Long Run Effects of Pollution: Evidence from the Acid Rain Program

Alan I. Barreca
Tulane University, IZA, and NBER

Matthew Neidell
Columbia University and NBER

Nicholas J. Sanders
Cornell University

Abstract

We estimate the long-run effects of pollution on adult mortality using quasi-experimental variation induced by the U.S. Acid Rain Program. The program caused an immediate reduction in sulfur dioxide (SO2) emissions at select coal power plants in 1995. We compare long-run changes in adult mortality in counties surrounding regulated power plants to counties that are farther away using an event-study framework with propensity score regression adjustment. We find a 2.5-3.8% reduction in overall mortality. Effects are largest for cardiovascular deaths, which supports the hypothesis that exposure to fine PM is the relevant mechanism. We find no detectable effects for vehicular deaths, which helps rule out endogenous sorting or economic effects as the main channels. These findings suggest short-run mortality estimates understate the benefit of pollution abatement policies.

1. Introduction

There is compelling quasi-experimental evidence pollution has a short-run effect on human health, especially with regards to infants and young children. We know considerably less about the long-run effects of exposure to pollution on adults, although results from observational studies suggest that such effects could be quite large.[1] Obtaining causal estimates of long-run effects is essential to properly evaluating the benefits to environmental regulation. There are two key challenges to addressing this empirical question. First, pollution is tied with omitted factors, like economic growth, that may have an independent effect on mortality. Second, individuals can mitigate the costs of pollution over time by migrating, leading to sample selection bias.

In this paper, we estimate the long-run effects of pollution on adults using a natural experiment generated by the Acid Rain Program (ARP), a cap and trade program designed to reduce sulfur dioxide (SO2) emissions from coal-fired power plants. A number of program-specific factors make the ARP an ideal setting to study the long-run health effects of pollution. First, the ARP facilitates a quasi-experimental research design that mitigates bias from omitted variables. The program only regulated the 110 most polluting power plants at the time of its implementation in 1995. Using semi-parametric distance estimates, Sanders and Barreca (2016) shows the ARP caused SO2 levels to fall detectibly in counties as far as 100 miles away from the regulated plants, meaning health effects can extend far beyond the range of standard potential economic spillovers from regulation. We compare health outcomes in these "nearby" (<100 miles) counties to farther counties to mitigate biases from confounders changing over time independent of SO2.

[1] Two of the most influential studies in setting air quality standards are the Harvard "six cities" study and the Cancer Prevention Study II (Pope et al., 2002). These cohort analyses follow individuals over a length of time and link average pollution exposure with the likelihood of death, adjusting for a handful of covariates. Pope et al. (2002) found that each 10-μg/m3 increase in fine particulate increased all-cause mortality risk by 4%.

Second, the potential for bias from sorting is relatively small since the health effects of the ARP are still unknown, and SO2 is a highly migratory pollutant. Although the health effects of SO2 are generally considered to be small, the ARP had the potential to significantly affect human health in the long-run via decrease in dangerous fine particulate matter (PM2.5), since SO2 and PM2.5 are naturally linked via an atmospheric conversion process. Little was known about PM2.5 at the time, so the potential for migration and other adaptive behavior with respect to *observed* pollution is limited. Existing evidence from the opening and closing of power plants suggest the change in amenity values occurs at only very close distances.[2] We omit counties housing plants from our analysis, meaning treatment comes only from counties within 100 miles of a plant but not directly next to a plant, something that is only identifiable using migratory, non-localized pollution.

Third, the potential for the ARP to affect human health through local economic conditions, as opposed to pollution, is limited. A common concern with using environmental regulations as a source of variation in pollution is that regulations can impact human health independent of pollution, e.g., via employment. With the ARP, compliance costs were much lower than expected due to rapid technological innovation in desulphurization and unexpected benefit from previous railroad deregulations, and as a result the economic effects of the ARP have been very limited (Schmalensee and Stavins, 2012).

We begin analysis in 1990, the year the ARP is first announced, and end in 2005, before the changes in the SO2 permit market potentially shifted the relationship between regulation and

[2] Using detailed micro-census data, Davis (2011) finds that power plant closures affect housing values within 2 miles of a power plant. Davis's (2011) results deal with natural gas power plants, though similar results may extend to coal plants. Although focused on toxic plants, Currie et al. (2015) find that plant closings affect housing values within half a mile of the plant. An additional relevant finding in Currie et al. is the discord between the degree of sorting and the extent of the health effects: housing values change within half a mile of the plant, but health effects extend beyond 1 mile. This suggests that sorting is imperfect, and likely to be more imperfect the broader the distance and more unknown the health effects.

ambient emissions.[3] We use restricted mortality data from the National Center for Health Statistics with information on cause of death and county of residence for the decedent. Our main outcome is the mortality rate for internal causes of death for individuals aged 35-64 in a given county and year.

Our research design compares changes in adult mortality for counties within 100 miles of regulated power plants (treatment group) to counties farther than 100 miles (control group) – to help reduce localized effects of the regulation itself, we omit any counties that actually contain regulated plans. All treatment identification depends on counties that received plant-related pollution but didn't house plants themselves. The core empirical model is an event study where we examine changes in mortality for adults ages 35-64 compared to their mortality in 1994, the year prior to the ARP's implementation. To address the possibility of dissimilar mortality trends, we restrict our sample to counties that are similar on observables at baseline using a p-score matching algorithm and control for flexible time effects by p-score based probability of treatment exposure.

We find a decline in mortality in treatment counties that, rather than being an immediate downward shift in mortality, manifests as a change in mortality trends that becomes statistically significant shortly after the ARP began. We view this as an indication of a change in long-run (rather than short-run) exposure to pollution effects. Consistent with effects operating through particular matter, we observe large effects for cardiovascular deaths but limited respiratory mortality. As a test for confounders, we show no consistent changes in vehicular mortality, which would correlate with shifts in economic or demographic conditions but should be

[3] The market price for SO2 increased to the point that it became cheaper to install scrubbers and reduce emissions, at which point the price for SO2 precipitously fell.

unrelated to pollution. In future versions, we plan to incorporate information on wind direction and pollution intensity of each power plant to move beyond simple treatment and control effects.

While our study is not the first quasi-experimental economic study to explore long run effects (Chen et al., 2013; Anderson, 2015), it offers several innovations. One, we explicitly estimate a dynamic relationship between the ARP and mortality in an event study framework. Two, our analysis focuses specifically on prime age working adults in the United States (as well as other age groups), thus providing an estimate for a group where mortality represents a significant loss in life expectancy and economic productivity. Three, our reduced form effect of the ARP is of direct interest to the recently enacted Clean Power Plan (CPP). The CPP, which aims to cut carbon emissions from the power plant sector, is likely to lead to significant reductions in PM2.5 emissions. Our estimates provide a potential benchmark for evaluating the potential mortality benefits from the CPP.

2. Background

2.1 The ARP

Concerns about acid rain in the 1970s spurred the Acid Deposition Act of 1980, a 10-year program that increased SO2 monitoring, investigated precipitation acidity levels, and established a monitor network to study levels of dry deposition (sulfuric acid deposited in the absence of precipitation). Lessons from the Acid Deposition Act led to the Acid Rain Program (ARP), a provision of the Clean Air Act Amendments of 1990. To reduce SO2 emissions, the EPA instituted a tradable permit system among SO2-polluting power plants, with the intended goal of reducing US SO2 levels to 50% of 1980 levels. The ARP had two phases. Phase I, which began in 1995, regulated the 110 power plants with the highest SO2 emissions. In 2000, Phase II

further limited emissions of Phase I plants and added more plants to program. Both Phase I and Phase II worked through an SO2-emission cap-and-trade system, where plants had the ability to bank and sell allocated permits. Phase I had a considerably larger impact on SO2 emissions than Phase II, since firms had the ability to smooth SO2 production by banking permits in advance of Phase II. As such, our quasi-experimental research design focuses on Phase I.

The EPA distributed SO2 allowances to 263 units at the 110 plants based on baseline (1985-1987) heat input (Stavins, 1998), calculated in British thermal units (BTUs). Each year, plants report SO2 emissions to the EPA for verification, holding a permit for each ton of SO2 produced. For plants polluting in excess of held permits, the EPA assigned a fee of $2,000 (adjusted for inflation) per ton of overage and required eventual accounting for any overages by purchasing sufficient permits. Plants could hold unused permits, and sell or transfer permits across years. As the program moved into Phase II, the EPA further restricted the total number of available annual pollution permits, with a final goal of 8.95 million permitted tons for electric utilities by 2010. The EPA reports the program achieved close to full program compliance, leading to substantial decline in SO2 emissions from regulated plants.

On many accounts, the ARP has been hailed a major success. Studies have estimated the benefits at over $100 billion per year with costs around $3 billion per year (Chestnut and Mills, 2005), making it one of the most cost effective environmental regulations. While much of this success was unanticipated (Schmalensee and Stavins, 2015), the success of the ARP has become the basis for many other cap and trade programs. However, the success of the ARP has mostly been inferred from back-of-the-envelope calculations, with limited direct quasi-experimental studies on the impacts of the program.[4]

[4] As one recent example, Sanders and Barreca (2016) show that the ARP had the unintended consequence of reducing agricultural yields by lowering ground level sulfur.

2.2 SO2 and PM2.5

Airborne SO2 can have negative human health consequences, including respiratory difficulty. Most of these effects are thought to be short-run, leading to increases in hospital admissions for outcomes such as asthma exacerbations. Airborne SO2 can also lead to acid rain, which is thought to have little impact on human health.

Much of the anticipated health effects are from the other pollutants correlated with SO2 levels. Through an atmospheric conversion process, SO2 produces sulfate particles, a portion of small-scale particulate matter (PM). PM2.5 penetrates deep into the lungs to enter the bloodstream, where it causes cardiovascular effects. Given this path of entry, PM2.5 is widely believed to lead to effects in both the short- and long-run. Given its diminutive size, PM2.5 also penetrates into buildings at a high rate, suggesting the ability to avoid it is quite limited.
Both SO2 and PM2.5 can travel great distances from the point source.

2.3 Previous evidence on long-run effects

To our knowledge, there are only two quasi-experimental papers focusing on the long-run effects of pollution. A paper by Chen et al. (2013) exploits the government policy of free winter heating north of the Huai River in China and restricted migration enforced by the government. However, the paper faces concerns over validity of the regression discontinuity design due to the river representing a major separation point. One such concern is that along the river is also a major mountain range such that there are considerable geological, geographical and climatological differences on opposite sides of the river. This issue is exacerbated by the small sample size that forces the authors to use a parametric regression discontinuity design. Further,

this study is less generalizable to the United States because of the extraordinarily high levels of pollution.

Anderson (2015) focuses on residents living upwind versus downwind from major highways. His analysis relies on strong assumptions about the decay of ultrafine particles, an unobserved variable, and incomplete evidence regarding sorting. For example, although he does not find statistically significant differences in housing values on opposite sides of highways, the estimates are in the direction consistent with sorting and are borderline significant. Furthermore, his study focuses on adults over age 75, a narrow subset of the total population.

3. Full Literature Review (in process)

4. Data and methods

4.1 Data sources

We use county-level cause of death mortality data from the Centers for Disease Control (CDC) as our measure of changes in mortality levels. Our current data cover 1990-2005, and represent the universe of US deaths with information on cause of death, classified by either ICD-9 (up to 1997) or ICD-10 (1998 and beyond) cause of death codes. Data include information on age at death, county of residence, sex, and race of the decedent. Our main outcome is the death rate per 100,000 population, which we construct using population data from the Surveillance, Epidemiology, and End Results (SEER) program. We focus on the death rate for internal causes, but also estimate impacts on respiratory-related and cardiovascular-related deaths to better isolate the causal channel. We perform robustness checks using motor vehicle fatalities as a test

of endogenous sorting and local economic effects. In most regressions, we limit analysis to those between 35 and 64 years of age, but we also estimate impacts on other age groups.

We obtain a list of all Phase I power plants from the EPA Air Markets Program Data. These data include per-plant SO2 emissions and British Thermal Unit (BTU) input in 1985, one of the primary measures of regulatory intensity in the ARP. We focus on a simplified identification strategy where we define treatment based on presence of a plant within 100 miles, with no differential levels of treatment intensity, which better parallels the standard difference-in-difference model.

Although we would ideally like to demonstrate a link between the ARP and PM2.5, the sparce reporting of PM2.5 data prevents us from doing so. Instead, we demonstrate a link between the ARP and SO2 using monitor-level pollution data. The EPA reports daily SO2 averages in parts-per-billion (ppb). We convert monitor-level measures to county-level estimates by calculating the distance between each monitor and each county centroid. Using all monitors within a range of a maximum distance d of the centroid, we collapse to the county-by-year level using weights equal to $1/d$. We set $d = 50$, though results are robust to alternate distances. In the off chance that weather patterns differed by treatment status after 1995, we also flexibly control for temperature and rainfall. Weather data are from the National Oceanic and Atmospheric Administration (NOAA), and we aggregate the station-day variables up to the county-year level using inverse distance squared weights, and a maximum distance of 100 miles, between county centroid and weather station.

To test for the possibility that the ARP impacts the local economy, we also estimate impacts on various county-year outcomes using data from Bureau of Economic Analysis Regional Data. These variables include: log income per capita, Medicare insurance payments per

capita, and local employment rates (calculated by dividing wage employment by total county population).

4.2 Treatment definitions

Our core treatment definition involves a binary classification, where we define a county as being treated if its geographic centroid lays within a 100 miles from a Phase I power plant. To potentially rule out local economic effects, we omit all counties that have an actual power plant --- we thus identify our effects entirely off mobile pollution in counties beyond those with the plant itself. Although this binary classification ignores the intensity and directionality of emissions, it allows for simple graphical presentations of our results. Panel A of Figure 1 shows a map of treatment areas using the binary classification. In future robustness checks, we will allow effects to vary by a) pollution intensity of the regulated plant, and b) wind direction to the plant. Table 1 shows summary statistics for the treatment and control counties using the binary classification of treatment at a distance of 100 miles.

4.3 Econometric model

To identify the effect of the ARP program on mortality and SO2 levels, we estimate the following equation:

$$y_{ct} = \beta_t + \Sigma_t \beta_t * \tau + \theta X_{ct} + \alpha_c + \varepsilon_{ct.}$$

The variable y is either log of the mortality rate or the SO2 level in county c in year t. The vector X_{ct} includes the weather variables to address the possibility of differential trends in treatment counties.[5] County fixed effects are captured by α_c, while the error term (ε_{ct}) includes an

[5] We control for precipitation and precipitation squared, as well as the fraction of year in one of six 10-degree Fahrenheit bins (<30 F, 30-40, 40-50, 50-60, 70-80, and >80 F), leaving 60-70 as the omitted category.

idiosyncratic component as well as a term clustered on the county to allow for arbitrary serial correlation within a county.

Our treatment variable, τ, is an indicator for whether county c is within 100 miles of a regulated power plant. The treatment variable varies separately by year (β_t), where we omit the year 1994.[6] This amounts to a difference-in-differences estimator where we compare the change in mortality in treatment areas over time to the change in mortality in control areas relative to the baseline year of 1994. For evidence of a long run effect, we expect the estimates to increase in years since the ARP began ($t \geq 1994$). We also obtain estimates separately by year for the years before the ARP ($t<1994$) to test the core identifying assumption.

Demographic and economic covariates could correlate with treatment if regulation of power plants coincides with migration of different groups. While SEER data provide information on race and age groups that we could include as controls, data are inter-Census estimates that can be problematic for smaller sample cells. They may also be ``bad controls'' in that they may change as a result of treatment via migration or mortality. Instead, we opt for a propensity-score methodology, where we use pre-period county-level information from the 1990 County Data Book to predict probability of treatment. We include a number of estimates of county-level urbanization, general health and economic effects, and demographic distribution. Specifically, we estimate a logit model of treatment using the following economic and demographic variables (in both linear and quadratic forms): land area, population density, average travel time to work, income per capita, population share by race (black, white, and Hispanic), population share by age group (< 5, 5-17. 18-20, 21-24, 25-34, 35-44, 45-54, 55-64,65-74, and > 75), active physicians, hospital beds, and births per capita, percentage civil employment, and share employed in health services, retail, FIRE, and manufacturing. Appendix Figure A-1 shows the distribution of

[6] We also present results where we focus solely on a post-treatment indicator.

estimated p-scores. We trim to a common distribution for treatments and controls, meaning we drop all counties outside the vertical red lines on Figure A-1. Panel B of Figure 1 shows the remaining counties used and their respective p-score ranges: darker shaded counties are counties the model predicts are more likely to be similar to treatment counties.

We control for such similarity in our regressions by including this p-score interacted with year fixed effects. In essence, this flexibly allows for counties of similar p-score (and thus similar 1990 traits) to have similar changes in mortality outcomes over time. As robustness checks, we allow the effect of the initial demographic variables to vary over time by interacting them with β_t as well. We also control for state-specific trends in two robustness checks. However, these controls potentially introduce bias given dynamic treatment effect (Wolfers, 2006).

5. Results

5.1. SO2 results

An assumption for obtaining long-run effects is that the decrease in SO2 was immediate and persistent. The immediately binding cap and fixed costs involved with complying with these regulations support the plausibility of such a pattern. As shown in Panel A of Figure 2, before the beginning of the ARP, ambient SO2 pollution levels trended downward, for both treatment and control areas. The introduction of the ARP coincides with an acute 25% drop in SO2 levels in 1995. Levels returned to trend after that, with a small decline in SO2 at the beginning of Phase II in 2000. Two features emerge from this graph. One, SO2 levels trend down over time in both treatment and control areas. Two, while SO2 levels in control areas evolve continuously around the start of the ARP, SO2 levels in treatment areas drop suddenly when the ARP began, and

return to trend after. This immediate drop, followed by a return to trend, is the identifying variation we use to estimate long run effects.

Panel A of Figure 3 plots the event study coefficient estimates from the above equation, which show differential SO2 levels across treatment and control areas, adjusting for county and year fixed effects, weather factors, and p-score covariates. Aligning with the raw data in Figure 2, prior to the ARP, there is no trending difference in SO2. Relative SO2 levels drop rapidly after the ARP and the difference remains flat with a few years after the start at a relatively 1 ppb lower level than before. These patterns are comparable to the raw plots, suggesting covariate adjustment has little effect on our estimates.

As further evidence of the persistence of SO2 changes, we present several scatter plots of plant-level SO2 over time by focusing on plant emissions. Appendix Figure A-2 shows a scatter plot of 1985 vs 1990 emissions. Given that the ARP was not yet enacted, the points lie on a 45 degree line, as expected. Performing the same plot for 1990 vs. 1995, however, we see a significant flattening of this line, indicating the high initial emitters experienced large drops after the ARP. Finally, comparing 1995 vs. 2000 and 2005 emissions, we see effects similar to our 1995 comparison, indicating relatively little change after the ARP was already in place.

5.2. Mortality results

Given the demonstrated "first stage" in SO2 reductions, we now explore the reduced form effect of the ARP on mortality. To capture effects on the working population, we currently focus our analysis on individuals age 35-64, though we later expand to other age groups of interest. All outcomes are inverse hyperbolic sine (IHS) of deaths per 100,000 of the relevant age-group population at the county level, meaning coefficients represent an approximate

percentage change in deaths per 100,000.[7] Panel B of Figure 2 shows the raw plot for all internal causes of death by treatment and control, weighted by population. Prior to the ARP, mortality in both areas is trending down in a parallel fashion. After the ARP, the mortality rate in treatment areas experiences a shift in trend, decreasing more rapidly and eventually passing below control areas. That the effect presents as a change in trend rather than a discontinuous break suggests the ARP reduced mortality gradually over time rather than in a single shift. This observation motivates our choice of regression design below.

In Panel B of Figure 3, we present the coefficients based on our main regression equation. For mortality results, we weight by county population in 1990 unless otherwise noted (we use 1990 values rather than current values as current population may be endogenous to treatment). We cluster standard errors at the county level. Consistent with the raw plots, we see no difference in trends before the ARP. After the ARP, we see a gradual and steady decrease in mortality in treatment areas (relative to controls), suggesting a long run effect of the ARP on mortality.

Panel A of Table 2 shows these results in a basic difference-in-difference regression form. Column 1 includes only county and year effects, Column 2 includes the control for p-score-by-year interactions, and Column 3 adds weather effects (matching the event study graphs above). Consistent with our visual demonstration, addition of covariates and p-score controls does little to change the main effect. Using Column 3 as our preferred specification, we find that, on average, internal mortality in the 35-64 population was 2.5% lower after the beginning of the ARP.

[7] The inverse hyperbolic sine function, IHS(x) = ln(x + sqrt(1 + x^2)), has properties similar to the natural log function with the advantage of being able to handle zeros.

An examination of the graphs shows a basic difference-in-difference does not accurately capture the effect of the ARP. Specifically, the mortality changes appear to be gradual over time rather than instantly discontinuous. To better model this effect, Panel B shows our adjusted difference-in-difference, which allows for two parts to the effect: a discontinuous break (Treatment X Post) and a shift in trend (Treatment X Post X Years Since ARP). The joint p-value shows the joint test of significance of these two coefficients, which in each case is below 5%. To help interpret the long-run effects, we calculate the combined estimated effect by 2005, the end of our sample period. Using Column 3, two factors emerge. First, the coefficient on Treatment X Post is economically and statistically insignificant – the ARP did not cause a large detectable *break* in mortality levels. However, the coefficient on the "shift in trend" variable is statistically significant at 1%, and suggests that in each year after the ARP, mortality among the relevant age group in treatment counties decreased 0.26% relative to control counties. By 2005, this translates to a 3.8% decrease in mortality, or an avoided 44,000 deaths in 2005. As with the earlier difference-in-difference specification, results are robust across the addition of covariates.

5.3. Results by cause of death

To test the hypothesis that most of the health effects arise due to particulate matter, we explore effects by cause of death. SO2 affects primarily respiratory health, whereas PM affects mortality through changes in cardiovascular functioning. Therefore, we expect to see a greater relative decrease in cardiovascular death, with limited respiratory deaths, if the effect is from PM rather than SO2 itself. Panels C and D of Figure 3 show precisely this pattern: the shifts in cardiovascular disease largely mirror those in all internal cause mortality, with little observable

effect for respiratory deaths, though results are much noisier. The pattern over time also follows the pattern for all internal deaths, further supporting a long-term effect.

Table 3 follows the design of Table 2, with a focus on the "Treatment X Post" and "Treatment X Post X Years Since ARP" model. Cardiovascular mortality follows a pattern similar to that of overall deaths with a small immediate effect and a growing long-term effect. By 2005, there is a 5.3% reduction in cardiovascular mortality, or an estimated 20,000 fewer deaths. Estimates are slightly more sensitive to the addition of p-score controls, but remain within a standard error of the "no controls" model.

5.4. Testing for sorting

A key assumption for our identification strategy is that people did not sort in response to the ARP by changing residential locations. This could complicate identification if, for example, the ARP caused more unhealthy groups to migrate out of treatment areas into control areas, which would decrease relative deaths in treatment areas despite no true improvement in health due to air quality shifts. We test this using several approaches. First, we use total population (currently the denominator in our equation) as the dependent variable. If we see significant changes in population levels in response to the ARP, this would be suggestive of sorting. Shown in Panel A of Figure 4, however, population moves smoothly both before and after the ARP, suggesting no observable aggregate migratory patterns.

The ARP still may have caused changes in the composition of the population, with unhealthy out-migrants being replaced by healthy incomers. Although we do not observe demographic characteristics in the treatment and control counties at an annual frequency, we perform several indirect assessments to test for compositional changes. First, we explore the

effect of the ARP on vehicular deaths. Vehicular deaths are caused, in part, by the safety of automobiles. Since automobile safety is a normal good, and this time period experienced rapid improvements in automobile safety (air bags, ABS breaks), a change in population composition should lead to a change in vehicular deaths. Shown in Panel B of Figure 4, we observe no shifts in trend or levels for vehicular deaths. In a related test, we explore the effect of the ARP on per capita income in Panel C of Figure 4. Consistent with the results for vehicular deaths, we do not find a consistent effect of the ARP on mean income levels.

In regression form, vehicular deaths regressions (Table 3) show the unexpected pattern that the instant effect is negative, while the lasting effect is positive. While both are individually significant, they are jointly significant at only 5.9%. A visual examination of the data helps explain this odd pattern: the event study for vehicular deaths shows that in only one year post-ARP is there a statistically significant (positive) difference, and for the majority of the time effects are noisy. The econometric model, in attempting to fit this pattern, applies what is an unlikely representation of the truth (hence the insignificant joint test), and a standard difference-in-difference result as in Panel A of Table 2 finds a noisy zero.

Last, we investigate the effect of the ARP on infant mortality. Although pollution has been shown to affect infant mortality, it cannot affect infant mortality in the long run because infants are only infants for one year. However, we could witness gradual changes in infant mortality in there is sorting. For example, race is an important predictor of mortality. If there is a trending significant change in the racial composition of an area, we expect to find a change in infant mortality following a similar patter. Shown in Panel D of Figure 4, we find no evidence of gradual effects on infant mortality, and any suggested effects are not building over time Together, these results suggest no observable sorting in response to the ARP, and we can

interpret our long-run results for mortality as a causal effect of the pollution reductions from the ARP. Table 5 shows regression results for different age groups: infants, the young (1-25), the above 65 population, and all. For all internal mortality, only the over 65 group has a detectable effect. Cardiovascular results for infants are very large and significant at 10%, but given the rarity of infant deaths classified as cardiovascular this may be statistical noise.

5.5. Robustness checks

We present a number of alternative specifications to our primary model in Table 4. Column 1 adds state-specific linear time trends, and Column 2 expands this to quadratic trends. This does decrease the size of our overall effect to around 2.7-2.8% by 2005. Column 3 includes an additional interaction between the percentage rural by county in 1990 and a year fixed effect: as many of our treatment counties are located in the Midwest with a higher rural population, some of the change in mortality may be due to changing patterns across rural and urban areas. Adding such controls does little to change our result. In Column 4, we omit propensity score controls in favor of controlling for demographic and economic covariates. As these covariates may be endogenous, we instead include 1990 levels interacted with a year fixed effect. For example, rather than control for income per capita, we include income per capita in 1990 X year FEs. We include income per capita, medicare payments per capita, share of population that is male, black, and of other race, and share of population between 35-44, 45-54, and 55-64. The result is similar. As a future additional check, to allow for directionality of treatment, we plan to incorporate wind data.

6. Conclusion

Using the Acid Rain Program as an effective shock to ambient air quality, we present some of the first quasi-experimental, large-scale evidence of long-run health effects of exposure to air pollution. The ARP provides a unique opportunity to study such effects: economic changes were limited given rapid adaptation of new technologies, and the migratory nature of SO2 and related particulate matter meant many of the pollution gains were far enough away from plants that local economic and demographic shifts are not a concern. Among the primary working population (ages 35-64), internal mortality decreased with the beginning of the ARP, with growing reductions over the following 10 years. Our estimates suggest that by 2005, baseline internal mortality rates were 3.8% percent lower in treated counties, resulting in 44,000 fewer deaths in 2005. We find no changes in vehicular mortality, suggesting limited shifts in general demographics correlated with general health behaviors.

References in process

Figure 1
Treatment and control counties

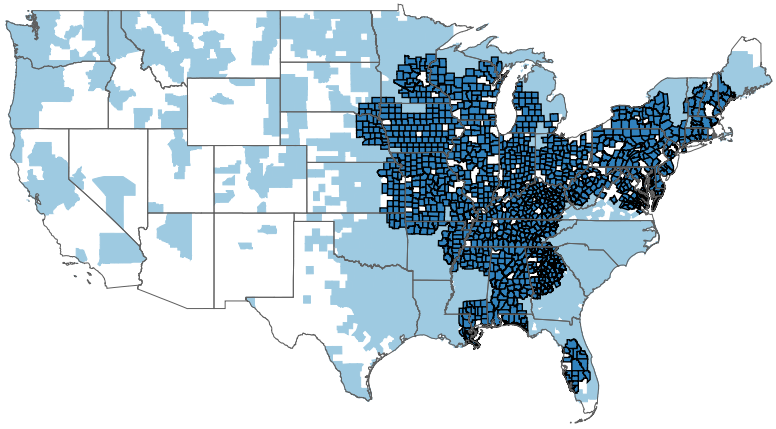

Counties by p-Score

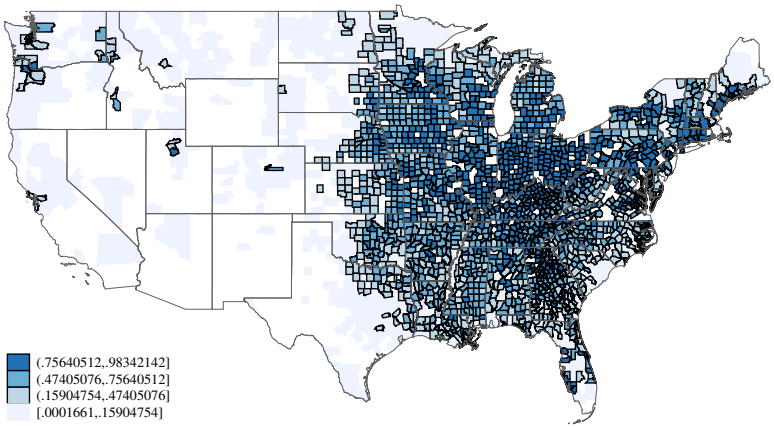

Figure 2
Treatment and Control Mean Comparisons (county FE-adjusted): Treatment = Within 100 miles of Phase I Plant

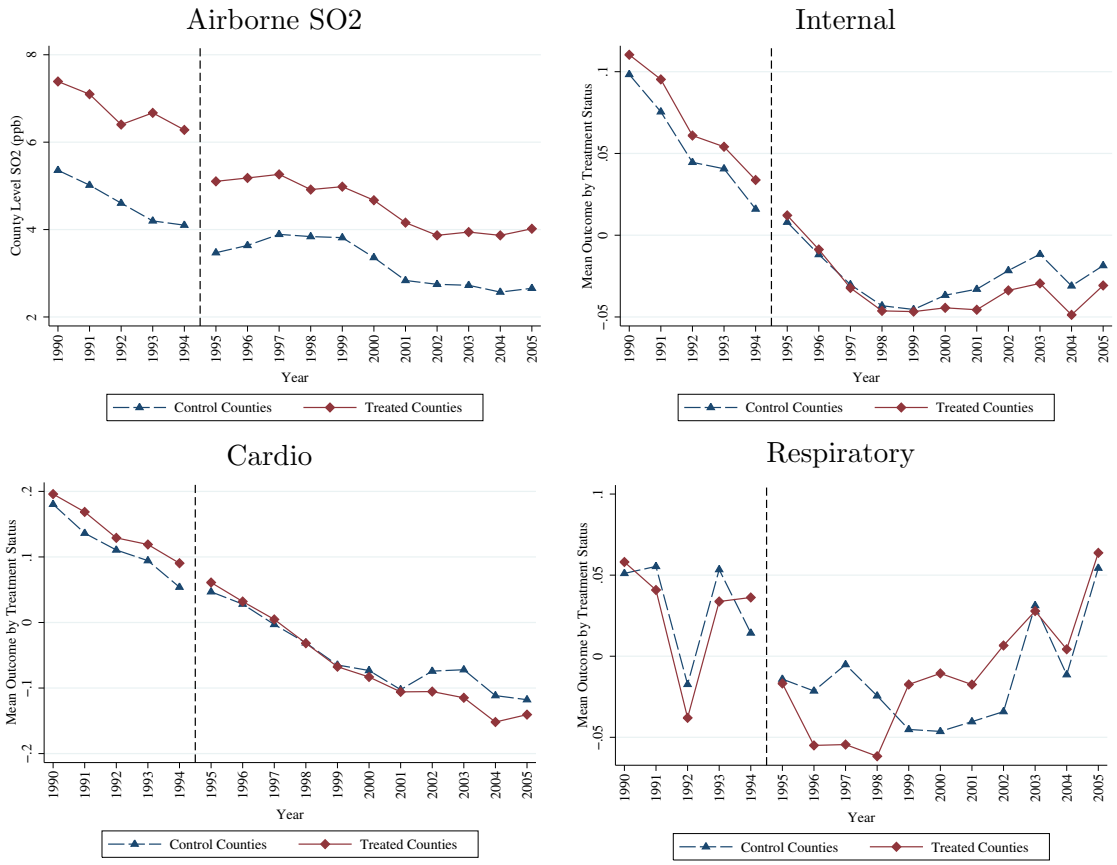

Notes:

Figure 3
Event Study by Outcome: Treatment = Within 100 miles of Phase I Plant

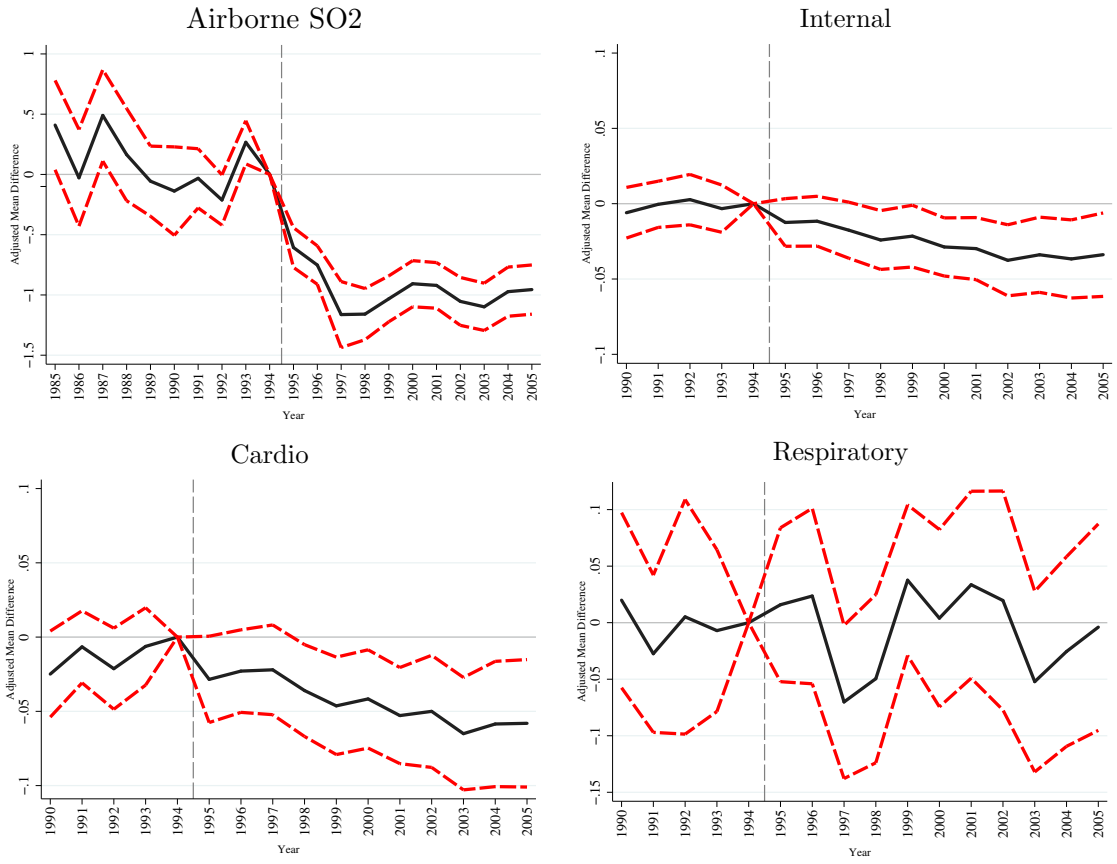

Notes:

Figure 4
Event Study by Outcome: Treatment = Within 100 miles of Phase I Plant

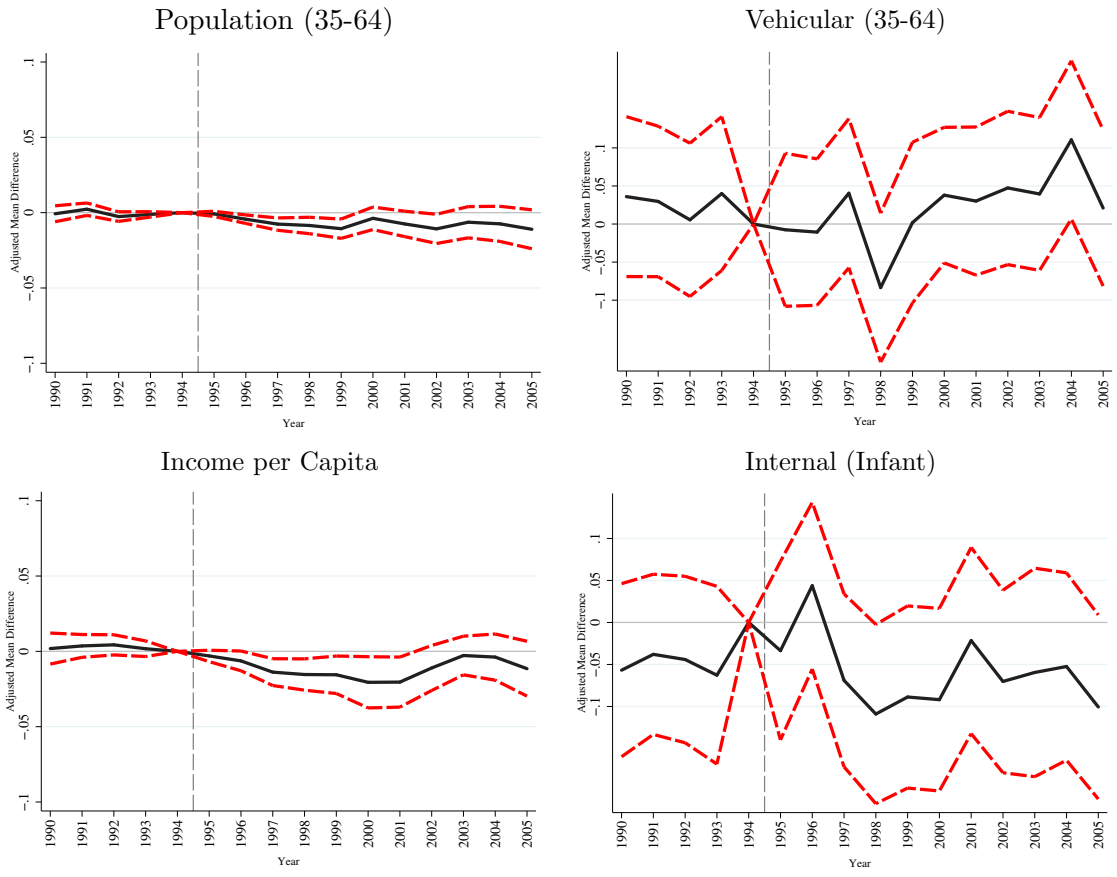

Table 1
Summary Statistics - 1990 Levels

	Control	Treatment
Male	0.49	0.48
Black	0.11	0.12
Age 0	0.02	0.02
Age 1-14	0.36	0.34
Age 15-34	0.33	0.34
Age 35-64	0.12	0.13
Age 65 and up	0.12	0.13
Income per Capita	24667.05	25551.42
Medicare per Capita	776.23	804.48
Wage Employment Rate	359.23	356.65
Population	73965.74	83242.85
Deaths		
Deaths per 100,000	736.96	848.56
Cardio per 100,000	340.00	396.33
Vehicular per 100,000	22.67	18.75

Table 2
Reduced Form Effect of Being Within 100 Miles of a Phase I Plant

	(1)	(2)	(3)
Panel A: Basic Difference-in-Difference			
Treatment X Post	-0.0231***	-0.0244***	-0.0248***
	(0.008)	(0.009)	(0.009)
Panel B: Difference-in-Difference With Trending Effect			
Treatment X Post	-0.0099**	-0.0082	-0.0089
	(0.004)	(0.006)	(0.006)
Treatment X Post X (Years since ARP)	-0.0022**	-0.0027**	-0.0026**
	(0.001)	(0.001)	(0.001)
Joint p-value	0.0092	0.0203	0.0159
Annual Effect by 2005	-0.0341	-0.0378	-0.0381
Annual Avoided deaths by 2005	-39060	-43386	-43668
County FEs	X	X	X
Prob. Weights		X	X
Weather			X
Counties	2435	2435	2435
Observations	38960	38960	38960

Table 3
Reduced Form Effect of Being Within 100 Miles of a Phase I Plant: Cause of Death

	(1)	(2)	(3)
Panel A: Cardiovascular			
Treatment X Post	-0.0058	-0.0052	-0.0073
	(0.008)	(0.010)	(0.010)
Treatment X Post X (Years since ARP)	-0.0052***	-0.0044***	-0.0041**
	(0.001)	(0.002)	(0.002)
Joint p-value	0.0002	0.0203	0.0218
Annual Effect by 2005	-0.0633	-0.0531	-0.0527
Annual Avoided deaths by 2005	-23865	-19932	-19785
Panel B: Vehicular			
Treatment X Post	-0.0388	-0.0479*	-0.0514*
	(0.025)	(0.028)	(0.028)
Treatment X Post X (Years since ARP)	0.0076**	0.0074**	0.0083**
	(0.003)	(0.004)	(0.004)
Joint p-value	0.0489	0.0951	0.0585
Annual Effect by 2005	0.0451	0.0341	0.0395
Annual Avoided deaths by 2005	2882	2186	2532
County FEs	X	X	X
Prob. Weights		X	X
Weather			X
Counties	2435	2435	2435
Observations	38960	38960	38960

Table 4
Reduced Form Effect of Being Within 100 Miles of a Phase I Plant: Robustness Checks

	State Trends	Quad State Trends	Share Rural by Year	Interacted Covariates
Panel A: All Internal				
Treatment X Post	-0.006	-0.0024	-0.0111**	-0.0105**
	(0.006)	(0.006)	(0.005)	(0.005)
Treatment X Post X (Years since ARP)	-0.0019**	-0.0023**	-0.0029***	-0.0020**
	(0.001)	(0.001)	(0.001)	(0.001)
Annual Effect by 2005	-0.0265	-0.0276	-0.0428	-0.0324
Annual Avoided deaths by 2005	-30210	-31487	-49267	-37111
Panel B: Cardiovascular				
Treatment X Post	-0.0027	-0.0091	-0.0112	-0.0064
	(0.009)	(0.010)	(0.008)	(0.008)
Treatment X Post X (Years since ARP)	-0.0028*	(-0.002)	-0.0050***	-0.0032***
	(0.001)	(0.002)	(0.001)	(0.001)
Annual Effect by 2005	-0.0334	-0.0317	-0.0658	-0.0419
Annual Avoided deaths by 2005	-12407	-11757	-24838	-15623

Table 5
Reduced Form Effect of Being Within 100 Miles of a Phase I Plant: Alternate Age Groups

	Infants	Young	Age 65 and up	All
Panel A: All Internal				
Treatment X Post	-0.0202	-0.0119	-0.0108***	-0.0066
	(0.024)	(0.017)	(0.004)	(0.007)
Panel B: Cardiovascular				
Treatment X Post	-0.1651*	0.0324	-0.0069	-0.0016
	(0.100)	(0.027)	(0.006)	(0.008)

Figure A-1
Propensity Score Overlap

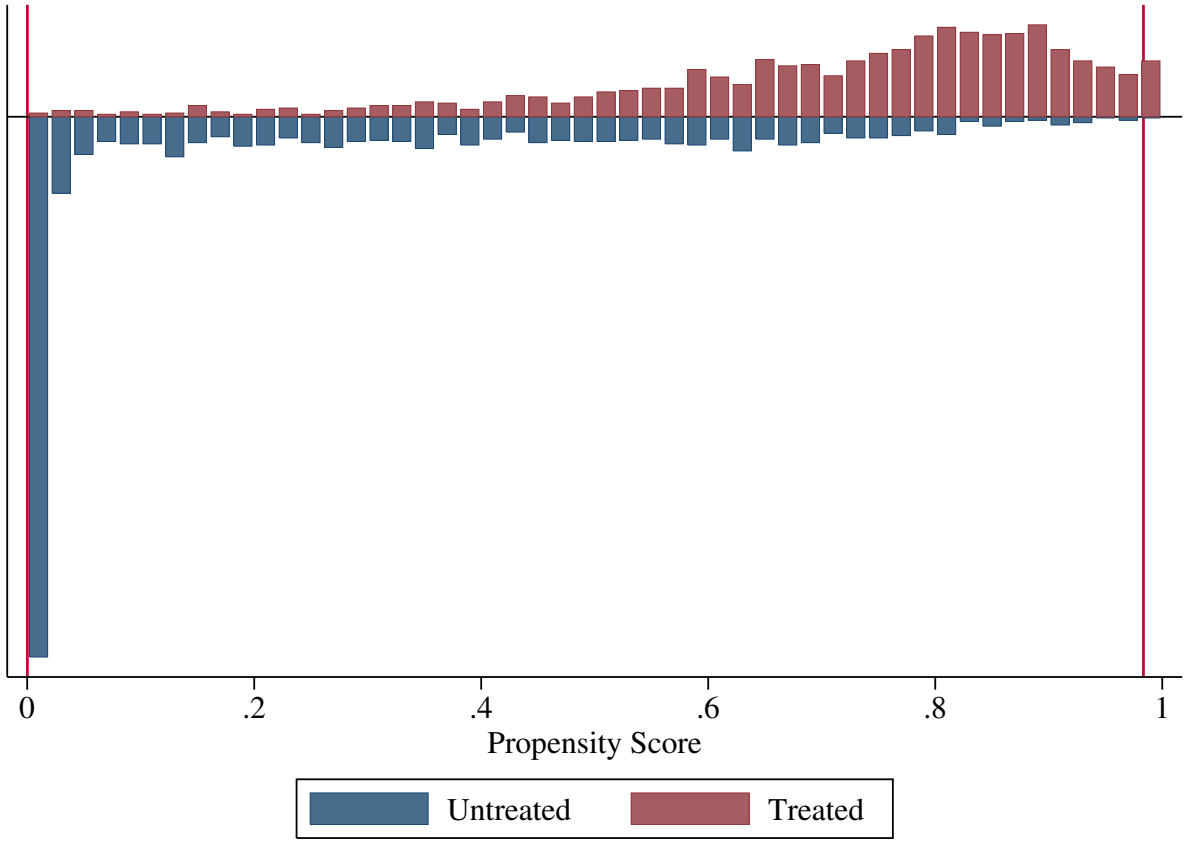

Figure A-2
SO2 Emissions from Phase I Power Plants by 1985 BTUs

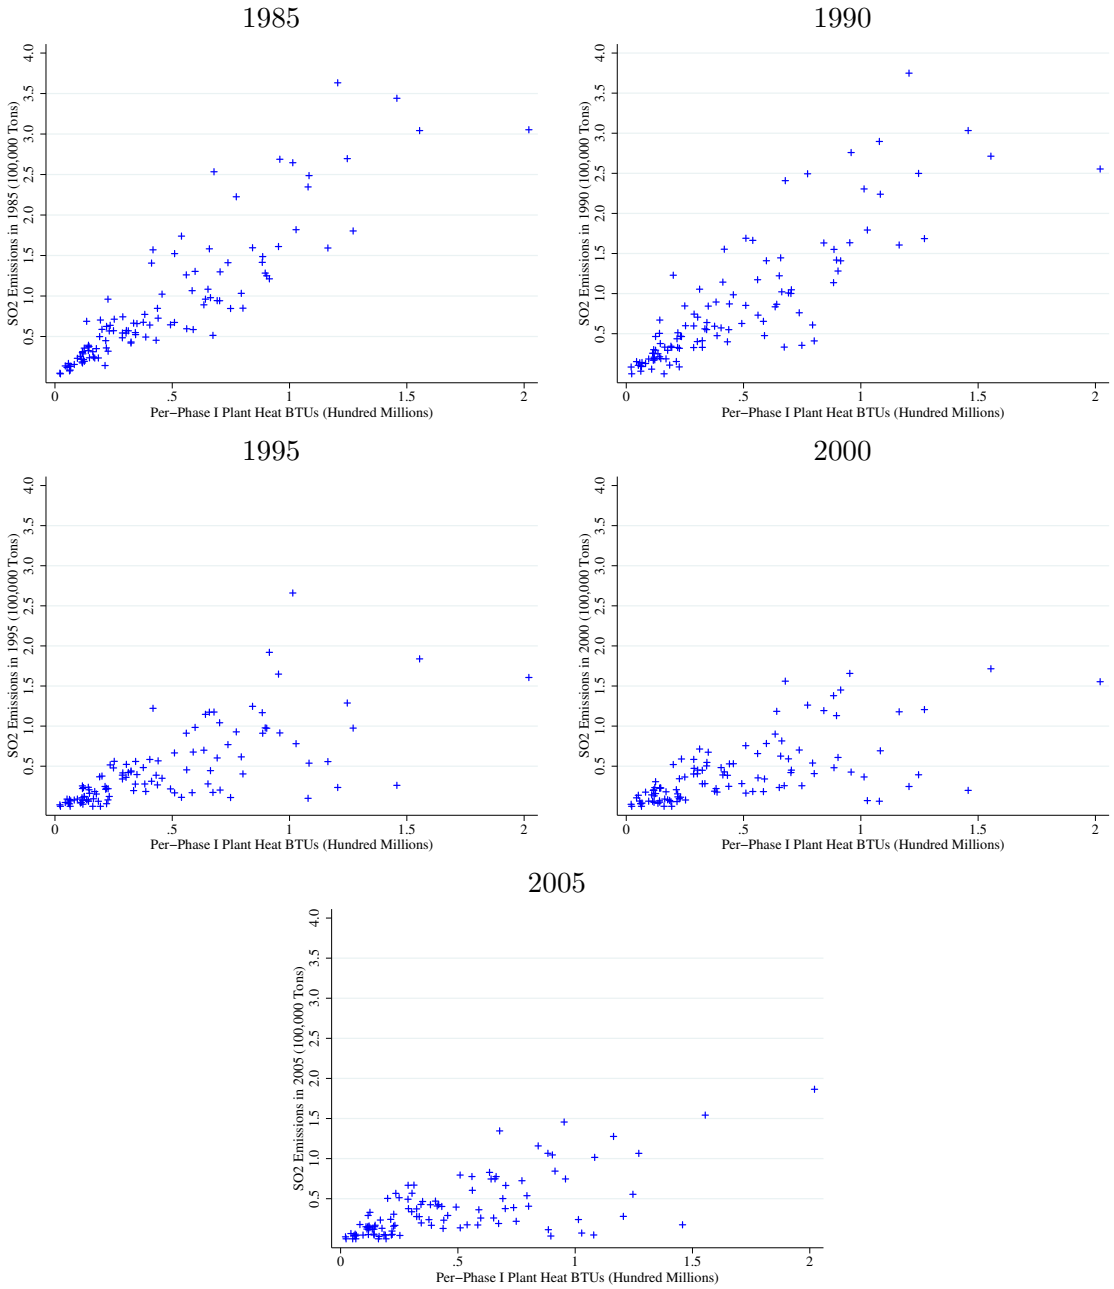

Notes: Point source emissions data from the EPA Air Markets Program Data. Horizontal axis shows reported heat input, in British Thermal Units (BTUs), in hundreds of millions, for each Phase I plant in 1985. Vertical axis shows reported SO2 emissions for each respective Phase I plant in the noted year. Each point represents a different Phase I plant in the given year.