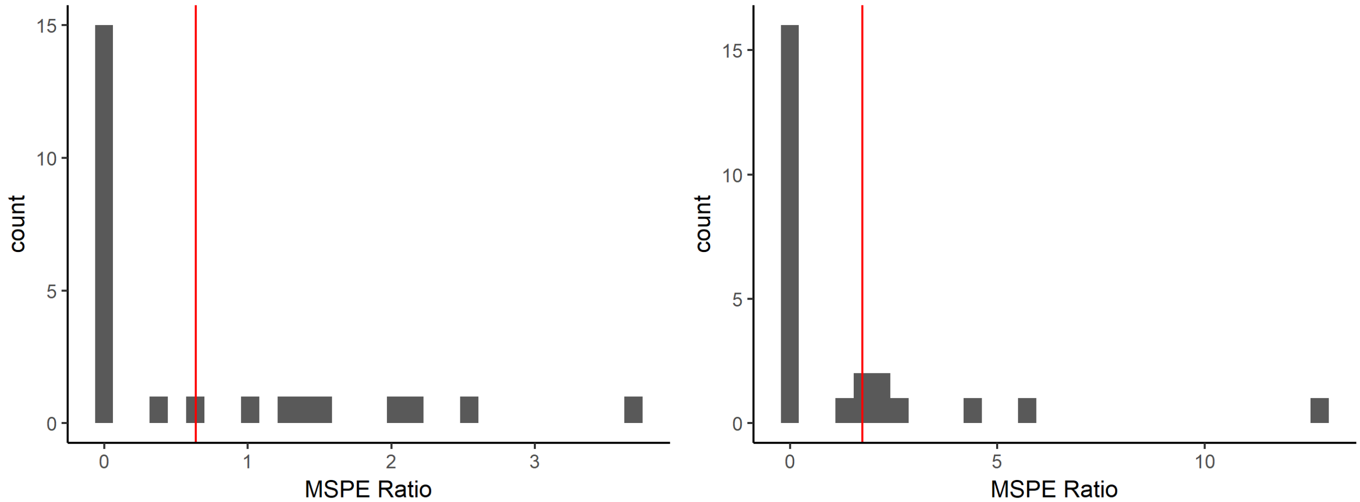

(a) NOx (b) SO_x

Notes: This figure shows the distribution of the post-treatment to pre-treatment MSPE ratios for California and the placebo treated states. The red line shows the MSPE ratio for California. The sample is limited to include only plants whose community have a minority share above the median of Californian plant-adjacent communities.

Table A1: Synthetic Control Weights

Full Sample				Low-Income Only				High-Minority Share Only			
NO_x		SO_x		NO_x		SO_x		NO_x		SO_x	
State	Weight	State	Weight	State	Weight	State	Weight	State	Weight	State	Weight
NV	0.0964	AZ	0.0267	AR	0.0329	ID	0.131	AL	0.487	AL	0.449
OR	0.127	ID	0.0832	NE	0.478	IL	0.16	MA	0.384	AZ	0.0256
RI	0.124	KS	0.323	NV	0.0782	KY	0.214	NV	0.129	CO	0.01
WA	0.0231	MA	0.001	OR	0.107	NE	0.0536			IL	0.215
WY	0.629	ME	0.0303	WA	0.304	OR	0.177			IN	0.0597
		MT	0.214			PA	0.221			NJ	0.001
		NH	0.0577			SC	0.0002			NV	0.145
		NV	0.0872			TX	0.0419			NY	0.001
		OR	0.0669			WA	0.001			SC	0.002
		WA	0.109							TX	0.0927
		WV	0.001							UT	0.0002
										VA	0.0001

Notes: This table shows the weights used to construct the synthetic California used as a counterfactual. Weights were chosen to minimize the mean square prediction error for the pre-treatment period.

Table A2: ATT Using The Propensity Score for Matching

Control Group	NO_x^*	SO_x	N-Treated	N-Controls
Western U.S.	-25.3	-0.377	86	69
	(14.9)	(0.314)		
Western U.S. (With Bias Adjustment)	-26.7^{*}	-0.356	86	69
	(14.4)	(0.306)		
Entire U.S.	-29.0^{***}	-0.488^{***}	86	662
	(9.01)	(0.243)		
Entire U.S. (With Bias Adjustment)	-32.9^{***}	-0.512^{**}	86	662
	(9.26)	(0.243)		

Notes: This table shows results from a matched difference-in-difference estimator where the nearest-neighbors are found using the propensity score instead of the Mahalanobis norm.

Table A3: Heterogeneous Treatment Effects: Propensity Score Matching

	Emissions (Tons/Year)			
	Western US NOx	Western US SOx	Entire US NOx	Entire US SOx
	(1)	(2)	(3)	(4)
Treat	−3,137.24	−4.30	1,688.96	−2,524.68
	(3,388.25)	(78.53)	(2,661.72)	(2,681.01)
Proportion Minority	166.60	−0.84	−7.90	145.09*
	(142.89)	(2.89)	(55.55)	(84.55)
Per-Capita Income	−2.59	−0.20	1.47	3.65*
	(2.24)	(0.15)	(1.74)	(2.13)
Treat× Proportion Minority	−134.71	−0.13	69.20	−101.74
	(146.66)	(3.45)	(107.98)	(108.63)
Treat× Per-Capita Income	2.73	0.14	1.28	−1.30
	(2.48)	(0.20)	(3.40)	(2.98)
Constant	−8.86	5.72	−39.15	−92.31
	(71.45)	(6.12)	(61.55)	(79.84)
Observations	155	155	748	748

Notes: This table shows estimates of the heterogenous treatment effects using the matches obtained from propensity score matching obtained from the estimator described in Table A2.

Table A4: Robustness Table for Heterogeneous Treatment Effects

	Dependent variable:									
	Western U.S.		Entire U.S.		No Western States		No 2012		Post Closing	
	NO_x	SOx	NO_x	SO_x	NO_x	SO_x	NO_x	SO_x	NO_x	SO_x
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treat	−4,011.57	−32.55	2,858.26	−2,296.46	3,900.56	−2,646.48	2,822.03	−289.64	1,557.43	−2,532.94*
	(3,065.11)	(68.92)	(2,735.31)	(2,738.06)	(3,248.61)	(3,163.29)	(2,822.15)	(3,108.21)	(2,185.34)	(1,530.68)
Proportion Minority	100.67	1.57	−13.63	146.18*	−23.97	163.32*	−37.55	80.64	5.57	124.05*
	(122.13)	(2.82)	(55.55)	(86.78)	(63.20)	(95.79)	(59.76)	(82.23)	(54.73)	(65.44)
Per-Capita Income	−3.64	−0.09	1.33	3.82*	1.64	4.46*	1.11	1.73	1.71	3.22*
	(3.99)	(0.09)	(1.89)	(2.17)	(2.01)	(2.44)	(2.11)	(1.86)	(1.45)	(1.80)
Treat× *Proportion Minority*	−174.54	−1.40	117.04	−92.09	156.74	−104.54	115.53	−10.58	63.74	−102.17*
	(133.04)	(3.02)	(111.17)	(111.01)	(129.89)	(126.23)	(114.74)	(126.28)	(88.75)	(61.84)
Treat× *Per-Capita Income*	4.15	0.06	1.66	−2.04	2.26	−2.43	2.32	0.51	0.08	−4.08**
	(4.06)	(0.09)	(3.48)	(2.97)	(3.84)	(3.39)	(3.72)	(3.15)	(2.61)	(1.86)
Constant	26.03	2.06	−12.94	−62.68	−19.04	−80.83	8.70	30.55	−51.09	−113.77**
	(81.30)	(3.34)	(71.09)	(80.83)	(77.50)	(90.88)	(77.61)	(83.25)	(61.19)	(54.07)
Observations	154	154	734	734	666	666	721	721	733	733

Notes: This table fives estimates of heterogeneous treatment effects for each of the robustness checks described in Table 5.

Table A5: ATT Robustness to Choice of Number of Neighbors

	$M = 2$	$M = 3$	$M = 4$	$M = 5$
Western U.S. NO_x	-23.2	-23.5	-25.7	-25.6
	(17.8)	(15.7)	(15.0)	(15.0)
Western U.S. SO_x	-0.716	-0.626	-0.634	-0.674
	(0.507)	(0.456)	(0.454)	(0.470)
Entire U.S. NO_X	-23.9	-25.7	-29.3	-28.9
	(9.73)	(9.35)	(9.89)	(9.82)
Entire US SO_x	-0.807	-1.22	0.124	0.305
	(0.916)	(0.680)	(2.25)	(2.33)

Notes: This table shows how estimates change with the number of plants used to construct the counterfactual outcome.

Table A6: Robustness to Number of Nearest Neighbors: Heterogeneous Treatment Effects

	$M = 2$	$M = 3$	$M = 4$	$M = 5$
$\mathbf{NO_x}$ Western U.S.				
$CA \times Prop\_Minority$	-126.2	-178.5	-121.1	-98.1
	(150.6)	(138.9)	(159.1)	(122.6)
$CA \times Per\_Cap\_Inc$	4.327	3.863	2.439	3.683
	(3.926)	(3.665)	(3.355)	(3.771)
$\mathbf{NO_x}$ Entire U.S.				
$CA \times Prop\_Minority$	113.8	117.4	131.0	108.1
	(89.75)	(109.1)	(111.2)	(110.0)
$CA \times Per\_Cap\_Inc$	1.325	1.625	2.397	1.594
	(3.216)	(3.432)	(3.535)	(3.390)
$\mathbf{SO_x}$ Western U.S.				
$CA \times Prop\_Minority$	1.231	-1.434	-0.076	-1.068
	(2.870)	(3.065)	(2.997)	(3.574)
$CA \times Per\_Cap\_Inc$	0.036	0.059	0.012	0.083
	(0.074)	(0.081)	(0.085)	(0.139)
$\mathbf{SO_x}$ Entire U.S.				
$CA \times Prop\_Minority$	-77.54	-86.49	-44.56	-60.18
	(105.7)	(108.2)	(119.4)	(106.2)
$CA \times Per\_Cap\_Inc$	-1.876	-2.008	-1.645	-1.072
	(2.725)	(2.925)	(2.788)	(2.830)

Notes: For the models of heterogenous treatment effects, this table shows how the estimates change with the number of plants used to construct the counterfactual outcome.