Journal of Public Economics 200 (2021) 104440

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

Long-run pollution exposure and mortality: Evidence from the Acid Rain Program ☆

Alan I. Barreca [a], Matthew Neidell [b], Nicholas J. Sanders [c,*]

[a] UCLA, IZA, and NBER, United States
[b] Columbia University, IZA, and NBER, United States
[c] Cornell University and NBER, United States

ARTICLE INFO

Article history:
Received 29 September 2020
Revised 17 March 2021
Accepted 26 April 2021
Available online 18 June 2021

Keywords:
Acid Rain Program
Pollution and mortality
Atmospheric transport model

ABSTRACT

We estimate the effects of long-run pollution exposure on mortality by exploiting the United States Acid Rain Program (ARP) as a natural experiment. We use a difference-in-differences design to compare changes in adult mortality over time driven by installations of sulfur controls on power plants, combined with a model of atmospheric pollution transport. We find that sulfur controls reduced pollution immediately, with smaller relative improvements in the following years. Mortality reductions started small and grew steadily, suggesting cumulative health effects over time. We also find persistent mortality effects for those 35–64 years of age, suggesting the ARP had large productivity gains for the working-age population.

1. Introduction

While the Environmental Protection Agency (EPA) attributes over 90 percent of the monetized benefits from environmental quality improvements to changes in long-run exposure (Dominici et al., 2014), nearly all quasi-experimental estimates linking pollution and health are based on short-run pollution changes. Exogenous, temporary shocks to pollution can help with causal identification (Graff Zivin and Neidell, 2012; Deschenes et al., 2017; Schlenker and Walker, 2016; Deryugina et al., 2019), but fail to identify longer-run cumulative effects, as they largely affect the timing of exposure without substantively affecting lifetime exposure. Behavioral responses to these shocks, such as residential sorting and employment changes, complicate identification of long-run effects (Walker, 2013; Banzhaf and Walsh, 2008). Therefore, even if an initial change in pollution is exogenous, exposure is likely to become endogenous over time. Quantifying effects from long-run exposure requires a long-lasting, exogenous change in ambient pollution without accompanying behavioral responses.

In this paper, we estimate the effects of long-run pollution exposure on mortality by exploiting the United States Acid Rain Program (ARP), a cap-and-trade regulation to control sulfur dioxide (SO2) emissions, which avoids many of the issues associated with identifying long-run effects. First, the ARP created an immediate and persistent drop in SO2, a precursor gas in the formation of particulate matter smaller than 2.5 micrometers (PM2.5)[1]; PM2.5 is a pollutant with detrimental effects on human health (see Environmental Protection Agency (2004) for a comprehensive review). This one-time drop in pollution allows for an event study analysis to explore dynamic effects.

Second, the ARP regulated only certain SO2-intensive coal plants. This lends to a design comparing changes in mortality over time in counties most impacted by pollution from regulated plants to changes in mortality in counties with little to no related exposure. This mitigates bias from national trends changing over time independent of pollution from coal-fired plants, such as business cycles, health care access, and technological advances.

Third, the vast distance both SO2 and PM2.5 travel once airborne subsumes many potential confounding general equilibrium

☆ We thank Joshua Graff Zivin, Douglas Miller, Reed Walker, Michael Anderson, Shinsuke Tanaka, and seminar participants at UC-Santa Barbara, Indiana University, University of Pittsburgh, Toulouse School of Economics, Lafayette College, University of Connecticut, Louisiana State University, Auburn University, the 2016 ASHEcon Biennial Conference, the 2017 Meetings of the Allied Social Sciences Association, and the Stanford Institute for Economic Policy Research Annual Postdoctoral Scholars Conference for many useful comments. Neidell received support from the Columbia Population Research Center, which is supported by the Eunice Kennedy Shriver National Institute of Child Health & Human Development of the National Institutes of Health under award number P2CHD058486. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

* Corresponding author.
E-mail address: njsanders@cornell.edu (N.J. Sanders).

[1] See, for example, Lippmann and Thurston (1996) and Hodan and Barnard (2004).

https://doi.org/10.1016/j.jpubeco.2021.104440
0047-2727/

effects. Households sorting in response to the economic effects of the ARP will not bias estimates if such sorting remains within the effective range of pollution transport from a given region.[2] Given existing evidence that finds such changes in housing amenity values occur at distances of less than 2 miles, long-distance residential sorting is likely to be limited.[3]

Fourth, the ARP generated limited economic spillovers. While many environmental regulations often lead to job loss (Greenstone, 2002; Walker, 2013), which can have independent effects on health (Sullivan and von Wachter, 2009), this issue is less relevant to the ARP as compliance costs and economic effects were low (Schmalensee and Stavins, 2017).[4] Moreover, the broad spread of emissions beyond plants themselves subsumes local economic effects, to the extent they exist, in the same way they deal with sorting. We verify the ARP had negligible effects on income, employment, or migration, suggesting economic effects are unlikely to bias estimates of health effects.

Our main research design involves an event study analysis spanning 1985 through 2005. Treatment begins in 1995 with enforcement of the ARP, and ends in 2005, before changes to the SO2 permit market potentially shifted the relationship between regulation and ambient emissions. To construct a measure of each county's exposure to treatment, we combine information on baseline plant efficiency, sulfur control installations (e.g., sulfur scrubbers), and atmospheric migration of SO2. Since evidence suggests PM2.5 drives mortality effects, we use a model of pollution transport to project how SO2 disperses to form PM2.5 across the country (Muller et al., 2014). This treatment variable provides a basis for a difference-in-differences model with varying exposure intensity. We estimate the dynamic, reduced-form effects of the ARP, tracing out the effects of the policy for up to 10 years, while also testing for pre-trends in mortality leading up to the ARP as a check on identifying assumptions.

Our results suggest reductions in long-term pollution concentrations from the ARP contributed to declining mortality rates. Ambient SO2 emissions exhibit a large drop in 1995, with considerably smaller reductions in the following years. Mortality, on the other hand, declines more steadily over time, suggesting a cumulative effect from lower sustained pollution levels. By 2005, a standard deviation increase in treatment exposure reduces all-age cardiorespiratory mortality rates by approximately 0.4 percent. Effects are largest for the age 35–64 population, in relative terms, with a standard deviation increase reducing annual cardiorespiratory mortality rate by 1.6 percent.

Our results build on recent quasi-experimental studies addressing long-run exposure to pollution (Chen et al., 2013; Ebenstein et al., 2017; Anderson, 2019). First, we provide a year-by-year event study of changes in mortality, providing suggestive evidence that reductions in cumulative pollution exposure drove health improvements over time. Second, we provide estimates for working age adults, a group where mortality represents a significant loss in life expectancy and economic productivity, meaning existing estimates of the value of statistical life (VSL) more readily apply.[5]

These estimates of the mortality effects of coal pollution are of direct interest to policies centered on emissions from coal plants. Global coal consumption has nearly doubled since the turn of the century with rapid growth in developing nations such as China and India[6]. In the United States, proposals targeting carbon emissions often focus heavily on the role of coal plants, with potentially large regulatory costs. Our estimates help in properly assessing the mortality benefits of reducing coal plant emissions.

2. Background

2.1. The Acid Rain Program

The ARP was a provision of the EPA Clean Air Act Amendments of 1990. The two-phase program regulates SO2-producing coal power plants, with the goal of reducing SO2 levels to 50 percent of 1980 levels. Phase I, beginning in 1995, regulated the 110 power plants with the highest SO2 emissions, as measured in 1985. Phase II, beginning in 2000, further limited emissions of Phase I plants and added all remaining coal plants to the program. Both Phase I and Phase II worked through an SO2-emission cap-and-trade system. Plants could bank permits across time and buy and sell permits with other plants. Phase I had a considerably larger impact on SO2 emissions than Phase II, as firms had the ability to smooth Phase II regulations by banking permits in advance. This motivates our quasi-experimental research design, where we focus on the timing of Phase I regulations.

The EPA distributed SO2 allowances to 263 heating units at the 110 plants based on baseline (1985–1987) heat input to minimize regulatory gaming (Stavins, 1998). Heat input is a measure of the amount of fossil fuel used to generate electricity, calculated in British thermal units (BTUs). Each year, plants report SO2 emissions to the EPA for verification. For plants polluting in excess of held permits, the EPA assigns an initial fee of $2,000 (adjusted for inflation) per ton of overage and requires accounting for overages by purchasing sufficient permits. 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 compliance, leading to substantial decline in SO2 emissions from regulated plants.

2.2. Health effects of SO2 and PM2.5

The primary aim of the ARP was reducing acid rain via regulating SO2. Acid rain has no known direct impacts of human health, but the regulation likely improved health due to reductions in associated pollutants.[7] SO2 exposure health effects are primarily short-run and respiratory-related, leading to increases in hospital admissions for outcomes such as asthma exacerbations.[8] Much of the anticipated lasting health effects from the ARP arise from other pollutants correlated with SO2 levels, namely PM2.5. Through an atmospheric conversion process, SO2 produces sulfate particles, a

[2] Although the ARP focused on SO2, where health effects are short-run and respiratory-related, the eventual conversion to $PM_{2.5}$ is what likely drives mortality. Both the conversion process and harmful effects of $PM_{2.5}$ were largely unknown at the time, making it unlikely people sorted directly on $PM_{2.5}$. To the extent the two pollutants are correlated, by sorting on SO2 people may indirectly sort on $PM_{2.5}$.

[3] Using detailed micro-census data, Davis (2011) finds that power plant openings 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. Focused on toxic plants, Currie et al. (2015) find that plant openings 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. Walker (2013) and Greenstone (2002) show that employment effects of regulation exist at the county level, though in both cases the regulations they consider are county-level rather than plant-level as in the ARP.

[4] Regulatory costs were lower than predicted due to rapid technological innovation in desulphurization and an unexpected increase in access to low-sulfur coal.

[5] Existing studies typically focus on compensating differentials for workers across industries. See Murphy and Topel (2006) for a discussion of VSL over the lifecycle.

[6] "China and India drive recent changes in world coal trade," Today in Energy, U.S. Energy Information Administration, Nov. 20, 2015.

[7] See https://www.epa.gov/acidrain/effects-acid-rain, accessed March 9, 2017.

[8] See Environmental Protection Agency (2009) for more information.

portion of small-scale PM2.5.[9] PM2.5 penetrates deep into the lungs and enters the bloodstream, where it can lead to negative cardiovascular and respiratory effects.[10] Extended exposure to PM2.5 may cause pulmonary and systemic oxidative stress and inflammation, which can lead to vascular dysfunction, atherosclerosis, and altered cardiac autonomic function (Brook et al., 2010). Effects from long-run PM2.5 exposure may accrue after sustained, low-level exposure, similar in spirit to cigarette smoking.[11]

3. Data and methods

Multiple stages link the ARP to reductions in ambient pollution concentrations. Boilers at coal-fired plants produce SO2 emissions, which can convert to PM2.5 and disperse to surrounding areas. The ARP, by leading to the adoption of abatement techniques, reduces SO2 emissions from a power plant, which ultimately reduces PM2.5 in surrounding areas. Since measures of ambient PM2.5 concentrations – the pollutant most likely responsible for any mortality effects – do not exist throughout this entire period, we instead focus primarily on the reduced form relationship between ARP regulation-induced changes and mortality, though we also present various "first stage" results to highlight the effect of the ARP on pollution, the main channel for health effects.

Some plants responded to the ARP by installing new technologies on their boilers, thereby reducing SO2 emissions and associated PM2.5 concentrations. To examine this relationship, we would ideally have detailed measures of 1) installation of boiler controls; 2) boiler-specific SO2 emissions; 3) a model linking boiler SO2 emissions to ambient PM2.5 concentrations; and 4) ambient PM2.5 concentrations. We possess data on items (1) and (3), and partial data on (2); the EPA provides boiler-level SO2 emissions data every 5 years prior to the ARP and every year after the ARP. There exist no data on (4) in the years prior to the ARP, as PM2.5 was not officially regulated by the EPA until 1997, and reliable data on ambient concentrations before even 1999 is scarce, though we do have SO2 ambient concentration data from the EPA monitoring network.

Within these data limitations, we construct a reduced form measure of exposure to pollution reductions using baseline plant efficiency, boiler control installs and the transport of emissions from these boilers. Given most pollution control installations, and the most dramatic decreases in SO2 output, happen only once during the period of our study (as we demonstrate below), this yields a treatment exposure we can use to explore dynamic mortality effects.

3.1. Data sources

We use county-level cause of death-coded mortality data from the Centers for Disease Control (CDC) as our main outcome variable. These data 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. Restricted data include information on age at death and county of residence. As no Phase I plants exist further west than approximately −96 degrees in latitude, we only include counties with centroids east of −100 degrees latitude.

We obtain a list of all ARP plants and associated boilers from the EPA Clean Air Markets Program Data. Information include boiler-level SO2 emissions in tons, plant location, existence of installed sulfur control technology, and the date of install. We observe the year of sulfur control technology installations and the annual SO2 emissions for all years 1995 and beyond, but prior to the ARP annual plant-level SO2 emissions are only available every five years (1980, 1985 and 1990). In our discussion of the first-stage relationship between boiler sulfur controls and emissions, we assume constant emissions from the last year of available data (e.g., we assign 1980 levels to 1981–1984, 1985 levels to 1986–1989, and 1990 levels to 1991–1994). Interpolating data has no effect on our main mortality estimates, since we employ a reduced form estimation strategy with boiler controls as the main independent variable.

We model the dispersion of boiler-level SO2 emissions across the country using the second iteration of the Air Pollution Emission Experiments and Policy analysis (APEEP) atmospheric transport model (Muller et al., 2014). This model provides a county-to-county matrix that converts SO2 emissions into expected contributions to PM2.5 concentrations, adjusting for information on topography, wind direction, and other such region-specific factors. Rather than incorporate the predicted PM2.5 concentrations directly in our treatment variable, we use the predicted PM2.5 conversion rate as a weight to model the dispersion of boiler-level emissions.[12] We construct our measure of treatment intensity by interacting these conversion rates by the expected intensity of the pollution reduction due to sulfur controls. We also use the APEEP matrix to generate predicted PM2.5 levels, for which data do not exist, so we can approximate the marginal effect of PM2.5 on mortality.

Given SO2 was the primary target of the ARP, we first examine the impact on ambient SO2 levels using measures from the EPA monitoring network. The EPA reports daily SO2 averages in parts-per-billion (ppb). We convert daily monitor-level measures to annual county-level measures by inverse distance weighting of monitors within 50 miles of a county centroid. We use an unbalanced panel of sensors given the limited number active across the entire sample period. Analysis using a balanced panel yields similar results but with fewer applicable counties.

To account for the possible confounding role of weather, we flexibly control for temperature and rainfall using data from the National Oceanic and Atmospheric Administration (NOAA). We aggregate station-day variables up to the county-year level using inverse distance squared weights, using all monitors within a maximum distance to 100 miles between county centroid and weather stations. We employ data from the County Data Book (CDB) and the regional economic accounts from the Bureau of Economic Analysis (BEA) to measure economic and demographic factors.

3.2. Measure of treatment

We define our treatment variable, weighted exposure to sulfur controls (SC), as a combination of baseline plant efficiency, sulfur control installations on boilers, and emission transport weights.

[9] See "Sulfur Dioxide Basics", provided by the Environmental Protection Agency (available online at https://www.epa.gov/so2-pollution/sulfur-dioxide-basics, accessed February 20, 2020).

[10] See Environmental Protection Agency (2004) for a comprehensive review. Given its diminutive size, $PM_{2.5}$ also penetrates into buildings at a high rate, suggesting the ability to avoid exposure is quite limited (Thatcher and Layton, 1995; Ozkaynak et al., 1996; and Vette et al., 2001). For a review of quasi-experimental evidence of the health effects of short-run exposure to pollutants such as SO_2 and $PM_{2.5}$, see Graff Zivin and Neidell (2013).

[11] Note the effects from long-run *exposure* are distinct from the long-term *effects*. For example, several studies estimate the effect from early childhood exposure on adult outcomes. While the outcomes are long-term, the exposure remains short-run. See Sanders (2012), Isen et al. (2017), and Bharadwaj et al. (2017).

[12] For example, the matrix conversion rate between a ton of SO2 released in Franklin County, MO and a microgram per cubic meter (μg/m3) of PM2.5 in the same county is 1.87 X 10e-5. The conversion for nearby St. Louis, MO County (FIPS 29189) is 1.09 X 10e-6, and for further away Wayne County, MI (FIPS 26163) is 5.7 X 10e-7. EPA data on boiler-level emissions are in tons, and the transport matrix program natively uses short tons, requiring a basic conversion of 1.102311 short tons per ton.

Since we expect control installations to have larger effects on boilers that were initially "dirtier", we incorporate heat input in 1985, the metric used for initial assignment of ARP permits, as a measure of plant efficiency. Higher heat input translates to less efficient electricity production, with more fuel required per unit of energy output, and thus an expected higher pollution level, all else constant. Therefore, a sulfur control installed on a more inefficient plant should yield a larger improvement in emissions. The incorporation of the transport model allows the effects of the installed sulfur controls to vary geographically, where this variation is driven by the SO2-PM2.5 conversion matrix by county. This yields the following equation:

$$SC_{ct} = \sum_i f(\text{PhaseI Sulfur Controls}_{it} * heat_{i,1985}) \quad (1)$$

SC_{ct} is the measure of exposure for county c in year t, which accounts for sulfur control installations on Phase I boiler i present in year t throughout the U.S.. f() is the atmospheric transport model that weights by the strength of the relationship between emissions at plant i and PM2.5 in county c, $\Sigma_i()$ denotes the sum of all Phase I boilers with sulfur-controlling technology installed as of the specified year, and $heat_{i,1985}$ is the initial, plant-specific heat input, in BTUs, in 1985. To simplify interpretation of coefficients, we standardize the SC variable to have mean zero and standard deviation of one for our regressions.

For an example of (1), we discuss here the treatment intensity Franklin County, MO would receive from the top three outside county contributors with upgrades present by 1995: Gibson, IN (FIPS 18051), Carroll, KY (FIPS 21041), and Warrick, IN (FIPS 18173). In practice, emissions from more counties contribute, but we focus on the top 3 for illustrative purposes. The conversion matrix assigns the following receiving weights for Franklin, MO: 3.55 X 10e-7 for Gibson, IN, 2.8 X 10e-7 for Carroll, KY, and 2.72 X 10e-7 for Warrick, IN. Gibson, IN had one Phase I sulfur control in 1995, on a boiler with an initial heat input of 34.7 million (MM) BTUs. Carroll, KY had one Phase I sulfur control in 1995, on a boiler with an initial heat input of 26.1 MM BTUs. Warrick, IN had two Phase I plants with sulfur controls in 1995, with initial heat inputs of 13.6 MM BTUs and 5.8 MM BTUs. If SC_{ct} were based on only these three outside counties, in 1995 Franklin County would receive a value of:

$$SC_{\text{Franklin},1995} = [34.7\text{MM} * 0.000000355 + 26.1\text{MM} * 0.00000028 + (13.6\text{MM} + 5.8\text{MM}) * 0.000000272$$

As the majority of SO2 (and modeled PM2.5) reductions occur in 1995, we focus on upgrades by 1995 for our event studies. In an alternate reduced form model, we allow for annual variation in upgrades. We also explore models in which we weight installed sulfur controls by initial baseline SO2 emissions, models in which we treat all upgrades equally regardless of plant size, and models that allow for controls installed on both Phase I and Phase II plants. These changes have minimal impact on our main findings.

Fig. 1 illustrates how the policy affected plant behavior, and how we operationalize this to address data limitations. Panel A of Fig. 1 shows the running tally of installed controls on Phase I boilers, by month and year. The majority of Phase I control installs occurred either just before, during, or shortly after 1995, and

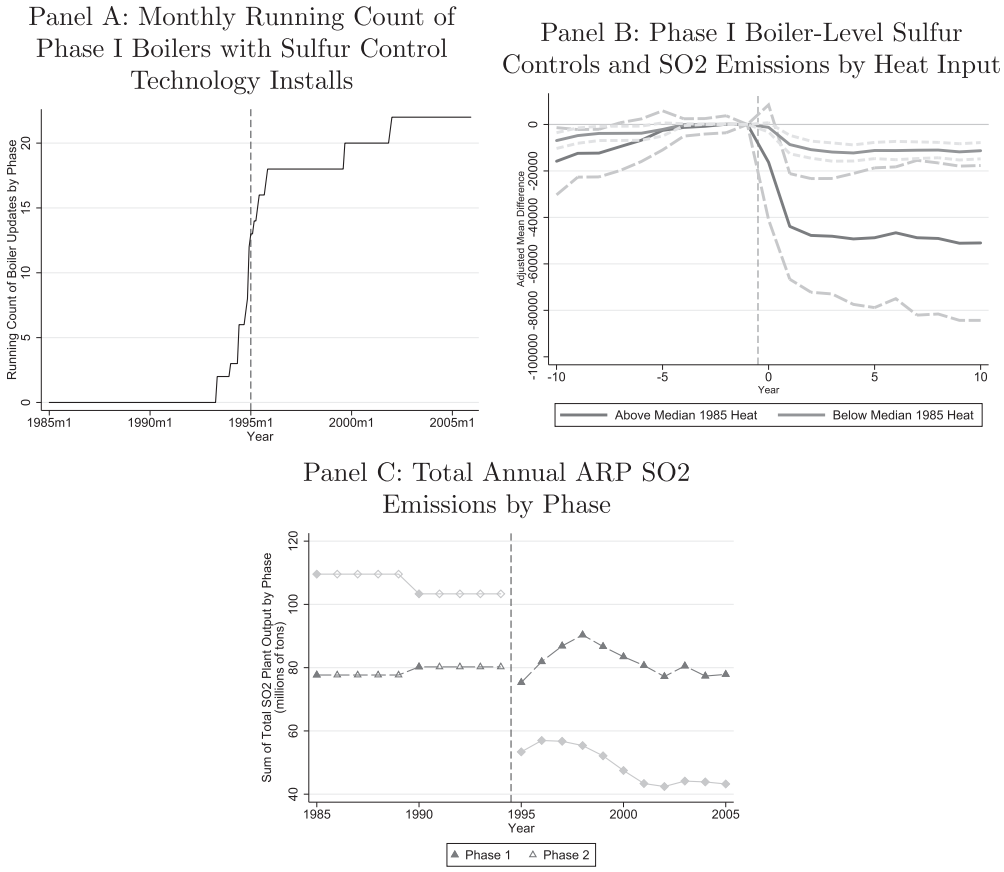

Fig. 1. Installation of sulfur controls and boiler emissions. Notes: See Section 3.2 for detailed discussion of each figure. All figures based on data from the EPA Clean Air Markets Acid Rain power plant data set. Dashed vertical line indicates the beginning of enforcement of the Acid Rain Program. Thick dashed lines of Panel B indicate 95% confidence intervals. In Panel C, data prior to 1995 are only available in 1985 and 1990. Hollow markers in 1986–1989 and 1991–1994 indicate imputation from most recent available power plant data — see Section 3.1 for details.

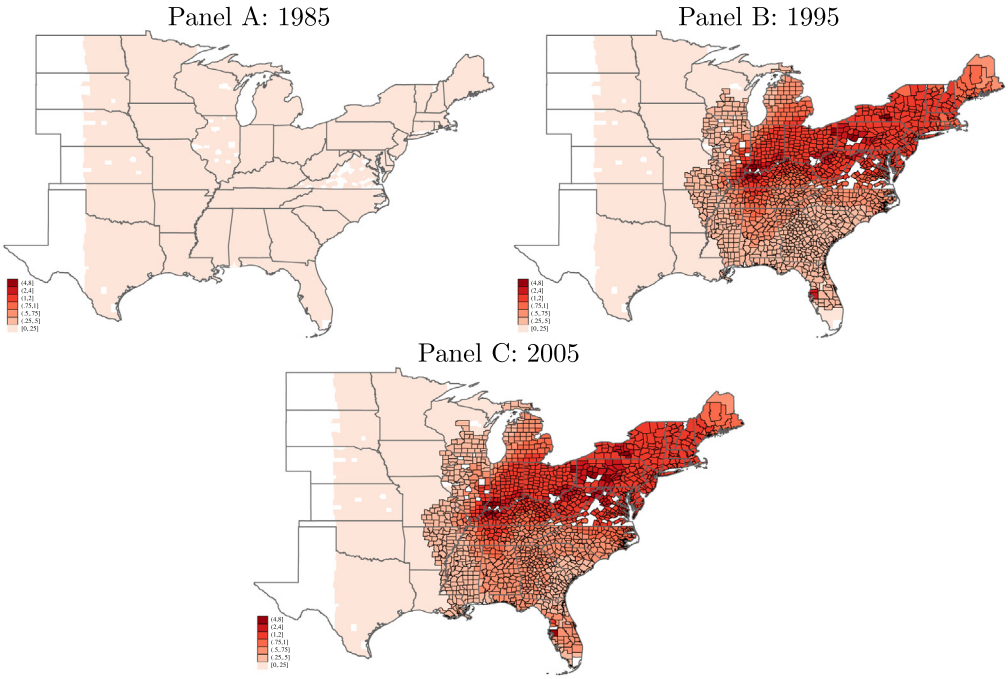

Fig. 2. Exposure to sulfur controls on phase I ARP boilers by selected year. Notes: Figures show county-level exposure to our measure of treatment: number of sulphur controls, weighted by initial heat output and the atmospheric transport matrix, which we describe in Section 3.2.

remain largely stable after that. While some installs occur before 1995, the clear majority happen in the first year of enforcement, and many of those that happen earlier occur in the final months of 1994, which we expect would have little effect on average pollution and mortality levels in that year.

To understand the impact of these sulfur control installations on emissions, Panel B of Fig. 1 shows an event study of Phase I boiler-level emissions, where the relevant treatment year (0) is the year in which plants install sulfur control technology.[13] The model, which controls for year and boiler fixed effects (with standard errors clustered at the boiler level), shows that a Phase I boiler with such technology installed saw immediate reductions in SO2 output. We split all Phase I plants with upgrades into above and below median level of heat input in 1985, effectively dividing the sample by boiler efficiency. For our treatment variable definition, we assume more inefficient boilers (higher heat input) will experience larger emissions reductions from installs of sulfur controls. Our estimates support this: the average boiler with above-median heat input reduced SO2 emissions by approximately 50,000 tons, while the average boiler below the median heat input reduced SO2 emissions by closer to 10,000 tons.

Panel C of Fig. 1 shows the trends in total boiler emissions by year and plant phase, which also illustrates that (1) the majority of reductions appear in 1995, and (2) effectively all reductions are due to Phase I plant behavior.[14] On net, Phase I plants reduced SO2 output by around 50 million tons in 1995, while Phase II plants saw largely no change, and even small increases early on.

Fig. 2 illustrates county-level values of our sulfur control exposure in 1985, 1995, and 2005. The figure highlights two important factors. First, exposure to Phase I plants with sulfur controls is zero in the beginning of our sample, with substantially higher levels across the country in 1995. Second, from 1995 to 2005, there is little change in county-level exposure to the number of boilers with sulfur controls, as most of the large additions occurred in the first year of the ARP.

3.3. Dynamic regression model

To identify the link between the ARP and mortality across time, we estimate the following equation:

$$y_{ct} = \lambda_t + \Sigma_t \beta_t * SC_c + \Sigma_t \Phi_t * Z_c + \gamma X_{ct} + \alpha_c + \varepsilon_{ct} \quad (2)$$

The variable y is the outcome of interest in county c in year t, often the inverse hyperbolic sine (IHS) of the age-adjusted mortality rate for four age groups (all ages, under 35, 35–64 and 65 and over).[15] The IHS function behaves similarly to the natural log function, but is robust to the inclusion of zero values, which occur when looking at mortality rates by subgroups and cause of death in smaller-population counties. For the denominator, we use county-level, age-specific annual population estimates from the Surveillance, Epidemiology, and End Results (SEER) Program. We focus on cardiorespiratory-related deaths, as this contains the outcomes most likely to be affected by PM2.5 exposure. We also present results for all-cause internal mortality in the Appendix, and perform a falsification test using external deaths (accidents, murders, etc.); such deaths may change due to potentially confounding factors such as selective migration and economic fluctuations, but not with pollution-related health improvements. We also estimate this model with ambient SO2 and modeled PM2.5 as the outcomes to demonstrate the "first stage" effect of the ARP on pollution levels.

The vector α_c captures county fixed effects, the vector λ_t captures year fixed effects, and the error term (ε_{ct}) includes an idiosyncratic component as well as a term clustered on the county to allow for arbitrary serial correlation within a county. We weight all regressions by county level age-specific population, using the population from the first year in our sample (1985) for all years

[13] Appendix Fig. A1 illustrates a similar idea, but splitting by Phase I and Phase II boilers rather than Phase I heat input, showing the majority of reductions are due to Phase I plants.

[14] We cannot directly verify the magnitude of change between 1994 and 1995, as plant-specific output data are not available between 1991 and 1994.

[15] See http://www.cdc.gov/nchs/data/statnt/statnt20.pdf (accessed February 2, 2020) for age adjustment using the 2000 U.S. Population.

to prevent endogeneity of weights (results are robust to alternate weight choices).

Our treatment variable, SC_c, is our normalized measure of county c's exposure to sulfur controls installed on Phase I boilers as of 1995. We consider how the treatment effect varies separately by year (β_t), omitting 1994 as the reference year. This amounts to a difference-in-differences estimator that compares the difference in mortality over time between areas with varied levels of boiler control exposure in each year relative to the baseline difference in 1994. The event study coefficients reflect the marginal effect of one standard deviation change in our treatment variable compared to the baseline year of 1994.

We interpret the pattern of estimates of the β_t coefficients as a test for mortality effects from long-run exposure. We can infer the ARP had no observable impact on health if all estimates of β_t are zero. If estimates reveal an immediate decline in mortality, similar to the observed decline in SO2 emissions, this suggests a short-run effect where the level of pollution in year t only impacts mortality in year t. If estimates increase in years since the beginning of the ARP in 1995 (e.g., $0 = \beta_{1995} < \beta_{1996} < \ldots < \beta_{2005}$), this suggests immediate reductions in emissions led to a building reduction in mortality from cumulative, long-run exposure.[16]

In addition to county and year fixed effects, we take several steps to address possible confounding variables. The vector X_{ct} includes weather variables (to address the possibility of differential weather patterns in treatment counties) and the inverse hyperbolic sine of relevant denominator population.[17] We also control for the economic and demographic traits of each county flexibly by interacting 1990 measures (Z_c) with year fixed effects (illustrated here with the summation across Φ_t). Interacting with year effects allows one-time observed economic and demographic variables to have time varying effects. We favor this approach over using annual measures of control variables because annual measures could be affected themselves by treatment, though we explore including annual values as controls in robustness checks. Specific variables (from the 1990 County Data Book) used for Z_c include age distribution, percent male, racial and ethnic makeup, income per capita, unemployment, land area (to reflect potential urban vs. rural differences in mortality trends), and percentage employed in various economic sectors (manufacturing, retail, public administration, health services, federal government employment, and state government employment).[18] These economic and demographic factors interacted with year fixed effects control for differential changes across time. We also estimate a model with state-year fixed effects as part of our sensitivity analysis.[19]

[16] Another possible explanation for this effect would be a concave pollution damage function (Pope et al., 2015), where initial pollution drops do little to change mortality, but additional improvements at lower baseline levels have larger marginal effects on mortality.

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

[18] Age distribution controls include percentage in the following year groupings, as provided in the County Data Book data set: 5–17, 18–20, 21–25, 25–34, 35–44, 45–54, 55–65, 65–74, and 85 and up. Racial and ethnic controls include percentage black and percentage Hispanic.

[19] While geographic trend controls are possible, controlling for county-specific time trends potentially biases estimates in the presence of a dynamic treatment effect (Wolfers, 2006). If we were solely examining the contemporaneous response of mortality to the ARP, our model could credibly include county-specific time trends. Instead we focus on estimating a dynamic treatment effect in which the ARP causes mortality to evolve over time. County-specific trends absorb not only different preexisting trends across counties, but also differences in the evolution of mortality between treatment and control counties subsequent to the ARP. That is, county specific trends partly reflect the dynamic nature of the response variable to the policy shock, soaking up the variation in which we are particularly interested.

Table 1
Summary statistics.

	1985–1994	1995–2005
Panel A: Age-adjusted mortality rate (per 100,000 population)		
Cardiorespiratory	516.14	430.35
External	59.75	56.74
Panel B: Economic and demographic factors		
Population (1,000)	78.13	86.72
Employment (1,000)	35.36	41.26
Income per capita	24859.89	29335.70
Income maintenance programs per capita	441.22	529.84
Medicare per capita	0.82	1.27
Other public medical programs per capita	609.82	1067.17
Panel C: Pollution and upgrade exposure		
Monitor-measured SO2 (ppb)	7.53	4.65
Phase 1 upgrades: heat weighted	12.10	176.13

Notes: All values excepting SO2 are from a balanced panel of 2,414 counties used in our primary analysis. SO2 are from an unbalanced panel of monitors within 50 miles of 1,090 counties, based on availability of data. Mortality data are from the National Vital Statistics System restricted cause of death files. Population data are from the Surveillance, Epidemiology, and End Results (SEER) projections. Economic variables are from the Bureau of Economic Analysis (BEA) Regional Economic Accounts. We transform relevant values to per-capita using BEA estimates of population, and adjust to 2010 dollars.

3.4. Summary statistics

Table 1 shows pre- and post-ARP means for various measures of mortality, economic outcomes, and our measures of treatment. Panel A shows that age-adjusted mortality rates for cardiorespiratory and external mortality is decreasing across the period. Panel B shows population, wage employment, income per capita, and government payments per capita are all increasing. Finally, Panel C shows ambient SO2 decreasing, while the (non-standardized) measure of exposure to sulfur controls on Phase I boilers rises substantially after 1995.

4. Dynamic results

4.1. Impact on pollution

Evidence indicates the decrease in coal plant-generated SO2 after the ARP was immediate and persistent (Environmental Protection Agency, 2006). To test if this translated to observable changes in SO2 in surrounding areas, we use ambient SO2 levels as the dependent variables in our dynamic reduced form model. Panel A of Fig. 3 presents mean ambient SO2 data over time, with counties split by the median level of our treatment variable to facilitate comparisons. SO2 levels trend continuously downward over time in low treatment counties. SO2 levels in high treatment counties follow a similar trend but with a sharp drop immediately after the ARP began.

Panel B of Fig. 3 shows results from estimating Eq. (2). Although raw data show SO2 levels generally trending downward prior to the ARP, the event study shows little difference in trending by treatment intensity. Relative SO2 levels drop rapidly after the ARP for counties with greater exposure to sulfur controls. The marginal impact of a standard deviation in weighted sulfur controls is approximately 0.2 ppb in 1995 (statistically significant at 1 percent). The marginal effect of our treatment measure on the reduction in SO2 grows in the later years of the ARP, though it takes an additional 10 years to see further reductions similar to the initial one-year drop in 1995. Column 1 of Appendix Table A1 shows the estimates and standard errors for this figure.

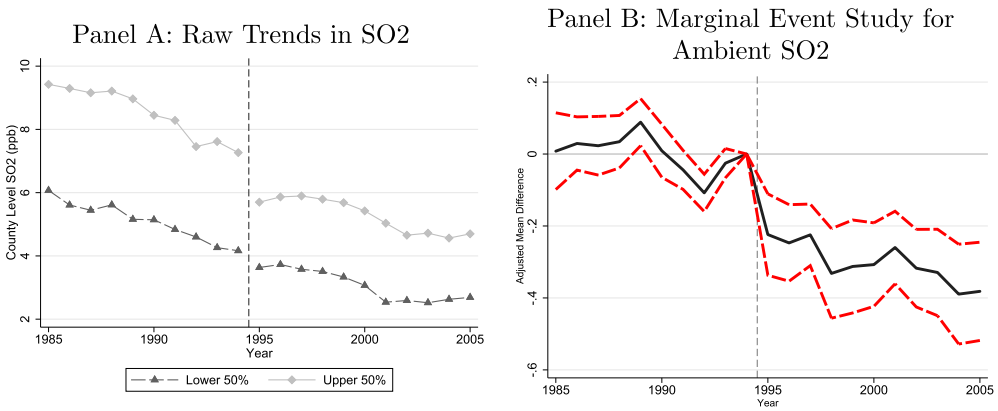

Fig. 3. Raw means by treatment intensity and regression-adjusted marginal event studies for SO2. Notes: Panel A is a raw trend in pollution collapsed by above and below the median of our measure of treatment, which we describe in Section 3.2. Panel B is an event study using a continuous measure of our treatment, as we describe in Eq. (2). Plotted values reflect the impact of a standard deviation in our normalized treatment, as compared to the base year of 1994. Dashed vertical line indicates the beginning of enforcement of the Acid Rain Program. Thick dashed lines indicate 95% confidence intervals.

We also present results for modeled PM2.5 using above and below split, shown in Appendix Fig. A2. Darker markers illustrate years in which true boiler emissions were unavailable and interpolated using data from surrounding years. Similar to the SO2 findings, we find an abrupt decrease in modeled PM2.5 after the ARP. The effect on modeled PM2.5 is largely flat during the 1995–2005 period. This immediate reduction in emissions, followed by smaller to no reductions in later years, underlies our identifying variation to separate the effects from short- and long-run pollution exposure.

4.2. Impact on mortality

Fig. 4 shows the raw plot for cardiorespiratory and external death rates by treatment intensity counties. Cardiorespiratory mortality rates decrease across time, with the gap between the two groups largely stable leading up to 1995. Following the ARP, mortality rates trend down relatively quicker in high treatment counties. External mortality rates also generally trend down but rise toward the end of the time period. Unlike cardiorespiratory mortality, external mortality moves in parallel fashion for both groups the entire time period. These patterns are consistent with the ARP having effects on mortality, but only for causes related to pollution exposure.

Fig. 4 also presents our main event study results for the two outcomes based on Eq. (2). As mortality outcomes are the inverse hyperbolic sine of deaths per 100,000 population, coefficients represent an approximate percentage change in mortality rates relative to 1994 from a standard deviation change in treatment exposure. There is no statistically significant difference in the years immediately preceding the ARP. After the ARP, estimates imply exposure to sulfur controls correlates with a gradual and steady decrease in mortality. While effects in 1995 are statistically indistinguishable from zero, by 2005, we estimate that cardiorespiratory mortality decreased by around 0.4 percent per standard deviation of treatment (statistically significant at 5%). Furthermore, the coefficient estimate in 2005 is significantly different from the estimate in 1995, suggesting mortality effects grew over time. Given the immediate drop in SO2 emissions after the ARP, the pattern of our mortality estimates support the notion of a long-term effect from cumulative exposure.

For external mortality, we find no statistically significant effect of treatment. Standard errors are sufficiently large such that we cannot fully reject the presence of effects comparable in size to the cardiorespiratory effect, but the absence of trends across the entire time period – in particular a trend that aligns with cardiorespiratory mortality – supports that external mortality is unrelated to the ARP. Appendix Fig. A3 shows our results for all-cause internal mortality – it follows a similar pattern to cardiorespiratory effects, though there are possible differential trends from 1985–1990.

Appendix Fig. A4 shows how different combinations of control variables affect our estimates. We find treatment correlates with similar decreases in mortality following the ARP, regardless of our choice of controls. With just basic year and county fixed effects, there are some suggestive differences in mortality between 1985 and 1990, though any such pre-trend disappears by 1991. Adding weather variables does little to change estimates. Once we control for population and year-effect interactions with demographic descriptors, pre-ARP estimates are near zero.

4.3. Effects by age

In Fig. 5, we present results by age to explore whether effects arise for more susceptible segments of the population or generalize more broadly across the population. We present event study results for all ages, under 35, 35–64, and 65 and older; columns 2–7 of Appendix Table A1 show the estimates and standard errors associated with these figures as well as those for all-cause internal mortality. Although estimates for the under-35 age group are noisy, there is no evidence to suggest an effect exists for this group. For both the 35–64 and over-65 group, we find a pattern of effects similar to our all age findings. For the 35–64 age group, a standard deviation increase in treatment reduced mortality by 1.6% by 2005, which is statistically different from both zero and the 1995 estimate at 1%. The effect in 2005 was smaller at 0.2% for the 65 and over group, though it reached as high as 0.4% in 2003. This smaller effect for the over-65 group comes from a much larger baseline mortality rate, and thus translates to a greater overall change in level of deaths for the oldest group. The suggestive evidence of a cumulative effect for those age 35–64 demonstrates potential for long-term exposure to pollution effects that extend into the prime working age population.[20]

4.4. Robustness checks

Appendix Fig. A6 shows results that estimate impacts on cardiorespiratory mortality for all ages using alternative specifications

[20] Appendix Fig. A5 replicates our falsification test using external mortality as the outcome but focuses on ages 35–64. We find little evidence of a change in trends or levels around the ARP, though estimates are sufficiently noisy.

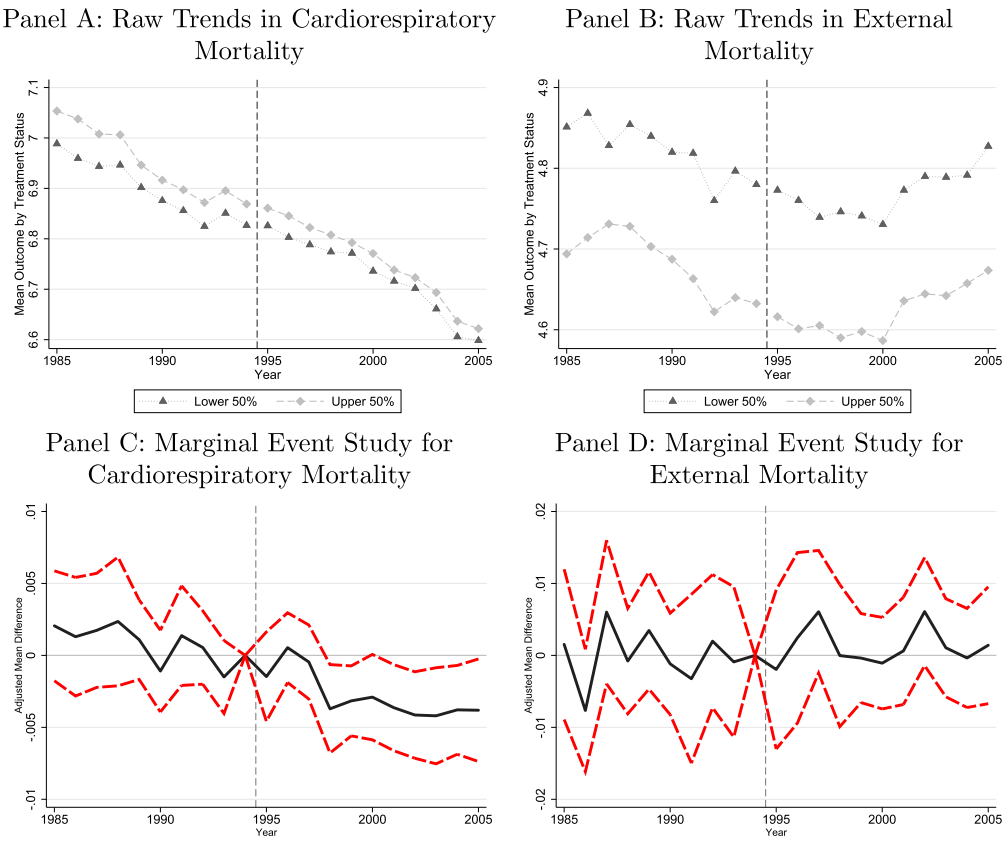

Fig. 4. Age-adjusted mortality for all ages: raw means by treatment intensity and regression-adjusted marginal event studies. Notes: Panels A and B are raw trends in the inverse hyperbolic sine of age-adjusted mortality rates, collapsed by above and below the median of our measure of treatment, which we describe in Section 3.2. Panels C and D are event studies using a continuous measure of our treatment, as we describe inEq. (2). Plotted values reflect the impact of a standard deviation in our normalized treatment, as compared to the base year of 1994. Dashed vertical line indicates the beginning of enforcement of the Acid Rain Program. Thick dashed lines indicate 95% confidence intervals.

to assess the robustness of our results. The figure shows 95 percent confidence intervals of our main estimates for context, along with event study results for 10 alternate models. We first assess sensitivity to our measures of population given that annual population at the county level is measured with error. In our first permutation, we replace the IHS of population with the CDB data on population in 1990 interacted with year fixed effects. In our second permutation, we omit our CDB age controls and control for the annual share in various age groups using the SEER population estimates (e.g., of the age 35–64, what percentage falls in the 35–44 category). In our third permutation, we use mortality rates that are not age-adjusted. Our fourth permutation uses annual population estimates as weights (rather than constant baseline population weights). In nearly all years and all models, our estimates lie within the 95% confidence intervals of our original estimate – the only exception is in the model without age adjustment where the early years' estimates are slightly higher.

We next explore the role of non-pollution changes which might also alter mortality. We do so in three ways. First, we re-run our analysis adding state-by-year fixed effects, which control for state-level policy changes that might drive health. Second, we omit from our analysis all counties that have a Phase I plant, which were the most heavily regulated under the policy. If economic shocks are localized around changes at the plant level, removing these counties from our sample removes local economic effects. Third, we add additional economic controls of income per capita, wage employment, Medicare per capita, income maintenance per capita, and additional public medical spending per capita, all from BEA regional economic accounts. In these cases, estimates remain within the confidence intervals of our initial estimates.

Given we use mortality rates per 100,000 population, and the wide range of county population sizes, we favor a model with weights to help with potential heteroskedasticity (Solon et al., 2015). To verify that results are not sensitive to our choice to use weights, we also replicate our main model without weights. As with our other robustness checks, the majority of the estimates fall within the 95% confidence interval of our main model, and the general pattern of effects remains the same. Our final test replaces the inverse hyperbolic sine of mortality rates with the natural log of mortality rates, with little change to our main estimates.

4.5. Assessing behavioral responses to the ARP

We further address the concerns that other factors might influence differential trends in health by treatment exposure by exploring impacts on economic outcomes. Appendix Fig. A7 shows four potentially indicative outcomes in our dynamic model: wage employment, income per capita, population, and a measure of migration using Internal Revenue Service (IRS) tax data. In each case, we repeat Eq. (2), with the small change that we omit our IHS of population on the right hand side (as population itself is one of our outcomes of interest). Panel A shows results for wage employment as measured in the regional economic accounts data from the BEA. While the occasional year is statistically different from 1994, there is no trend or pattern to the variation. Panel B shows income per capita, provided in the same data set. Again, there does not appear to be a distinct pattern either before or after the ARP. Panel C shows total population as provided in the SEER data. Despite statistical significance in occasional years, no discernable pattern arises.

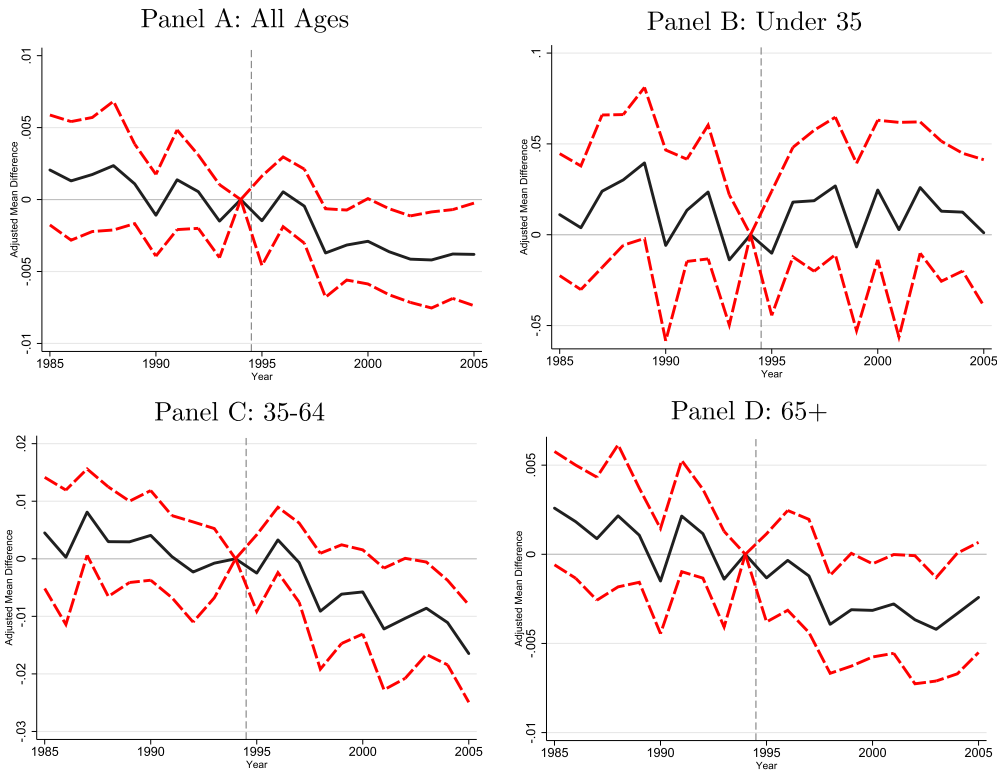

Fig. 5. Regression-adjusted marginal event studies for age-adjusted cardiorespiratory mortality by age group. Notes: All graphs are event studies using a continuous measure of our treatment, as we describe in Eq. (2). Plotted values reflect the impact of a standard deviation in our normalized treatment, as compared to the base year of 1994. Dashed vertical line indicates the beginning of enforcement of the Acid Rain Program. Thick dashed lines indicate 95% confidence intervals.

Finally, Panel D shows estimates for migration as the outcome variable. We label this "Net Income Exemptions" per 100,000 population because it is based on IRS estimates of county-to-county migration, measured by using the number of exemptions listed on tax forms. These data are only available from 1990 onward. To construct net exemptions, we subtract outgoing exemptions from incoming exemptions, thus measuring the total net flow of individuals – a positive number means a net gain in population via in migration. We find limited evidence of statistically significant effects, though we find a U-shaped pattern that bottoms out in 1994. These effects are economically small, implying a net inflow change of no more than 50 persons per 100,000 per standard deviation in treatment. Even under the strong assumption that none of these additional individuals died that year, this would decrease our treatment effect by 0.05%, which is one tenth the magnitude of our treatment effect of 0.5% in the later years.[21]

5. Static results

5.1. Static regression model

To quantify the per-unit effect of pollution, we take the following steps. First, we simplify our model as follows:

$$y_{ct} = \lambda_t^s + wSC_{ct} + \Sigma_t \Phi^s c * Z_c + \gamma^s X_{ct} + \alpha_c^s + \varepsilon_{ct}^s. \quad (3)$$

This is similar to Eq. (2), with the modification that ω is the *average* annual effect of our reduced-form emissions control installation on mortality, and SC_{ct} varies by year rather than being a constant interacted with year fixed effects. The superscript *s* indicates coefficients vary from the dynamic model (2). Second, in addition to using mortality as the dependent variable to obtain a reduced-form effect, we also estimate (3) using model-predicted PM2.5 as the dependent variable to obtain a "first stage" estimate. Third, we scale the reduced form coefficient by the first stage coefficient to obtain an estimate of the marginal effect of PM2.5 on mortality.

5.2. Mortality effects

While the dynamic effects of the ARP are the focus of our paper, we also estimate the effects for plant upgrades separately pooling pre- and post-treatment periods. Table 2 shows the OLS estimate of the reduced form effect. This is similar to our dynamic event study model, but assumes a common effect of treatment regardless of years of exposure and allows treatment intensity to adjust yearly. Here, cardiorespiratory mortality results are statistically significant for all age groups we consider. Column 1 shows the effect of a standard deviation in treatment is around 0.3% for all ages and 65 and older, and 0.8% for the 35–64 group. These estimates are smaller than later-year dynamic estimates (2000–2005), illustrating the critical nature of exploring our dynamic effects: a simple, static value across time fails to account for the accumulating health capital and thus potentially underestimates the true health gains of the program. Column 2 shows results for external mortality. As with our dynamic results, the effects on external mortality are economically small and statistically insignificant, though standard errors are large.

5.3. Estimating effects per unit of PM2.5

A lack of true data on PM2.5 until several years into the ARP means we cannot directly estimate a per-unit effect of PM2.5.

[21] Based on our pre-ARP mean of approximately 516 cardiorespiratory deaths per 100,000 population, the initial mortality rate would be 516/100,000. Adding an additional 50 people, none of whom die, would reduce the mortality rate to 516/100,050, a difference of 2.6*10e-6 percentage points or a relative decrease in mortality risk of 0.05%.

Table 2
Reduced form estimates: exposure to sulfur controls and age-adjusted mortality.

	1 Cardiorespiratory	2 External
Panel A: All Ages		
Weighted Plant Sulfur Controls	−0.003***	0.002
	(0.001)	(0.002)
Panel B: 35–64		
Weighted Plant Sulfur Controls	−0.008***	-0.001
	(0.003)	(0.004)
Panel C: 65+		
Weighted Plant Sulfur Controls	−0.003**	0.001
	(0.001)	(0.004)

Notes: Outcome is inverse hyperbolic sine (IHS) of internal age-adjusted mortality rate per 100,000 age-specific population. Regressions span 1985–2005 and include the IHS of age-specific population, county fixed effects, weather controls, and 1990 county demographic and economic characteristics interacted with year fixed effects (see Section 3.3). We cluster all standard errors at the county level and weight by county population in 1985. We derive the reduced form parameter estimates use APEEP2-weighted sulfur control technology upgrades at the power plant level (see Section 3), standardized to mean zero and standard deviation one — all coefficients show the impact of standard deviation increase in treatment.

Table 3
Estimate of change in cardiorespiratory deaths by year using 1994 baseline mortality rates and marginal annual reduced form values.

	1 Under 35	2 35–64	3 65+
Change in Deaths in 1995	−100	−732	−6546
Change in Deaths in 1996	177	966	−1689
Change in Deaths in 1997	184	−195	−6000
Change in Deaths in 1998	264	−2685	−19443
Change in Deaths in 1999	−66	−1810	−15388
Change in Deaths in 2000	242	−1697	−15569
Change in Deaths in 2001	27	−3592	−13779
Change in Deaths in 2002	255	−3048	−18167
Change in Deaths in 2003	127	−2532	−20829
Change in Deaths in 2004	122	−3266	−16425
Change in Deaths in 2005	11	−4852	−11970

Notes: We derive annual mortality reductions by taking 1994 mortality rates by county and age group, and multiplying by annual estimates of mortality reductions and sulfur upgrade exposure metrics (see Section 3) for each county. We then sum county-level predicted mortality changes to generate total expected changes for all counties in our sample. We perform estimates separately for the under 35, age 35–64, and age 65 and over groups.

Instead, using the data available from the APEEP transport model, we construct a measure of annual modeled PM2.5 arising from ARP-associated plants. This allows a rough instrumental variable approximation of the effect of PM2.5 on mortality by dividing our reduced form estimate by this "first stage" measure of how much ARP sulfur reductions affected modeled PM2.5.

Prior research and our own investigation shows the APEEP model accurately estimates shifts in regional PM2.5. For an in-depth analysis of the model's predictive ability of ambient emissions, see Tschofen et al. (2019) and Sergi et al. (2020). For our own exploration, we obtained PM2.5 data from EPA monitor readings, available from 1999 onward. We then aggregate daily monitor-level measures to annual county-level estimates. We plot the model-predicted levels of PM2.5, based on our sulfur output measures, against monitor-measured PM2.5 for all years in our data with overlap (1999–2005), in Appendix Fig. A8. Since we are most interested in the model's ability to predict changes in PM2.5, we also regressed monitor PM2.5 on our modeled PM2.5 and county fixed effects. The results show each additional unit of our modeled PM2.5 correlates with an additional 0.9 units of monitor-measured PM2.5. This suggests that, in our scenario, changes in predicted PM2.5 are a good approximation of ambient PM2.5 changes.

Based on a model similar to Eq. (3), we estimate the relationship between modeled PM2.5 and our measure of treatment. The relationship here is partly mechanical because we derive PM2.5 values using the same APEEP transport matrix we use to weight our measure of boiler sulfur controls. With that in mind, we find a standard deviation increase in sulfur controls lowers modeled PM2.5 by 0.72 micrograms per cubic meter, statistically significant at 1%. Dividing our 35–64 age group reduced form estimate by this gives an approximation of an IV result of 1.1% mortality reduction per unit of decreased baseline PM2.5.

5.4. Choice of treatment metric robustness checks

Appendix Table A2 shows how the estimates for all ages change under various different choices of treatment intensity. In each case, the treatment variable is some combination of boiler sulfur controls weighted by the dispersion matrix, normalized to mean zero and standard deviation one. Columns 1–4 gradually add weather, demographic, and economic controls, following the models in Appendix Fig. A3. Column 5 includes both Phase I and Phase II boiler sulfur controls. Columns 6 and 7 replace the heat-weighted boiler control installations with a boiler control installations weighted by 1985 sulfur emissions. Finally, Columns 7 and 8 weight boilers only by the dispersion matrix and ignore differences in plant size or initial traits. In each case, all of the predictive health effects in the reduced form are from Phase I sulfur controls, and correspond to a reduction of approximately 0.3% per standard deviation increase in controls. Appendix Fig. A9 shows how our event study results look under the alternate weighting structures, with little change in the overall pattern.

6. Benefit valuation

To gauge the magnitude of the estimated mortality effects, we calculate the avoided deaths per year in three different age groups: under 35, 35–64, and 65+. We take the 1994 age-appropriate mortality rates of each county in our sample, multiply by the annual reduced form estimate for each year, and then multiply by the county-specific value of the reduced form variable (in this case, our measure of SC_{ct}).[22] We then aggregate up to totals for all counties in our sample for a single age group/year value. Table 3 shows the results, indicating the number of avoided deaths increasing over time within each age group. The number of avoided deaths is generally higher in the 65+ age group despite percentage changes being smaller, which is consistent with the fact that baseline mortality rates increase with age. Based on our estimates, there was effectively no change in the under 35 mortality group. There were up to 5,000 fewer annual cardiorespiratory deaths for the 35–64 age group (highest in 2005), and up to 21,000 fewer annual cardiorespiratory deaths for the 65 and older age group (highest in 2003).

Differential numbers of avoided deaths by age also suggest valuing mortality improvements using a constant VSL is inappropriate; the loss in life expectancy is considerably larger for younger ages. We instead value mortality improvements using age-specific VSL estimates from Aldy and Smyth (2014) and compare them to valuations using a constant VSL.[23] Using age-specific VSLs, we calculate a mortality benefit that ranges from $51 billion to $109 billion in 2005 alone, depending on which age value we choose for each age range.[24] If instead we use a constant VSL of $4.9 million, we estimate

[22] For the purposes of these calculations, we use a non-standardized version of our treatment variable.

[23] We use the median VSL per life year based on their Fig. 3, and the implied constant VSL of $4.9million.

[24] For choosing a specific year for the age range, we use 2 different values for the 35–64 age group and the 65+ age group: 35 and 65 and 64 and 85. The range is largely driven by the choice for the older age group.

mortality benefits of $82 billion in 2005. Our finding of a large mortality effect for a younger age group suggests previous estimates of mortality benefits using constant VSLs were significantly undercounted.

7. Conclusion

The Acid Rain Program caused rapid and lasting improvements in ambient air quality, with health benefits that accumulated gradually over time. Using variation in installation of sulfur control technologies on regulated boilers, combined with a model of atmospheric transport and conversion between SO2 and PM2.5, we show cardiorespiratory mortality rates decreased in a manner consistent with long-run cumulative pollution exposure being an important health input. While emissions declined rapidly, mortality rates among those 35 and older decreased gradually. For the 35–64 age group, reductions in mortality risk from a standard deviation increase in sulfur control exposure grew from undetectable in the first year of the ARP to 1.6% 10 years later.

Our findings provide unique evidence that long-run exposure to pollution has detrimental effects on middle-aged individuals. These effects significantly increase the total mortality benefits from the ARP. We estimate the value of benefits in terms of annual lives saved as ranging from $51 billion to $109 billion by 2005, far exceeding estimated program costs of around $3 billion per year (Chestnut and Mills, 2005). The size of these mortality benefits makes the ARP stand out as one of the most cost-effective environmental regulations to date.

While we focus on the mortality effects from long-run exposure, there remain the full range of health and human capital outcomes potentially affected by the improvement in air quality. Quasi-experimental research shows effects from short-run exposure on educational outcomes (Sanders, 2012, Isen et al., 2017, Bharadwaj et al., 2017), worker productivity (Graff Zivin and Neidell, 2012; Chang et al., 2016, 2019), and morbidity (Schlenker and Walker, 2016). As we find mortality benefits from long-run exposure are quite different from short-run exposure, the same may be true for these other, important outcomes. Further understanding the effect from long-run exposure on these outcomes represents a fruitful area for future research.

Appendix

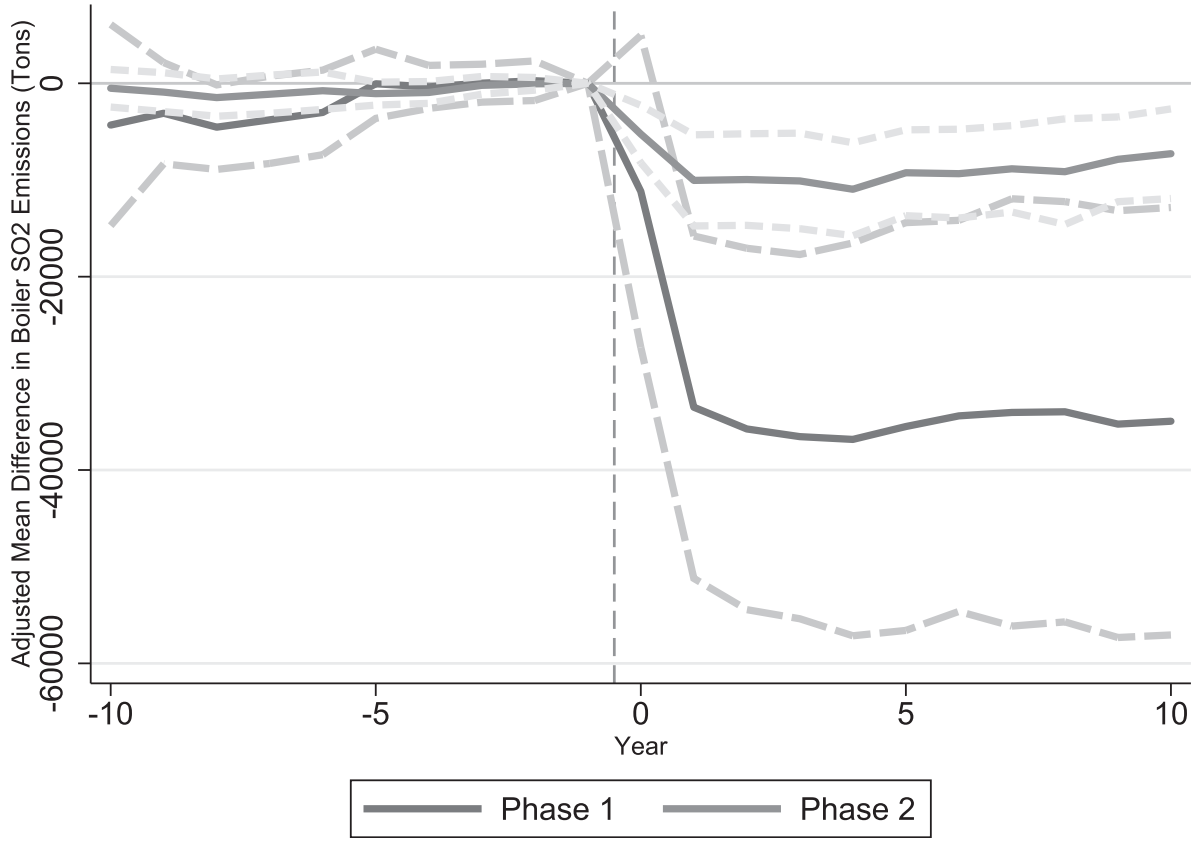

Fig. A1. Average impact of sulfur controls by acid rain program phases I and II. Notes: Figure based on data from the EPA Clean Air Markets Acid Rain power plant data set. Dashed vertical line indicates the installation of new sulfur controls at the boiler level. Thick dashed lines indicate 95% confidence intervals. See Section 3.2 for detailed discussion.

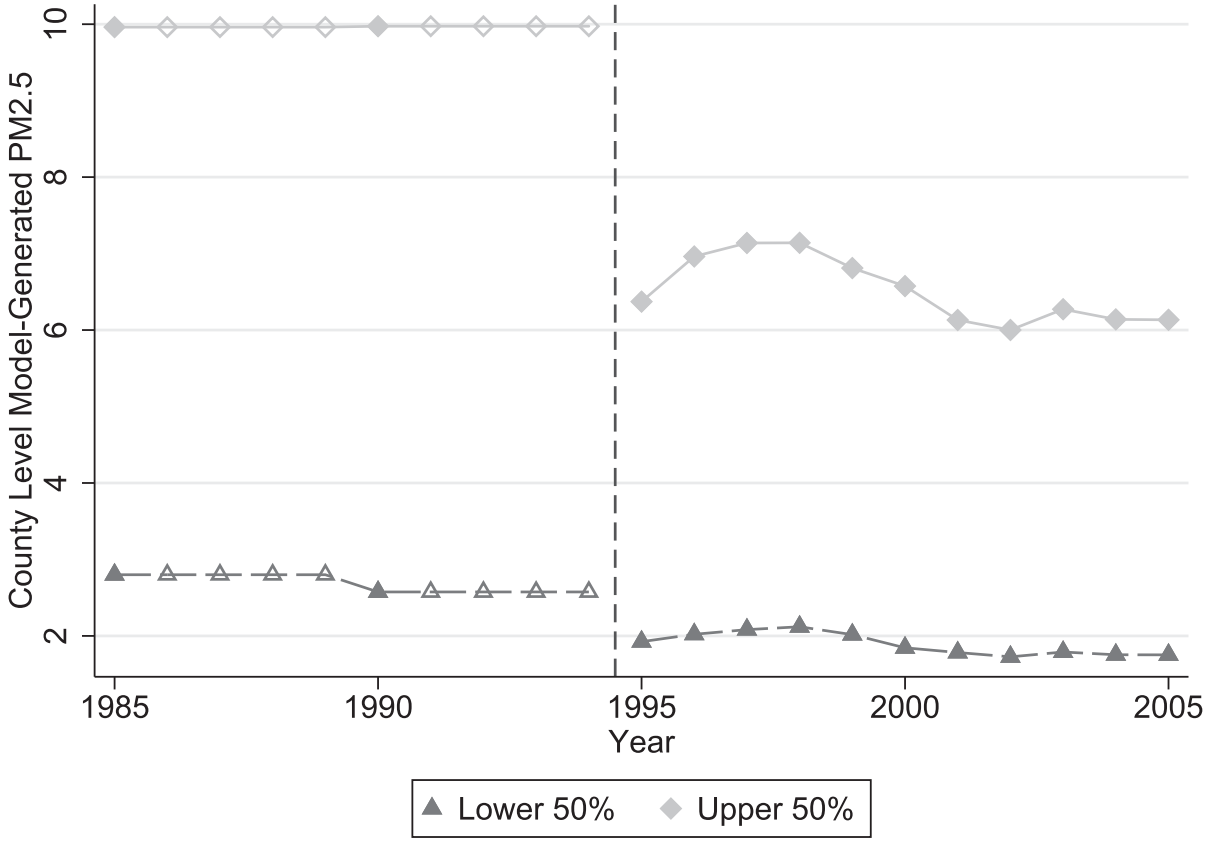

Fig. A2. Raw means by treatment intensity for model-estimated PM2.5. Notes: Figure shows a raw trend in model estimates of PM2.5 pollution, collapsed by above and below the median of our measure of treatment, sulfur control upgrades active in a given year, which we describe in Section 3.2. We generate PM2.5 values using the APEEP transport model and data from the EPA on plant-specific SO2 emissions. Data prior to 1995 are only available in 1985 and 1990. Hollow markers in 1986–1989 and 1991–1994 indicate imputation from most recent available power plant data | see Section 3.1 for details.

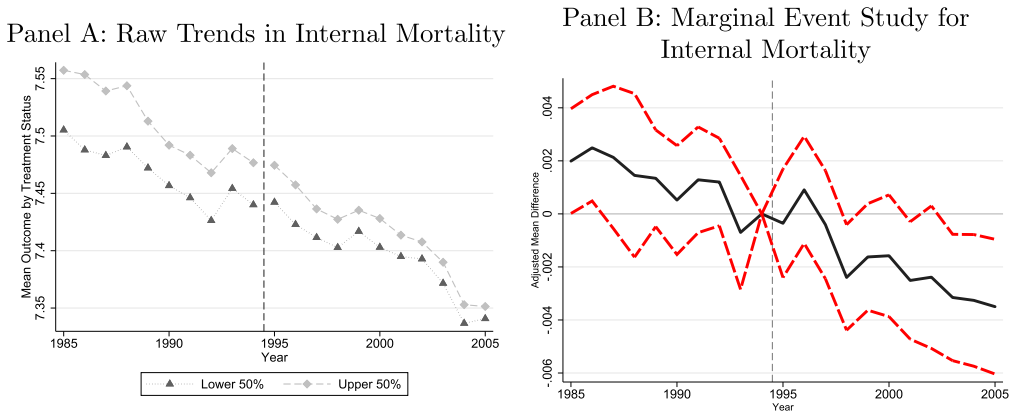

Fig. A3. Age-adjusted internal mortality for all ages: raw means by treatment intensity and regression-adjusted marginal event studies. Notes: Panel A is raw trends in the inverse hyperbolic sine of age-adjusted internal mortality rates, collapsed by above and below the median of our measure of treatment, which we describe in Section 3.2. Panel B is an event study using a continuous measure of our treatment, as we describe in equation 2. Plotted values reflect the impact of a standard deviation in our normalized treatment, as compared to the base year of 1994. Dashed vertical line indicates the beginning of enforcement of the Acid Rain Program. Thick dashed lines indicate 95% confidence intervals.

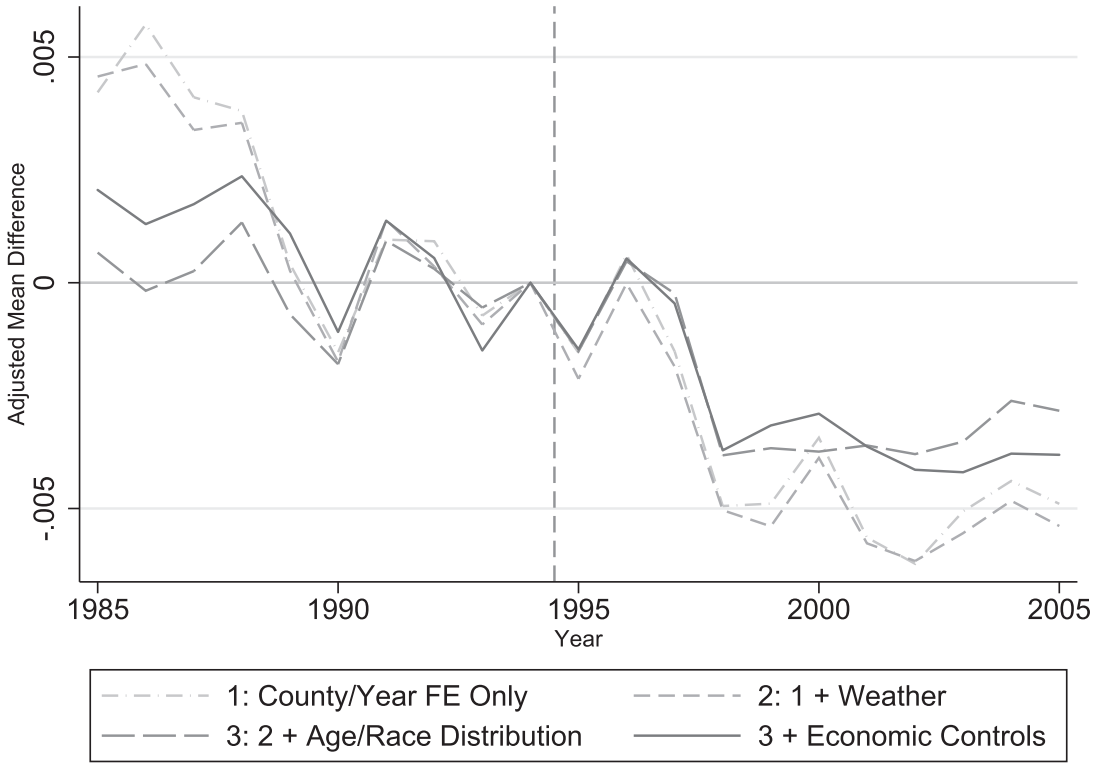

Fig. A4. Impact of adding covariates to regression-adjusted marginal event study of cardiorespiratory mortality for all ages. Notes: Figure shows the impact of adding various different sets of covariates to Eq. (2), as the legend describes –– see Section 4.2 for details.

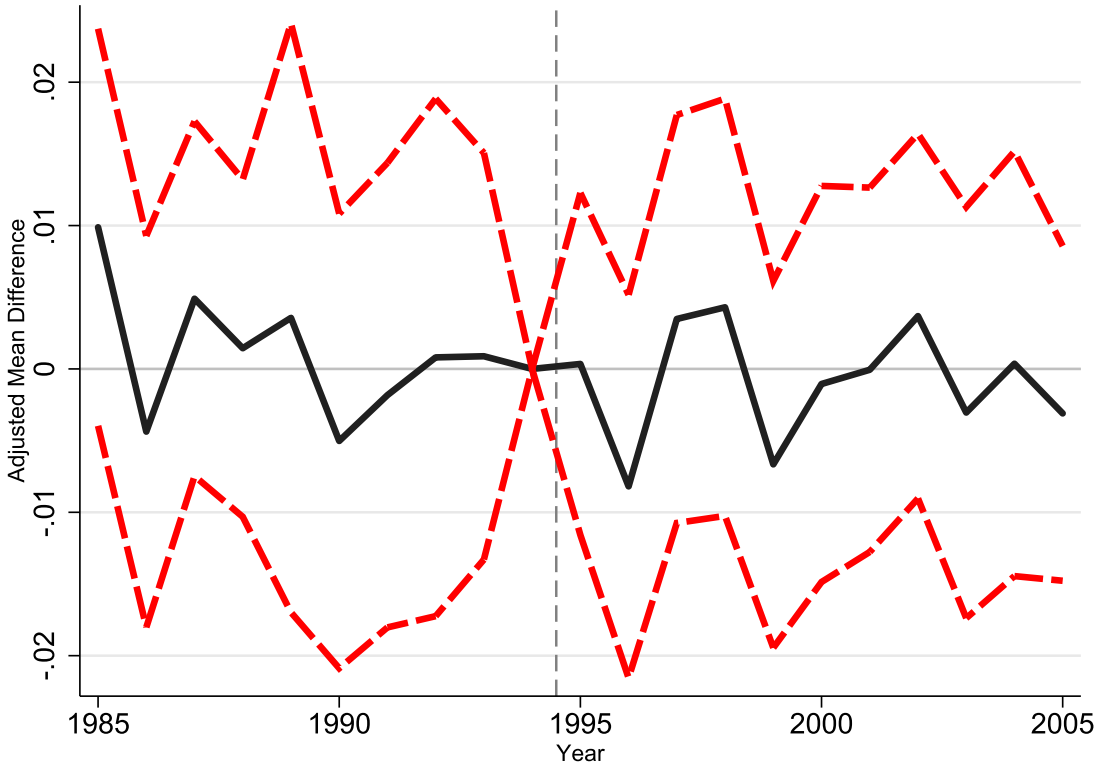

Fig. A5. Age-adjusted external mortality for 35–64: regression-adjusted marginal event study. Notes: Figure is an event study using a continuous measure of our treatment, as we describe in Eq. (2). Plotted values reflect the impact of a standard deviation in our normalized treatment, as compared to the base year of 1994. Dashed vertical line indicates the beginning of enforcement of the Acid Rain Program. Thick dashed lines indicate 95% confidence intervals.

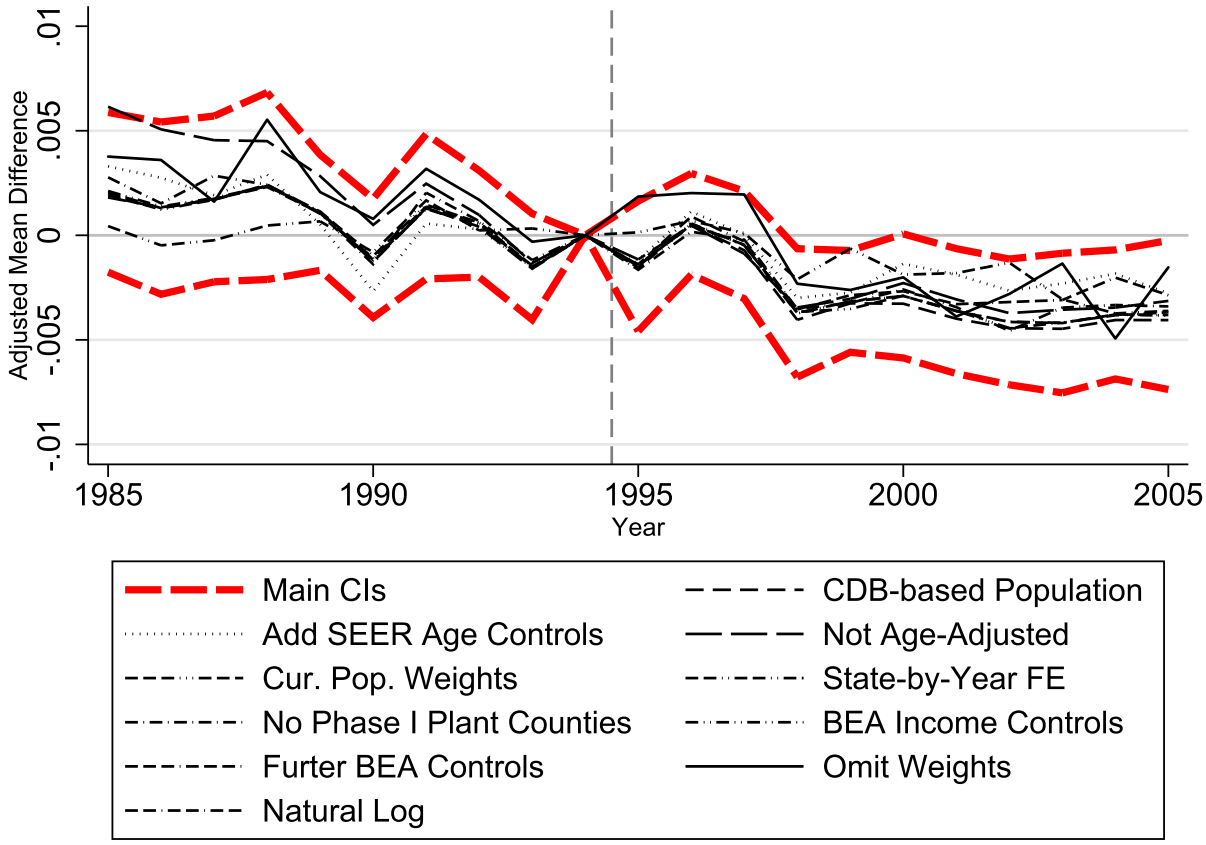

Fig. A6. Alternative specifications for regression-adjusted marginal event study of cardiorespiratory mortality for all ages. Notes: Dashed thick lines indicate 95% confidence intervals of our main estimates from Fig. 4C. Each additional line represents the coefficients from a model we estimate with the changes we note in the figure legend. We detail each model change in Section 4.4.

Fig. A7. Regression-adjusted marginal event studies: economic and population outcomes. Notes: Figures are event studies using a continuous measure of our treatment, as we describe in Eq. (2), but without controls for population. Plotted values reflect the impact of a standard deviation in our normalized treatment, as compared to the base year of 1994. Dashed vertical line indicates the beginning of enforcement of the Acid Rain Program. See Section 4.5 for details. Thick dashed lines indicate 95% confidence intervals.

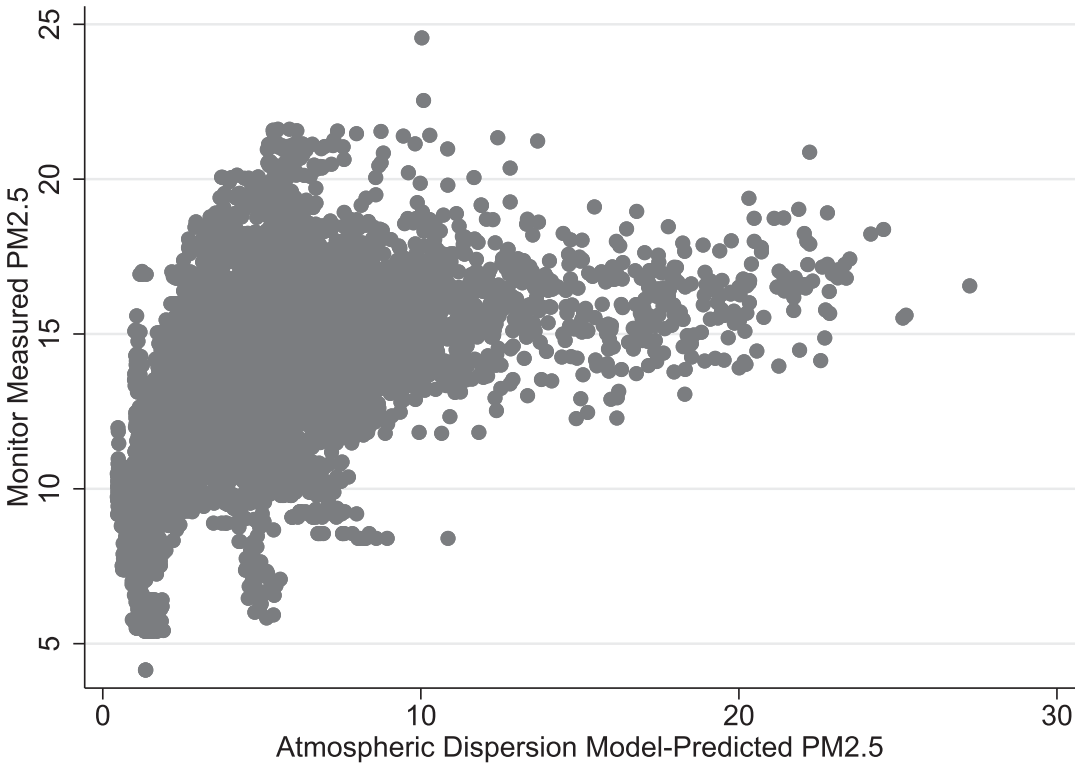

Fig. A8. Correlation between APEEP2 model-based PM2.5 and monitor-based PM2.5. Notes: Figure shows the relationship between our atmospheric transport model estimates of ARP-related PM2.5 generation and PM2.5 levels, aggregated to the county level, from monitors within 50 miles of a county centroid. Monitor data are not widely available prior to 1999. We use an unbalanced panel of monitors within 50 miles of 1,822 counties to maximize coverage.

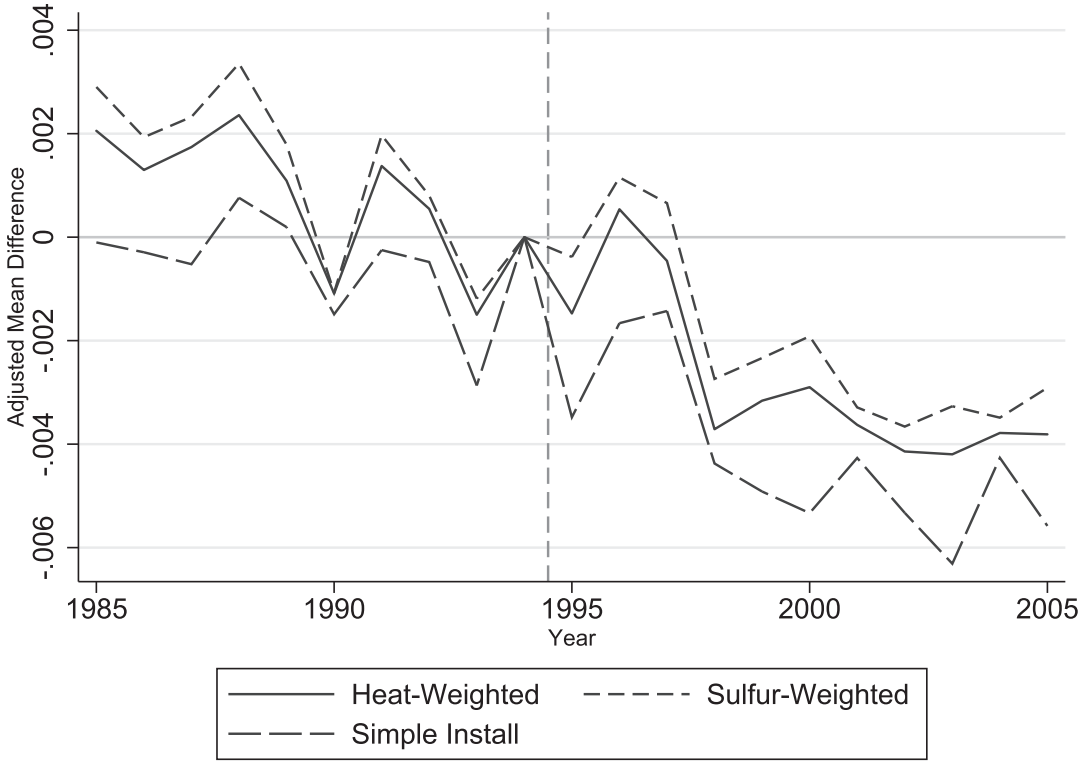

Fig. A9. Alternative weighting structures for measures of ARP exposure intensity. Notes: Figure shows different weighting structures for treatment intensity and impacts on all-age cardiorespiratory mortality. The first weights by boiler-specific heat output in 1985, our main specification corresponding to Panel C of Fig. 4. The second replaces heat output with sulfur emissions in 1985. The third treats all installs equally, with no additional weighting.

Table A1
Reduced form of exposure to upgraded ARP boilers.

	1	2	3	4	5	6	7
	Internal				Cardiorespiratory		
	SO2	All	35–64	65+	All	35–64	65+
Weighted Plant Sulfur Controls X 1985	0.008	0.002**	0.005**	0.001	0.002	0.004	0.003
	(0.054)	(0.001)	(0.002)	(0.001)	(0.002)	(0.005)	(0.002)
Weighted Plant Sulfur Controls X 1986	0.029	0.002**	0.004	0.002**	0.001	0.000	0.002
	(0.038)	(0.001)	(0.003)	(0.001)	(0.002)	(0.006)	(0.002)
Weighted Plant Sulfur Controls X 1987	0.023	0.002	0.006***	0.001	0.002	0.008**	0.001
	(0.042)	(0.001)	(0.002)	(0.001)	(0.002)	(0.004)	(0.002)
Weighted Plant Sulfur Controls X 1988	0.034	0.001	0.002	0.001	0.002	0.003	0.002
	(0.037)	(0.002)	(0.002)	(0.001)	(0.002)	(0.005)	(0.002)
Weighted Plant Sulfur Controls X 1989	0.089***	0.001	0.003*	0.001	0.001	0.003	0.001
	(0.033)	(0.001)	(0.002)	(0.001)	(0.001)	(0.004)	(0.001)
Weighted Plant Sulfur Controls X 1990	0.009	0.001	0.002	−0.000	−0.001	0.004	−0.002
	(0.038)	(0.001)	(0.002)	(0.001)	(0.001)	(0.004)	(0.002)
Weighted Plant Sulfur Controls X 1991	−0.044	0.001	0.000	0.002	0.001	0.000	0.002
	(0.028)	(0.001)	(0.002)	(0.001)	(0.002)	(0.004)	(0.002)
Weighted Plant Sulfur Controls X 1992	−0.108***	0.001	−0.003	0.002**	0.001	−0.002	0.001
	(0.026)	(0.001)	(0.002)	(0.001)	(0.001)	(0.004)	(0.001)
Weighted Plant Sulfur Controls X 1993	−0.025	−0.001	−0.001	−0.001	−0.001	−0.001	−0.001
	(0.021)	(0.001)	(0.002)	(0.001)	(0.001)	(0.003)	(0.001)
Omitted (1994)	-	-	-	-	-	-	-
	-	-	-	-	-	-	-
Weighted Plant Sulfur Controls X 1995	−0.224***	−0.000	−0.001	−0.000	−0.001	−0.002	−0.001
	(0.058)	(0.001)	(0.002)	(0.001)	(0.002)	(0.003)	(0.001)
Weighted Plant Sulfur Controls X 1996	−0.247***	0.001	−0.000	0.001	0.001	0.003	0.000
	(0.054)	(0.001)	(0.002)	(0.001)	(0.001)	(0.003)	(0.001)
Weighted Plant Sulfur Controls X 1997	−0.225***	−0.000	−0.003	−0.000	−0.000	−0.001	−0.001
	(0.044)	(0.001)	(0.002)	(0.001)	(0.001)	(0.004)	(0.002)
Weighted Plant Sulfur Controls X 1998	−0.332***	−0.002**	−0.006**	−0.002**	−0.004**	−0.009*	−0.004***
	(0.063)	(0.001)	(0.003)	(0.001)	(0.002)	(0.005)	(0.001)
Weighted Plant Sulfur Controls X 1999	−0.313***	−0.002	−0.005**	−0.001	−0.003**	−0.006	−0.003*
	(0.066)	(0.001)	(0.002)	(0.001)	(0.001)	(0.004)	(0.002)
Weighted Plant Sulfur Controls X 2000	−0.307***	−0.002	−0.002	−0.002	−0.003*	−0.006	−0.003**
	(0.059)	(0.001)	(0.002)	(0.001)	(0.002)	(0.004)	(0.001)
Weighted Plant Sulfur Controls X 2001	−0.260***	−0.003**	−0.007***	−0.001	−0.004**	−0.012**	−0.003**
	(0.051)	(0.001)	(0.002)	(0.001)	(0.002)	(0.005)	(0.001)
Weighted Plant Sulfur Controls X 2002	−0.317***	−0.002*	−0.004**	−0.002	−0.004***	−0.010*	−0.004**
	(0.055)	(0.001)	(0.002)	(0.001)	(0.002)	(0.005)	(0.002)
Weighted Plant Sulfur Controls X 2003	−0.329***	−0.003***	−0.007***	−0.003**	−0.004**	−0.009**	−0.004***
	(0.061)	(0.001)	(0.002)	(0.001)	(0.002)	(0.004)	(0.001)
Weighted Plant Sulfur Controls X 2004	−0.389***	−0.003***	−0.008***	−0.002*	−0.004**	−0.011***	−0.003*
	(0.071)	(0.001)	(0.002)	(0.001)	(0.002)	(0.004)	(0.002)
Weighted Plant Sulfur Controls X 2005	−0.382***	−0.003***	−0.010***	−0.002	−0.004**	−0.016***	−0.002
	(0.070)	(0.001)	(0.003)	(0.001)	(0.002)	(0.004)	(0.002)
p-value of 1995 = 2005	0.000	0.005	0.000	0.174	0.100	0.000	0.423
Counties	1.090	2.414	2.414	2.414	2.414	2.414	2,414
Observations	22.890	50.694	50.694	50.694	50.694	50.694	50,694

Notes: Table shows coefficients corresponding to graphs from Fig. 5, as well as all internal mortality by various age groups. p-value is for the test of equality between the coefficient for 1995 and 2005.

Table A2
Static reduced form of all-age cardiorespiratory mortality with alternate specifications and standardized treatment measures.

	1	2	3	4	5	6	7	8	9
Heat-Weighted Update (Phase 1)	−0.006***	−0.006***	−0.002*	−0.003***	−0.004***				
	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)				
Heat-Weighted Update (Phase 2)					0.001				
					(0.001)				
Sulfur-Weighted Update (Phase 1)						−0.003***	−0.003**		
						(0.001)	(0.001)		
Sulfur-Weighted Update (Phase 2)							−0.001		
							(0.000)		
Weighted Update (Phase 1)								−0.003**	−0.003**
								(0.001)	(0.001)
Weighted Update (Phase 2)									0.001
									(0.001)
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Weather	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age Distribution Controls	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Economic Controls	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	50,694	50,694	50,694	50,694	50,694	50,694	50,694	50,694	50,694
Counties	2,414	2,414	2,414	2,414	2,414	2,414	2,414	2,414	2,414

Notes: Each column corresponds to a different regression. Columns 1 through 4 use our main treatment with different choices of covariates, corresponding to the various lines in Appendix Fig. A4. Column 5 adds weighted upgrades to Phase II power plants. Columns 6 and 7 replicate Columns 4 and 5 but replacing heat input with 1985 boiler-specific SO2 output. Columns 8 and 9 replicate Columns 4 and 5 but treating all plant controls the same.

References

Aldy, Joseph, Smyth, Seamus, 2014. Heterogeneity in the Value of Life. NBER Working Paper 20206.

Almond, Douglas, Currie, Janet, 2011. Killing me softly: the fetal origins hypothesis. J. Econ. Perspect. 25, 153–172.

Anderson, Michael, 2019. As the wind blows: the effects of long-term exposure to air pollution on mortality. J. Eur. Econ. Assoc. 18 (4), 1886–1927.

Angrist, Joshua, Pischke, Jorn-Steffen, 2009. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press, Princeton, NJ.

Banzhaf, H. Spencer, Walsh, Randall P., 2008. Do people vote with their feet? An empirical test of tiebout's mechanism. Am. Econ. Rev. 98, 843–863.

Bharadwaj, Prashant, Zivin, Joshua Graff, Gibson, Matthew, Neilson, Christopher A., 2017. Gray matters: Fetal pollution exposure and human capital formation. J. Assoc. Environ. Resour. Econ. 4 (2), 505–542.

Brook, R.D., Rajagopalan, S., Pope, C.A. 3rd, Brook, J.R., Bhatnagar, A., Diez-Roux, A.V., Holguin, F., Hong, Y., Luepker, R.V., Mittleman, M.A., Peters, A., Siscovick, D., Smith, S.C. Jr, Whitsel, L., Kaufman, J.D., on behalf of the American Heart Association Council on Epidemiology and Prevention, Council on the Kidney in Cardiovascular Disease, and Council on Nutrition, Physical Activity and Metabolism, 2010. Particulate matter air pollution and cardiovascular disease: an update to the scientific statement from the American Heart Association. Circulation,121:2331–2378.

Card, David, DellaVigna, Stefano, 2017. "What Do Editors Maximize? Evidence from Four Leading Economics Journals". NBER Working Paper 23282.

Case, Anne, Deaton, Angus, 2015. Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century. Proc. Natl. Acad. Sci. 112 (49), 15078–15083.

Chang, Tom, Zivin, Joshua Graff, Gross, Tal, Neidell, Matthew, 2016. Particulate pollution and the productivity of pear packers. Am. Econ. J.: Econ. Policy 8 (3), 141–169.

Chang, Tom, Zivin, Joshua Graff, Gross, Tal, Neidell, Matthew, 2019. The effect of pollution on worker productivity: evidence from call-center workers in China. Am. Econ. J.: Appl. Econ. 11 (1), 151–172.

Chen, Yuyu, Ebenstein, Avraham, Greenstone, Michael, Li, Hongbin, 2013. Evidence on the impact of sustained exposure to air pollution on life expectancy from China's Huai River policy. Proc. Natl. Acad. Sci., 110(32): 12936-12941

Chestnut, Lauraine, Mills, David, 2005. A fresh look at the benefits and costs of the US acid rain program. J. Environ. Manage. 77, 25–266.

Currie, Janet, Davis, Lucas, Greenstone, Michael, Walker, Reed, 2015. Environmental health risks and housing values: evidence from 1600 toxic plant openings and closings. Am. Econ. Rev., 105(2): 678-709.

Deryugina, Tatyana, Heutel, Garth, Miller, Nolan H., Molitor, David, Reif, Julian, 2019. The mortality and medical costs of air pollution: Evidence from changes in wind direction. Am. Econ. Rev. 109 (12), 4178–4219.

Deschenes, Olivier, Greenstone, Michael, Shapiro, Joseph, 2017. Defensive investments and the demand for air quality: evidence from the NOx budget program. Am. Econ. Rev. 107 (10), 2958–2989.

Di, Qian, Wang, Yan, Zanobetti, Antonella, Wang, Yun, Koutrakis, Petros, Choirat, Christine, Dominici, Francesca, Schwartz, Joel D., 2017. Air pollution and mortality in the Medicare population. N. Engl. J. Med. 376 (26), 2513–2522.

Dominici, Francesca, Greenstone, Michael, Sunstein, Cass R., 2014. Particulate matter matters. Science 344 (6181), 257–259.

Dockery, Douglas W., Arden Pope, C., Xu, Xiping, Spengler, John D., Ware, James H., Fay, Martha E., Ferris, Jr., Benjamin G., Speizer, Franke E., 1993. An association between air pollution and mortality in six U.S. cities. New Engl. J. Med., 329(24): 1753-1759.

Ebenstein, Avraham, Fan, Maoyong, Greenstone, Michael, He, Guojun, Zhou, Maigeng, 2017. New evidence on the impact of sustained exposure to air pollution on life expectancy from China's Huai River Policy. Proc. Natl. Acad. Sci. 114, no. 39 (2017): 10384-10389

Environmental Protection Agency, 2004. Air Quality Criteria for Particulate Matter. National Center for Environmental Assessment, Research Triangle Park, NC.

Environmental Protection Agency, 2006. Acid rain program: 2005 progress report. EPA-430-R-06-015, Clean Air Markets Division, Office of Air and Radiation.

Environmental Protection Agency, 2009. Risk and Exposure Assessment to Support the Review of the SO2 Primary National Ambient Air Quality Standards: Final Report. Research Triangle Park, NC.

Graff Zivin, Joshua, Neidell, Matthew, 2012. The impact of pollution on worker productivity. Am. Econ. Rev., 102(7): 3652-3673.

Graff Zivin, Joshua, Neidell, Matthew, 2013. Environment, health, and human capital. J. Econ. Literature, 51(3): 689-730.

Greenstone, Michael, 2002. The impacts of environmental regulations on industrial activity: evidence from the 1970 and 1977 clean air act amendments and the census of manufactures. J. Political Econ. 110 (6), 1175–1219.

Hodan, William M., Barnard, William R., 2004. Evaluating the Contribution of PM2. 5 Precursor Gases and Re-entrained Road Emissions to Mobile Source PM2.5 Particulate Matter Emissions. MACTEC Federal Programs, Research Triangle Park, NC.

Isen, Adam, Rossin-Slater, Maya, Walker, Reed, 2017. Every breath you take – every dollar you'll make: The long-term consequences of the clean air act of 1970. J. Political Econ. 125 (3), 848–902.

Lippmann, Morton, Thurston, George D., 1996. Sulfate concentrations as an indicator of ambient particulate matter air pollution for health risk evaluations. J. Expo. Anal. Environ. Epidemiol. 6 (2), 123–146.

Molloy, Raven, Smith, Christopher L., Wozniak, Abigail, 2011. Internal migration in the United States. J. Econ. Perspect. 25 (3), 173–196.

Muller, N.Z., 2014. Boosting GDP growth by accounting for the environment: including air pollution and greenhouse gas damages increases estimated U.S. growth. Science 345 (6199), 873–874.

Murphy, Kevin, Topel, Robert, 2006. The value of health and longevity. J. Political Econ. 114 (5), 871–904.

Ozkaynak, H., Xue, J., Spengler, J., Wallace, L., Pellizzari, E., Jenkins, P., 1996. Personal exposure to airborne particles and metals: results from the Particle TEAM study in Riverside, California. J. Expo. Anal. Environ. Epidemiol. 6 (1), 57–78.

Pope III, C. Arden, Cropper, Maureen, Coggins, Jay, Cohen, Aaron, 2015. Health benefits of air pollution abatement policy: Role of the shape of the concentration-response function. J. Air & Waste Manage. Assoc., 65(5), 516-522.

Sanders, Nicholas J., 2012. What doesn't kill you makes you weaker: prenatal pollution exposure and educational outcomes. J. Hum. Resour. 47 (2), 826–850.

Schlenker, Wolfram, Walker, Reed, 2016. Airports, air pollution, and contemporaneous health. Rev. Econ. Stud. 83 (2), 768–809.

Schmalensee, Richard, Stavins, Robert, 2017. Lessons learned from three decades of experience with cap-and-trade. Rev. Environ. Econ. Policy 11 (1), 59–79.
Stavins, Robert N., 1998. "What can we learn from the grand policy experiment? Lessons from SO2 allowance trading. J. Econ. Perspect., 12(3): 69–88.
Sergi, Brian, Azevedo, Inês, Davis, Steven J., Muller, Nicholas Z., 2020. Regional and county flows of particulate matter damage in the US. Environ. Res. Lett. 15, (10) 104073.
Solon, Gary, Haider, Steven J., Wooldridge, Jeffrey M., 2015. What are we weighting for?. J. Hum. Resour. 50 (3), 301–316.
Sullivan, Daniel, von Wachter, Till, 2009. Job displacement and mortality: An analysis using administrative data. Quart. J. Econ. 124 (3), 1265–1306.
Thatcher, T.L., Layton, D.W., 1995. Deposition, resuspension, and penetration of particles within a residence. Atmos. Environ. 29 (13), 1487–1497.
Tschofen, Peter, Azevedo, Inês L., Muller, Nicholas Z., 2019. Fine particulate matter damages and value added in the US economy. Proc. Natl. Acad. Sci. 116 (40), 19857–19862.
Vette, A.F., Rea, A.W., Lawless, P.A., Rodes, C.E., Evans, G., Highsmith, V.R., Sheldon, L., 2001. Characterization of indoor-outdoor aerosol concentration relationships during the Fresno PM exposure studies. Aerosol Sci. Technol. 34 (1), 118–126.
Walker, Reed, 2013. The transitional costs of sectoral reallocation: evidence from the clean air act and the workforce. Quart. J. Econ. 128 (4), 1787–1835.
Wolfers, Justin, 2006. Did unilateral divorce raise divorce rates? A reconciliation and new results. Am. Econ. Rev. 96 (5), 1802–1820.