

NBER WORKING PAPER SERIES

LONG-RUN POLLUTION EXPOSURE AND ADULT MORTALITY:
EVIDENCE FROM THE ACID RAIN PROGRAM

Alan I. Barreca
Matthew Neidell
Nicholas J. Sanders

Working Paper 23524
<http://www.nber.org/papers/w23524>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2017

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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Alan I. Barreca, Matthew Neidell, and Nicholas J. Sanders. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Long-Run Pollution Exposure and Adult Mortality: Evidence from the Acid Rain Program
Alan I. Barreca, Matthew Neidell, and Nicholas J. Sanders
NBER Working Paper No. 23524
June 2017
JEL No. I10,Q51,Q53

ABSTRACT

Though over 90 percent of benefits from environmental quality improvements are attributed to long-term exposure, nearly all quasi-experimental evidence on the effects of pollution on health exploits changes in short-term exposure. Quantifying long-run exposure impacts requires a lasting, exogenous change in ambient pollution. Even if the initial change in pollution is exogenous, the long-run nature allows more time for economic agents to respond to changes in pollution, resulting in endogenous pollution exposure. We estimate the effects of long-run pollution exposure on mortality among adults by exploiting the United States Acid Rain Program (ARP) as a natural experiment. The ARP, which regulated emissions from coal power plants, created a permanent change in pollution across vast distances, enabling us to define broad treatment areas to subsume many potential confounding effects. We use a difference-in-differences design, comparing changes in mortality over time in counties “near” regulated plants to changes in mortality in similar counties “far” from the plants. We find relative mortality in treatment counties decreased after the introduction of the ARP, with mortality improvements growing steadily over time in both economic and statistical significance. The ARP had no significant effect on residential sorting or employment, helping rule out selection or economic mechanisms. Analysis by cause of death supports the role of fine particulate matter as the relevant pollutant.

Alan I. Barreca
Department of Economics
Tulane University
206 Tilton Hall
New Orleans, LA 70118
and NBER
abarreca@tulane.edu

Nicholas J. Sanders
Department of Policy Analysis and Management
Cornell University
MVR Hall, Room 105
Ithaca, NY 14853
and NBER
njsanders@cornell.edu

Matthew Neidell
Department of Health Policy and Management
Columbia University
722 W 168th Street, 6th Floor
New York, NY 10032
and NBER
mn2191@columbia.edu

1. Introduction

Though over 90 percent of benefits from environmental quality improvements are attributed to long-term exposure (Dominici et al., 2014), nearly all quasi-experimental research on the effects of pollution exposure on health exploits changes in short-run exposure.¹ Quantifying the effects of long-run exposure presents unique challenges because it requires a long-lasting, exogenous change in ambient pollution. Temporary, unexpected changes in pollution, while powerful in short-run estimates, can affect timing of exposure without substantively affecting long-run exposure (Graff Zivin and Neidell 2012; Deschenes et al. 2012; Schlenker and Walker 2016; Deryugina et al. 2016). Furthermore, economic agents may respond to the changes in pollution, especially in the long run, resulting in changes in labor and housing markets, such as job loss and residential sorting (Walker, 2013; Banzhaf and Walsh, 2008). Even if an initial change in pollution is exogenous, pollution exposure is thus likely to be endogenous in the long-run. This lack of quasi-experimental evidence on the health effects from accumulated, long-run pollution exposure leaves an important gap in the literature.

In this study, we estimate the effects of long-run pollution exposure on mortality among adults using the United States Acid Rain Program (ARP), a cap-and-trade program to control sulfur dioxide (SO₂) emissions, as a natural experiment. Four factors make the ARP an ideal setting for our study. First, the ARP created a permanent change in pollution, necessary for identifying effects of long-run exposure. Once implemented, the ARP improved SO₂ almost immediately, and improvements persist to date. SO₂ is a precursor gas in the formation of particulate matter smaller than 2.5 micrometers (PM_{2.5}),² which has much more detrimental

¹ Two important exceptions are Chen et al. (2013) and Anderson (2015). We discuss these in detail below.

² See Hodan and Barnard (2004).

effects on human health than SO₂ itself (see Environmental Protection Agency (2004) for a comprehensive review).

Second, the ARP initially regulated only the 110 most SO₂-intensive coal plants, leaving some areas affected by the regulation and others not. This lends to a difference-in-differences design comparing changes in mortality over time in counties “near” regulated plants to changes in mortality in counties “far” from the plants.³ This mitigates bias from national trends that change over time independent of proximity to coal-fired power plants, such as business cycles, health care access, and technological advances.

Third, given the vast distance both SO₂ and PM_{2.5} can travel once airborne, we define broad treatment areas up to 100 miles from the plant, subsuming many potential confounding general equilibrium effects.⁴ For example, households sorting in response to the economic effects of the ARP will not bias estimates if such sorting remains within a 100-mile radius of regulated power plants. Given existing evidence on the opening and closing of power plants finds changes in housing amenity values occur at distances of less than 2 miles, a 100-mile buffer seems adequate for capturing sorting.⁵ Even so, we directly test for sorting using various population characteristics, and find no evidence of sorting in response to the ARP.⁶

³ Regulatory status, announced in 1990 and imposed in 1995, used emissions levels from 1985 to avoid concerns of regulatory gaming.

⁴ We use flexible estimated effects at different distances to choose this treatment cutoff.

⁵ 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.

⁶ 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}.

Fourth, the potential impact of the environmental regulation on economic conditions is likely small in this context. Workers in regulated areas often experience 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, 2015).⁷ Moreover, our broad treatment areas subsume local economic effects, to the extent they exist, in the same way they deal with sorting. Indeed, we find the ARP had negligible effects on several labor market outcomes, suggesting that employment effects are unlikely to bias estimates of health effects. As a further test, since our broad treatment area includes multiple counties, we estimate models omitting counties containing regulated plants.

Our research design compares across-time differences in adult mortality for counties within 100 miles of regulated power plants (our treatment group) to counties beyond 100 miles (our control group). Our core population of interest is adults ages 35-64, though we explore impacts on other age groups as well. We begin analysis in 1990, when the government announced the ARP; define post-treatment as beginning in 1995, the year the ARP began; and end in 2005, before changes to the SO₂ permit market potentially shifted the relationship between regulation and ambient emissions. This allows us to test for pre-trends in mortality differences for the 5 years leading up to the ARP, and for changing mortality up to 10 years after the policy. Given regulated coal plants are largely located in the Midwest and Northeast, we restrict our sample to counties with similar geographic, economic, and demographic observables at baseline using a propensity-score matching algorithm. We also control for propensity-score based flexible time effects to better address differential trends that might correlate with treatment.

⁷ Regulatory costs were lower than predicted due to rapid technological innovation in desulphurization and an unexpected increase in access to low-sulfur coal.

Overall, we find evidence consistent with lasting reduced mortality effects from decreases in extended, long-run pollution exposure. We show that SO₂ emissions in treatment counties drop immediately after the ARP begins, and the reductions persist in the following years. We then show mortality in treatment counties decreased (relative to control counties) after the introduction of the ARP, with relative mortality improvements growing steadily over time in both economic and statistical significance. By 2005, mortality rates among the age 35-64 population are 5 percent lower in treated counties. Though the policy regulated SO₂, we contend health effects operate through PM_{2.5} rather than SO₂. Data limitations on PM_{2.5} for the years surrounding the ARP preclude a direct test of this hypothesis. Nonetheless, analyses by cause of death, showing large and statistically significant effects on cardiovascular deaths, suggest PM_{2.5} is the responsible pollutant. As a test of endogenous sorting and changing economic conditions, we show no discernible differences between treatment and control counties in employment, personal income, migration, and population composition either before or after the ARP.

Our study offers important innovations beyond recent quasi-experimental studies addressing potential long-run exposure (Chen et al., 2013; Anderson, 2015). First, prior research relies on cross-sectional designs, which may conflate the effects from short-run and long-run exposure. We provide a year-by-year event study of changes in adult mortality over time, allowing us to directly disentangle the short-run exposure effects from long-run exposure effects. Second, we focus on working age adults, a group considered a more health-robust segment of the population. The working age range is one where mortality represents a significant loss in life expectancy and economic productivity, so existing estimates of the value of statistical life (VSL) more readily apply.⁸

⁸ Existing studies typically focus on compensating differentials for workers across industries. See Murphy and Topel (2006) for a discussion of VSL over the lifecycle.

Beyond helping to better understand the relationship between pollution and adult health, our reduced form estimates are of direct interest to policies centered on emissions from coal power plants. Global coal consumption has nearly doubled since the turn of the century with rapid growth in developing nations such as China and India⁹. In the United States, the Clean Power Plan focuses heavily on reducing emissions from coal power plants, with potentially large regulatory costs. To properly assess the economic consequences from such changes, our estimates provide a benchmark for evaluating potential mortality benefits.

2. Background

2.1 The Acid Rain Program

The ARP was a provision of the Environmental Protection Agency (EPA) Clean Air Act Amendments of 1990.¹⁰ The two-phase program regulates SO₂-producing coal power plants, with the goal of reducing US 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 the Phase II regulations by banking permits in advance. As such, our quasi-experimental research design focuses on Phase I.

⁹ "China and India drive recent changes in world coal trade," Today in Energy, U.S. Energy Information Administration, Nov. 20, 2015.

¹⁰ The predecessor to the ARP was 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).

The EPA distributed SO₂ allowances to 263 heating units at the 110 plants based on baseline (1985-1987) heat input (Stavins, 1998), calculated in British thermal units (BTUs).¹¹ 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 program compliance, leading to substantial decline in SO₂ emissions from regulated plants.

2.2 Health effects of SO₂ and PM_{2.5}

Though the primary aim of the ARP was SO₂'s role in formation of acid rain, which has no known direct impacts of human health, there are potential health gains from the regulation due to improvement in associated pollutants, including direct SO₂ exposure.¹² Effects are primarily short-run and respiratory-related, leading to increases in hospital admissions for outcomes such as asthma exacerbations.¹³

Much of the anticipated lasting health effects from the ARP arise from other pollutants correlated with SO₂ levels, namely fine particulate matter (PM_{2.5}). Through an atmospheric conversion process, SO₂ produces sulfate particles, a portion of small-scale PM_{2.5}. PM_{2.5} penetrates deep into the lungs and enters the bloodstream, where it can lead to negative

¹¹ Heat input is a measure of the amount of fossil fuel used to generate electricity.

¹² See <https://www.epa.gov/acidrain/effects-acid-rain>, accessed March 9, 2017.

¹³ See Environmental Protection Agency (2009) for more information.

cardiovascular and respiratory effects.¹⁴ For a review of quasi-experimental evidence of the health effects of short-run exposure to pollutants such as SO₂ and PM_{2.5}, see Graff Zivin and Neidell (2013).

Effects from long-run PM exposure may accrue after sustained, low-level exposure, similar in spirit to cigarette smoking;¹⁵ extended exposure to PM 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).¹⁶ Quantifying such long-run effects is challenging as residential sorting on environmental quality can bias estimates. Quasi-experimental studies based on transitory shocks cannot directly address such bias – an unexpected change in pollution may disrupt the housing market equilibrium in the short-run, but effects become muted or disappear in the long-run as people adjust to this change by, for example, moving in or out of the affected area. This complicates separating long-run exposure effects of pollution from characteristics of a changing population, shifts in economic growth, and other health-related adjustments that come with lasting changes to the local economy. Banzhaf and Walsh (2008) provide evidence supporting environmental Tiebout sorting, demonstrating that areas with exogenous changes in air quality then experience changes in both size and composition of the population.

Two recent quasi-experimental papers attempt to address confounding from sorting using cross-sectional designs. Chen et al. (2013) exploit the government policy of free winter heating north of the Huai River in China, which led to persistent higher total suspended particle

¹⁴ 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 it is quite limited (Thatcher and Layton, 1995; Ozkaynak et al., 1996; and Vette et al., 2001).

¹⁵ 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).

¹⁶ The biological pathways differ between cigarette smoking and PM_{2.5} exposure, however, given the different substances involved.

emissions to the north. Restricted migration enforced by the government strengthens the research design, as it precludes sorting across the river. The authors find life expectancy for those living north of the river is 5.5 years shorter than those south of the river.

Anderson (2015) exploits wind patterns along major highways in Los Angeles as a source of differential variation in exposure to pollution. His analysis focuses on 3-year mortality rates for adults over age 75, comparing residents living upwind versus downwind from major highways, showing a doubling in downwind exposure leads to a 3.6 to 6.8 percent increase in mortality for this age group.

We offer two primary innovations over these quasi-experimental studies exploring long-run exposure effects. First, we explicitly estimate the dynamic relationship between pollution changes and mortality, allowing observation of differential mortality effects immediately and up to 10 years into the ARP. Prior studies using cross-sectional analyses provide estimates that represent a confluence of short-, medium-, and long-run effects, while we trace effects over time with graphical evidence. Our framework tests whether effects of the ARP's pollution reduction are rapid or take time to build, and whether effects are transient or lasting. We observe monotonically increasing and lasting improvements in mortality over time, supporting the hypothesis that long-run cumulative exposure is an important health input.

Second, we focus our analysis explicitly on working-age adults, a group typically considered a more health-robust segment of the population. Much of the earlier work focused on more vulnerable groups, such as infant and the elderly, and we show results for these groups as well for completeness. The working-age range is one where mortality represents a significant loss in life expectancy and economic productivity. Commonly used estimates of the value of

statistical life (VSL), typically based on hedonic wage regressions for the working-age population, more readily apply to this segment of the population.

3. Data and methods

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. Data include information on age at death and county of residence – we use the restricted data since county identifiers are not publicly available for counties with less than 100,000 people, which make up a large share of our sample.

We specify the main dependent variable as the inverse hyperbolic sine (IHS) of deaths per 100,000 residents aged 35-64.¹⁷ For the denominator, we use 1990 population from the decennial census. We favor a stable population denominator because between-census population estimates impose trending, which may confound dynamic effects given we are testing for both breaks in trend (short-run effects) and changes in trend (long-run effects). For these death estimates to be true indications of changes in health, we must assume county-level population changes do not correlate with the ARP; we show this assumption holds using county-specific migration data. To further test this concern, we also show results using annually adjusted

¹⁷ $IHS(x) = \ln(x + \sqrt{1 + x^2})$. The use of the IHS function is approximate to the log in that, for small changes, one can interpret marginal effects as percentage changes. However, the IHS of zero is mathematically defined. See Card and DellaVigna (2017) for a recent application of the IHS. We favor the IHS model as some counties have zeros in mortality in some years, which becomes especially prevalent when we explore particular causes of death. We also test results using the more standard log for all internal deaths (where zeros are uncommon), and find comparable results. Were we to use the log function, using deaths per constant (1990) population would be equivalent to log of raw deaths. This is a result of including county fixed effects and the mathematical properties of log functions where $\ln(a/b) = \ln(a) - \ln(b)$. The IHS function does not have this property, so $IHS(a/b) \neq IHS(a) - IHS(b)$ and using death rates with a constant denominator yields slightly different results than raw deaths. Our results using IHS of raw deaths, however, are similar to our main results.

population data from the Surveillance, Epidemiology, and End Results (SEER) program, which generates between-census estimates of population sizes by various age groupings. For population-adjusted estimates to be a true indication of health changes, county-level inter-census estimates must not correlate with the ARP. In practice, both sets of estimates are similar.

We focus on the death rate for internal causes, and further estimate impacts on respiratory-related and cardiovascular-related deaths to explore the causal channel.¹⁸ We also perform a robustness check using external deaths (accidents, murders, etc.) as a falsification test; such deaths may change with potentially confounding factors such as selective migration and economic fluctuations, but not with pollution-related health improvements. In most regressions, we limit analysis to those between 35 and 64 years of age, but show internal death results for other age groups as well.

Sparse data prevent us from showing a direct comparison between pre- and post-ARP levels of $\text{PM}_{2.5}$.¹⁹ Instead, we demonstrate a link between the ARP and SO_2 using monitor-level pollution data. Since SO_2 can convert to $\text{PM}_{2.5}$, evidence of a link between the ARP and SO_2 demonstrates proof of concept for an effect on $\text{PM}_{2.5}$. The EPA reports daily SO_2 averages in parts-per-billion (ppb). We convert daily monitor-level measures to annual 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 county 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. We used an unbalanced panel of sensors given the limited number that are active across the entire sample period, but focus our SO_2 results on a balanced panel of counties.

¹⁸ Since there was a major change in adult AIDS-related mortality in 1995 with the introduction of new drug technologies, we omit internal deaths from AIDS to improve precision. Results including AIDS are similar, which we expect given rollout of the new drug should not correlate with location of ARP-regulated power plants.

¹⁹ $\text{PM}_{2.5}$ was not officially regulated by the EPA until 1997, so monitoring before 1997 is insufficient in scope.

In the off chance adverse weather patterns differed by treatment status after 1995, we flexibly control for temperature and rainfall using data from the National Oceanic and Atmospheric Administration (NOAA). We aggregate the station-day variables up to the county-year level using inverse distance squared weights. Given the smaller number of available weather stations across the time period, we expand the maximum distance to 100 miles between county centroid and weather stations.

To test for ARP impacts on the local economy, we collect annual, county-level data from various sources. We use county-level wage employment and total income (including wages and government payments) from the Bureau of Economic Analysis (BEA) Regional Data to explore labor market effects. The relevant data are from the Local Area Personal Income files, which provide annual economic and demographic information at the county level. Our first measure, wage employment, describes total county-level employment. The BEA derives estimates using information from unemployment insurance, and then adds information from uncovered industries such as some non-profits, students, and farmers. Personal income is income received from all sources, including wages and government transfers, for all county residents.²⁰ In each case, we use the IHS of raw values rather than per-capita estimates, given the imposed trending of population between census estimates mentioned above.

We use county-level population migration data from the Internal Revenue Service (IRS) to explore possible shifts in population. The IRS data include the number of exemptions claimed on tax returns that both entered and exited each given county in a given year, which serves as a proxy for migration.²¹ We focus on net migration (the difference between inflows and outflows) divided by baseline population to align with our mortality estimates.

²⁰ This variable omits realized and unrealized capital gains and losses.

²¹ See Molloy et al (2011). These data do not permit an analysis of net migration for specific age groups.

3.2 Treatment definitions

To define treatment areas, we link the mortality and SO₂ data with proximity to Phase I power plants. We obtain a list of all Phase I plants from the EPA Air Markets Program Data, which includes the GPS coordinates of the plant. As no Phase I plants exist further west than approximately -96 degrees in longitude, we limit counties to those with centroids east of -100 degrees longitude. Our core treatment definition involves a binary classification of treatment. A county is "treated" if its geographic centroid falls within a given distance from any Phase I plant. We focus our baseline analysis on 100 miles. We base this distance choice off flexible, reduced form distance estimates we explore in Section 4.6. Although this binary classification ignores the intensity and directionality of emissions, it allows simple graphical presentations of our results, and our use of propensity-score to better align treatment and control counties. Further, our measure of emissions is already a proxy for the true pollutant of interest, PM_{2.5}, which makes us further favor a more general binary model.

Figure 1 shows a map of treatment areas using the binary classification. Table 1 shows summary statistics for the treatment and control counties using the binary classification of treatment. Note both Figure 1 and Table 1 are net of propensity-score sample trimming, which we discuss below.

3.3 Econometric model

To identify the link between the ARP, SO₂ levels, and mortality across time, we estimate the following equation:

$$y_{ct} = \beta_t + \sum_i \beta_i * \tau + \sum_i \delta_i * psc_c + \theta X_{ct} + \alpha_c + \varepsilon_{ct}.$$

The variable y is either the IHS of the mortality rate for the population aged 35-64 or the SO_2 level in county c in year t . We use a fixed age group over time to focus specifically on the prime working age population.²² The variable p_{sc} is a flexible use of a propensity-score estimate reflecting the probability of treatment, which we discuss in detail below. The vector X_{ct} includes weather variables to address the possibility of differential weather patterns in treatment counties.²³ The vector α_c captures county fixed effects, while 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 1990 age group-specific population.

Our treatment variable, τ , is an indicator for whether county c is within 100 miles of a regulated Phase I power plant. The treatment effect varies separately by year (β_t), omitting the year 1994. This amounts to an annual difference-in-differences estimator where we compare the difference in mortality between treatment and control areas in each year relative to the baseline difference in 1994.

The dynamic pattern of the estimates is critical for testing the importance of long-run exposure. To illustrate this point, consider two individuals with identical mortality risk in the absence of any differences in short- and long-run pollution exposure. Assume there is a downward shift in the pollution level in year t for only one of the individuals such that pollution decreases by amount q at year t (1995 for our purposes), and remains lower by q in all future years. In year $t+n$, this “treated” individual experiences $n*q$ less cumulative (or “long-run”) pollution along with q less immediate (or “short-run”) exposure, which would lower the mortality risk of this individual compared to the counterfactual “untreated” individual. With only

²² This approach, which is similar in spirit to Case and Deaton (2015), assumes the age distribution within a county is uncorrelated with treatment status. Results using fixed cohorts rather than fixed age groups are similar to our main estimates.

²³ We control for precipitation and precipitation squared, as well as the fraction of each 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.

mortality data in year $(t+n)$, one cannot disentangle short-term exposure from long-run exposure. However, with data on mortality risk for years t through $t+n$, one can separately identify both effects based on the changes in mortality risk over time. Specifically, if only a short-run mortality effect from exposure exists (e.g., a "shock" effect), the differential in mortality would decline at year t only and remain level in the following years.²⁴ If long-run (cumulative) exposure is an important health input, the relative mortality risk for the treated individual will continue to decrease in each subsequent year.²⁵

Applying this to our context, we can infer the ARP had no observable impact on health if all estimates of β_t are zero. If estimates are non-zero and constant from 1995 onward, this represents 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., $\beta_{1995} < \beta_{1996} < \dots < \beta_{2005}$), this presents evidence of an effect from cumulative, long-run exposure.

A core identifying assumption of our model is that treatment and control counties would experience a similar mortality pattern over time absent the ARP's implementation. As one test of this "parallel trends" assumption, we explore whether trends in mortality were similar before the ARP by obtaining estimates for each year leading up to implementation.

To control for the possibility that demographic and economic trends correlate with proximity to regulated plants, we estimate a propensity-score (p-score) of treatment, using pre-treatment county-level information from the 1990 County Data Book. We bin the p-score into

²⁴ If there are latent effects, there should be no short-run changes, with all improvements coming in some time $t > 1$; while we cannot rule this possibility out, these latent effects are more likely to arise during early childhood when critical development periods occur (see, e.g., Almond and Currie, 2011), whereas we focus largely on adults.

²⁵ Note that if relative pollution continued to decline after t , it would complicate separating the cumulative effects from the continued decrease in pollution using our framework. Our estimates suggest the majority of the SO₂ reductions from the ARP for the treatment group occurred in the years directly following 1995.

deciles and interact decile indicators with year fixed effects.²⁶ In essence, this flexibly allows treatment and control counties of similar p-score (and thus similar 1990 traits) to have similar background changes in mortality outcomes over time.

We estimate the p-score from a logit model with the treatment indicator as the outcome and the following economic and demographic variables as predictors of treatment (in linear, quadratic, and interacted forms): land area in square miles, total population, population share classified as black, population share classified as Hispanic, population share classified as male, population share by County Data Book age group range (5-17, 18-21, 21-24, 25-34, 35-44, 45-54, 55-64, 65-74, and > 75, with under 5 as the omitted group), income per capita, civilian unemployment, civilian employment in manufacturing, retail, public administration, health services, and employment by federal and state government.²⁷ Appendix Figure A-1 shows the distribution of estimated p-scores in bins of 0.1 for both treatment and control groups. We trim extremes of treatments and controls by dropping counties with p-scores below 0.01 and above 0.99 (removing 337 total counties), and test both more extreme and zero trimming as robustness checks.

Three factors drive us to opt for this p-score method over directly controlling for county-year demographic or economic covariates or county-specific time trends.²⁸ First, the available county-year demographic data are interpolated between censuses, which results in noisily measured, trending variables that can hinder our analysis. Second, controlling for county-year

²⁶ The choice between the number of n-tiles for binning faces a trade-off between increased flexibility (more bins) and desire to have treatment and control counties in each bin (fewer bins). Using both a smaller and larger number of bins yields similar results.

²⁷ We follow the reasoning of Imbens (2015) for deciding which covariates to include; theory serves only as a basic guide as to what covariates are likely to be important, and the larger goal is an accurate approximation of the conditional expectation. Our choice of these particular covariates was driven by the aim to equate pre-trends in mortality across treatment and control groups.

²⁸ As a robustness check, we allow the effect of the initial demographic variables to vary over time by interacting them with year fixed effects as well.

variables can introduce over-controlling bias if treatment changes mortality by otherwise changing the composition of the relevant population (Angrist and Pischke, 2009). Third, controlling for county-specific time trends potentially biases estimates in the presence of a dynamic treatment effect (Wolfers, 2006). Trend controls may help identification by accounting for pre-existing differences, but may also remove the signal of the policy. The p-score bin-by-year variables flexibly account for pre-trends without picking up the dynamic treatment effect and, as we show below, remove any concern of pre-existing trends. We also explore an alternative model specification that interacts 1990 covariate values individually with year fixed effects. This is more flexible in that it allows each covariate to have differential impacts by year, but less flexible in that it does not allow for higher order or interactive effects of these variables.

4. Results

We begin our results with a visual demonstration of our event study, and then follow with tables showing regression coefficients and standard errors.

4.1. SO₂ results

Evidence indicates the decrease in SO₂ after the ARP's implementation was both immediate and persistent. Panel A of Figure 2 presents SO₂ data by treatment status absent controls. Two features emerge: (1) SO₂ levels trend down over time in both treatment and control areas, and (2) SO₂ levels in control areas evolve continuously across the start of the ARP, while SO₂ levels in treatment areas drop sharply and return to trend at new, lower levels. Panel B of Figure 2 presents the event study coefficient estimates, which show differential SO₂ levels across treatment and control areas, adjusting for county and year fixed effects, weather factors, and p-

score interactions. In all event-study graphs, dashed lines indicate 95 percent confidence intervals. Aligning with the raw data, prior to the ARP, there is no trending difference in SO₂. Relative SO₂ levels drop rapidly after the ARP, and largely stabilize by 1999 at a 1 ppb lower level relative to before. The reduction in short-run exposure is mostly constant in all years after the ARP so we can isolate the effects of short- and long-run exposure. This immediate and stable reduction in SO₂ is our identifying variation to estimate the effects of long-run exposure; individuals in treated counties experienced lower cumulative exposure to pollution in the years after the ARP's implementation, relative to the counterfactual control counties.

4.2. Mortality results

We first focus our analysis on individuals age 35-64. Primary outcomes are IHS of death rates per 100,000 (of 1990 population) for the relevant age-group population at the county level. Panel A of Figure 3 shows the raw plot for internal deaths by treatment and control. Prior to the ARP, mortality in both areas is trending upward in a parallel fashion. Note the upward trending is due to the increase in age group population size and not from a worsening of health.²⁹ After the ARP, internal mortality in treatment areas experiences a subtly diminished trend relative to 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.

In Panel B of Figure 3, we present the event study using coefficients and standard errors from our main regression equation. Coefficients represent an approximate percentage change for the difference between treatment and control counties as compared to the 1994 difference. Consistent with the raw plots, we see no statistically or economically significant difference in trends before the ARP; no coefficients estimating the difference in internal mortality between

²⁹ Recall that we use the constant 1990 population in the denominator of our death rate.

treatment and control groups from 1990-1993 are statistically or economically different from the 1994 difference (omitted year). After the ARP, the estimates imply a gradual and steady decrease in mortality in treatment areas relative to controls, with effects becoming statistically significant by 1997. In 1995, the first year of ARP implementation, the estimate is a decrease in mortality of 1.3 percent, though not statistically significant at conventional levels. This estimate increases to a statistically significant 5.1 percent in 2005. Using 1994 mortality levels as a baseline, this translates to approximately 7,300 fewer deaths in 2005 for the relevant age group in treatment counties.³⁰ A gradually increasing coefficient, rather than a single shift, implies cumulative exposure to the ARP is an important health input.

4.3. Results by cause of death

As evidence supporting $PM_{2.5}$ as the driving force behind observed mortality impacts, we explore effects by cause of death. SO_2 primarily affects respiratory health, whereas $PM_{2.5}$ affects mortality through changes in cardiovascular functioning as well as respiratory. If we find evidence of an impact on cardiovascular deaths, this supports at least some role of $PM_{2.5}$. Panels A and B of Figure 4 show changes in cardiovascular mortality and respiratory mortality (respectively) as primary cause of death, which largely mirror those for internal cause mortality. Results for respiratory deaths, while much larger in some years, are noisier and exhibit less of a reliable pattern, though there are large and statistically significant values in the later sample years. Overall, the pattern of differential effect by cause of death supports $PM_{2.5}$ exposure as a causal mechanism.

³⁰ To calculate these values, we first estimate an average change in mortality rates for treatment counties for each year, which we take from our event-study results. We then multiply total deaths in 1994 by the estimated percentage change in deaths in 2005, as compared to 1994.

4.4. Effects by other age groups

Although we focus our estimates on those aged 35-64, effects may also exist for other groups. In Figure 5, we present event study results for 4 additional age groups: infants (age 0), the very young (1-14), the young working population (15-34) and the above 65 population.³¹

Existing evidence supports the presence of mortality effects for infants (see, e.g., Chay and Greenstone (2003) and Currie and Neidell (2005)). However, such effects would not increase with time; because infants only remain so for one year, there are no true "long-run" cumulative health effects where improvements build beyond year 1. Consistent with this, infant mortality follows a similar pattern to SO₂ (see Figure 2), with a downward shift after the ARP. Although the estimates are noisy, the effect for this age group is not increasing with time.

For both the very young and working young adult age groups, we see no statistically significant change or pattern in mortality rates across the ARP. This may suggest mortality among this generally healthy group is unaffected in the 10 years following the ARP, but estimates are noisy and confidence intervals contain potentially economically significant effects. We also cannot assess the potential for effects beyond 10 years. For the over 65 group, we observe similar effect sizes as for our core sample of 35-64 year olds, with estimates statistically significant across all years and an effect around 4 percent by 2005. Given the larger mortality rate for this age group, this translates to approximately 26,000 fewer deaths in 2005.

4.5. Testing for sorting and economic effects

If the ARP affected local economic activity or population characteristics, this would present a potential threat to identifying the direct health impacts of pollution. As one

³¹ Due to different baseline rates of mortality across age groups, the sizes of coefficients and confidence intervals vary considerably by age group. As a result, the y-axis is not constant across figures.

hypothetical example, workers may lose jobs, which might bias estimates towards finding the ARP-related pollution reductions worsened health. Alternatively, wealthier populations may relocate near clean plants given the improved air quality, biasing estimates towards finding the ARP improved health. The broad treatment areas of our design are much larger than the range at which these economic effects generally take place (see, for example, Currie, Davis, Greenstone, and Walker 2015); our treatment area extends 100 miles from plants while evidence suggests housing market effects exist within 1-2 miles. We nonetheless probe this possibility given the importance of this assumption.

We test effects on local economic activity by repeating our event studies using total wage employment and personal income as outcomes. For economic activity to fully explain observed mortality changes, not only would these outcomes need to be economically and statistically significant, but they must also follow the pattern observed for mortality: no trends prior to the ARP, and a building divergence after the ARP. Neither is apparent for either employment or income. Panel A of Figure 6 shows no trend in total wage employment differences prior to the ARP, and the pattern remains flat after the ARP. All estimates are small in magnitude and statistically insignificant. In Panel B, the ARP shows little correlation with personal income. We observe a small drop-off in income in the last year of our sample, but this estimate is not statistically significant. These results suggest changes in economic activity do not explain observed post-ARP changes in mortality, which is consistent with prior findings that the ARP's economic effects were quite small (Schmalensee and Stavins 2015).

We next test for migratory sorting using two additional outcome variables: net-migration and mortality from external causes. If the ARP led to selective migration, where less healthy people moved to control counties or more healthy people moved to treatment counties (or a

combination of both), this could cause a decline in mortality rates that correlates spuriously with changes in pollution. We test for this selection issue using migration data from the IRS as the dependent variable. Panel C of Figure 6 shows net migration (inflows minus outflows) of listed exemptions, expressed as a share of 1990 population.³² The estimates are statistically insignificant and do not follow a trend like that of observed mortality effects.

While net migration is informative about net changes in population, it cannot address issues with changes in population composition. For example, higher income individuals could displace lower income individuals with no net change in population. Unfortunately, data on population composition at the county-year level are not available except where such measures are interpolated between surveys. As an indirect test of this hypothesis, we examine if the ARP affected deaths from external causes. Such deaths are sensitive to demographic changes and local economic activity, but not pollution exposure. Panel D of Figure 6 illustrates that external mortality rates largely do not change after the ARP's implementation. There is a single statistically significant difference in 2005, but the external mortality estimates do not display a monotonic change in mortality risk like we observe for internal causes of death, suggesting population composition does not explain our main results.

4.6. Robustness checks

We present a series of results in Table 2 from alternative specifications to assess the robustness of our results, focusing on the models using all internal deaths. Column 1 includes only county and year effects, column 2 adds controls for weather, and column 3 replaces general year effects with our interacted propensity-score year effects, which matches our event study graphs. Adding weather has a minimal impact on estimates. Adding propensity-score decile-by-

³² We do not use the IHS or log conversion here because net flows can be negative.

year controls reduces the size of our estimates, though they still fall within approximately a standard error of the estimates in column 1. Adjusting for propensity-score differences removes any evidence of pre-treatment differences in treatment and control groups. Appendix Table A-1 shows estimates used in all other event studies (cause of death, age groups, and economic/population effects).

Table 3 expands on robustness checks by considering alternate regression models. In column 1, we control for the log of age-specific population using the SEER between-census estimates. Estimates are statistically unchanged but slightly smaller in absolute magnitude, now reaching around 4.5 percent in 2005. Column 2 uses SEER data to control for the age distribution by year: percentage of the population of ages 35-44, 45-54, and 55-64 (rescaled for these three categories to sum to 1). Estimates are relatively stable to including the age distribution, which we take as evidence population and age makeup shifts are exogenous to treatment and/or our baseline population and age distribution variables in the propensity-score control for any consequential differences. In column 3, we use the log function rather than the IHS to transform changes in mortality. Coefficients are unchanged, though there is a small loss in the number of observations: a few county/year cells contain zero deaths for the 35-64 year-old age group.³³

In column 4, rather than include the p-score decile indicators interacted with year dummies, we interact the 1990 baseline covariates used to construct the p-score with year fixed effects. This is more flexible in that each variable can have differential effects by year, but less flexible in that all effects are linear (while the propensity-score allowed for quadratic and interacted terms). Results are again stable to these changes.

³³ We do not present results using the log function for all of our results because of a greater fraction of zeros for certain dependent variables. For example, less than 1% of observations had zero internal deaths, whereas nearly 5% had zero external deaths.

In columns 5 and 6 we explore different assumptions regarding how we trim by the propensity-score. In column 5, we use all counties (no trimming), and in column 6 we trim further by cutting all counties with a p-score below 0.05 and above 0.95. Results are quantitatively similar.

Finally, in column 7, we drop all counties that contain a treatment plant. This addresses the possibility that economic and behavioral effects are confined with the political boundaries of the county where the power plant resides. Estimates are virtually identical to the full sample of counties.

4.7. Variations by treatment distance

Our primary results use a treatment distance of 100 miles as the cutoff. Given the absence of reliable PM_{2.5} data at the time, limited availability of balanced SO₂ monitors, and a variable conversion rate between SO₂ and PM_{2.5}, we cannot directly observe the distance range of ground-level changes to our pollutant of interest or draw inference from SO₂ dispersion. Instead, we arrived at the 100-mile distance cutoff by running our basic regression model for internal mortality using a variety of different distance ranges in defining treatment. Specifically, we estimate 8 regression models where we vary the treatment group (counties within 0-25 miles, 25-50 miles, ... 175-200 miles) but keep the control group constant (counties beyond 200 miles), allowing us to assess the marginal contribution of each treatment group. We limit this regression model to a simple pre-test/post-test design rather than allow dynamic effects for ease of exposition, but otherwise follow the model of column 2 in Table 2.³⁴

³⁴ We also produced the full event study results using various definitions of treatment and control (available upon request), which were consistent with these results. As use of a propensity-score requires a definition of treatment, which would result in 8 propensity-score models, we instead opt for the more simplistic model controlling for county and year fixed effects and weather variables.

The treatment effects, plotted in Figure 7, show a slightly non-monotonic relationship between treatment distance and mortality effects. The estimate for the 0-25 mile treatment group is roughly 4.6 percent, but increases to 12.5 percent for 25-50 miles. The estimate decreases to 6.0 percent for 50-75 miles, 4.7 percent for 75-100 miles, and around 2.7 percent at 100-125 miles. All estimates using treatment distances less than 75 miles are statistically significant, with the estimate using 75-100 miles for the treatment group reaching borderline significance.³⁵ This figure supports our choice of 100 miles as the treatment boundary.

4.8. Benefit valuation

To gauge the magnitude of the estimated mortality effects, we calculate the avoided deaths per year in 10-year age bins for deaths occurring in people over age 35, the ages at which we consistently observe statistically and economically significant decreases in mortality.³⁶ Table 4 shows the results, indicating the number of avoided deaths increasing over time within each age group. The number of avoided deaths is higher in the older age groups (within any given year), which is consistent with the fact that mortality risk increases with age.

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. Using age-specific VSLs, we calculate a mortality benefit of \$134 billion in 2005 alone. If

³⁵ The possible effects for 100-125 and 125-150 miles suggest a potential treatment definition of 0-150 miles for our baseline specification (and beyond 150 miles as control). We estimated such models and found generally comparable results (available upon request).

³⁶ We produce the avoided deaths per age bin by re-estimating our model separately for each age bin.

instead we use a constant VSL of \$4.9 million³⁷, we estimate mortality benefits of \$164 billion in 2005, which is nearly 25 percent larger. Interestingly, when we compare benefits using age-specific VSL estimates to previous model-based estimates of the mortality benefits from the ARP, the results are quite close in magnitude. Previous simulation-based studies estimate the benefits at over \$100 billion per year using a constant VSL of roughly \$6 million applied to all ages (Chestnut and Mills, 2005).³⁸ The similarity between estimates is coincidental – we find larger decreases in mortality across all age groups but value them using a smaller, age-appropriate VSL.

5. Conclusion

The Acid Rain Program caused lasting improvements in ambient air quality, with mortality benefits that accrued to adults gradually over time. We show internal mortality rates among the working-age population (ages 35-64) decreased in a manner consistent with extended, long-run pollution exposure being an important health input. While emissions reductions were rapid, mortality rates decreased only slightly in the initial years following the ARP. The reductions in mortality risk grew to 5 percent after 10 years, suggesting health effects accrued after a change in sustained exposure.

Importantly, we document mortality effects for an age group not typically addressed -- previous work largely focuses on traditionally more susceptible groups, such as infants and the elderly. Our findings suggest long-run exposure to pollution may have detrimental effects on everyone, regardless of underlying health capital. Furthermore, they support appropriateness of valuing the mortality improvements using a VSL based on labor market studies. Regardless of

³⁷ This is the implied constant VSL from Aldy and Smyth (2014).

³⁸ The authors allow the VSL to vary over time, but it is constant across all age groups at a given point in time.

the method used to calculate the mortality benefits, the ARP stands out as one of the most cost-effective environmental regulations to date. We estimate the value of benefits in terms of annual lives saved at \$134 billion by 2005, compared to previously estimated program costs of around \$3 billion per year (Chestnut and Mills, 2005). Although data limitations preclude us from obtaining estimates of the marginal effects of $PM_{2.5}$, our results show that understanding the link between long-term pollution improvements and mortality plays an important role in evaluating efficiency of other programs that affect $PM_{2.5}$ concentrations.

We focus on the mortality effects from long-run exposure, but 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. 2016a, 2016b), 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.

References

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

Almond, Douglas and Janet Currie (2011). "Killing Me Softly: The Fetal Origins Hypothesis," *Journal of Economic Perspectives*, 25, 153-172.

Anderson, Michael (2015). "As the Wind Blows: The Effects of Long-Term Exposure to Air Pollution on Mortality." NBER Working Paper 21546.

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

Banzhaf, H. Spencer, and Randall P. Walsh. 2008. "Do People Vote with Their Feet? An Empirical Test of Tiebout's Mechanism." *American Economic Review* 98: 843-63.

Bharadwaj, Prashant, Joshua Graff Zivin, Matthew Gibson, Christopher A. Neilson (2017). "Gray Matters: Fetal Pollution Exposure and Human Capital Formation." *Journal of the Association of Environmental and Resource Economists*, 4(2): 505-542..

Brook RD, Rajagopalan S, Pope CA 3rd, Brook JR, Bhatnagar A, Diez-Roux AV, Holguin F, Hong Y, Luepker RV, Mittleman MA, Peters A, Siscovick D, Smith SC Jr, Whitsel L, Kaufman JD; 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 and Stefano DellaVigna (2017). "What Do Editors Maximize? Evidence from Four Leading Economics Journals." NBER Working Paper 23282.

Case, Anne and Angus Deaton (2015). "Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century." *Proceedings of the National Academy of Sciences*, 112(49): 15078-15083.

Chang, Tom, Joshua Graff Zivin, Tal Gross, and Matthew Neidell (2016a). "The Effect of Pollution on Worker Productivity: Evidence from Call-Center Workers in China." NBER Working Paper 22328.

Chang, Tom, Joshua Graff Zivin, Tal Gross, and Matthew Neidell (2016b). "Particulate Pollution and the Productivity of Pear Packers." *American Economic Journal: Economic Policy*, 8(3): 141-169.

Chen, Yuyu, Avraham Ebenstein, Michael Greenstone, and Hongbin Li (2013). "Evidence on the impact of sustained exposure to air pollution on life expectancy from China's Huai River policy." *Proceedings of the National Academy of Sciences*, 110(32): 12936-12941.

Chestnut, Lauraine and David Mills (2005). "A fresh look at the benefits and costs of the US acid rain program." *Journal of Environmental Management* 77:25-266.

Currie, Janet, Lucas Davis, Michael Greenstone, and Reed Walker (2015). "Environmental Health Risks and Housing Values: Evidence from 1600 Toxic Plant Openings and Closings." *American Economic Review*, 105(2): 678-709.

Deryugina, Tatyana, Garth Heutel, Nolan H. Miller, David Molitor, and Julian Reif (2016). "The Mortality and Medical Costs of Air Pollution: Evidence from Changes in Wind Direction." NBER Working Paper 22796.

Deschenes, Olivier, Michael Greenstone, and Joseph Shapiro (2012). "Defensive Investments and the Demand for Air Quality: Evidence from the NOx Budget Program." NBER Working Paper 18267.

Dominici, Francesca, Michael Greenstone and Cass R. Sunstein (2014). "Particulate Matter Matters." *Science*, 344 (6181): 257-259.

Environmental Protection Agency (2004). [Air Quality Criteria for Particulate Matter](#). National Center for Environmental Assessment. Research Triangle Park, NC.

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 and Matthew Neidell (2012). "The Impact of Pollution on Worker Productivity." *American Economic Review*, 102(7).

Graff Zivin, Joshua and Matthew Neidell (2013). "Environment, Health, and Human Capital," *Journal of Economic 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." *Journal of Political Economy*, 110(6): 1175-1219.

Hodan, William M., and William R. Barnard (2004). "Evaluating the Contribution of PM_{2.5} Precursor Gases and Re-entrained Road Emissions to Mobile Source PM_{2.5} Particulate Matter Emissions." MACTEC Federal Programs, Research Triangle Park, NC.

Imbens, Guido (2015). "Matching Methods in Practice: Three Examples." *Journal of Human Resources*. 50(2): 373-419.

Isen, Adam, Maya Rossin-Slater, and Reed Walker (2017). "Every Breath You Take – Every Dollar You'll Make: The Long-Term Consequences of the Clean Air Act of 1970." *Journal of Political Economy*, 125(3): 848-902..

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak. 2011. "Internal Migration in the United States." *Journal of Economic Perspectives*, 25(3): 173-96.

Murphy, Kevin and Robert Topel (2006). "The Value of Health and Longevity." *Journal of Political Economy*. 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." *Journal of Exposure Analysis and Environmental Epidemiology*, 6(1), 57-78.

Sanders, Nicholas J. (2012). "What Doesn't Kill you Makes you Weaker: Prenatal Pollution Exposure and Educational Outcomes." *Journal of Human Resources*. 47(2): 826-850.

Schlenker, Wolfram and Reed Walker (2016). "Airports, Air Pollution, and Contemporaneous Health." *Review of Economic Studies*, 83(2): 768-809.

Schmalensee, Richard and Robert Stavins (2015). "Lessons Learned from Three Decades of Experience with Cap-and-Trade." NBER Working Paper 21742.

Stavins, Robert N. (1998). "What Can We Learn from the Grand Policy Experiment? Lessons from SO2 Allowance Trading." *Journal of Economic Perspectives*, 12(3): 69–88.

Sullivan, Daniel and Till von Wachter (2009). "Job Displacement and Mortality: An Analysis using Administrative Data." *Quarterly Journal of Economics*, 124(3): 1265-1306.


Thatcher, T. L., & Layton, D. W. (1995). "Deposition, resuspension, and penetration of particles within a residence." *Atmospheric Environment*, 29(13), 1487-1497.

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 Science & Technology*, 34(1), 118-126.

Walker, Reed (2013). "The Transitional Costs of Sectoral Reallocation: Evidence From the Clean Air Act and the Workforce." *Quarterly Journal of Economics*, 128 (4): 1787-1835.


Wolfers, Justin (2006). "Did Unilateral Divorce Raise Divorce Rates? A Reconciliation and New Results." *American Economic Review*, 96(5): 1802-1820.

Figure 1
Treatment and Control Counties




Notes: Map shows all treatment (dark) and control (light) counties included in our analysis. White counties indicate omission, either due to lack of data or trimming of highest and lowest percentile by propensity score.

Figure 2
Sulfur Dioxide by Treatment Status, 1990-2005




Notes: Outcome is SO₂ in parts per billion. We include an unbalanced panel of SO₂ data to expand the number of available counties for analysis (but a balanced panel of counties). Panel A shows mean SO₂ by treatment and control counties, weighted by number of SO₂ sensors in 1990. Panel B shows our event study, where each coefficient in event study represents an interaction between treatment (within 100 miles of a regulated plant) and year, with 1994 as the omitted year range. Regressions span 1990-2005 and include county fixed effects, weather controls, and decile-binned p-score indicators interacted with year fixed effects. We trim the sample on 1% p-scores. We cluster standard errors at the county level.

Figure 3
Internal Mortality for Ages 35-64, 1990-2005




Notes: Outcome is inverse hyperbolic sine (IHS) of internal mortality rate per 100,000 age-specific population, using 1990 population as a constant denominator. Panel A shows IHS of mean mortality by treatment and control counties, weighted by 1990 age-specific population. Panel B shows our event study, where each coefficient in event study represents an interaction between treatment (within 100 miles of a regulated plant) and year, with 1994 as the omitted year range. Regressions span 1990-2005 and include county fixed effects, weather controls, and decile-binned p-score indicators interacted with year fixed effects. We trim the sample on 1% p-scores. We cluster standard errors at the county level.

Figure 4
Event Study by Primary Cause of Death, 1990-2005




Notes: Outcome is inverse hyperbolic sine (IHS) of cause-specific mortality rate per 100,000 age-specific population, using 1990 population as a constant denominator. Panel A shows deaths classified as cardiovascular. Panel B shows deaths classified as respiratory. Both graphs show our event study, where each coefficient in event study represents an interaction between treatment (within 100 miles of a regulated plant) and year, with 1994 as the omitted year range. Regressions span 1990-2005 and include county fixed effects, weather controls, and decile-binned p-score indicators interacted with year fixed effects. We trim the sample on 1% p-scores. We cluster standard errors at the county level.

Figure 5
Event Study by Age Group, 1990-2005




Notes: Outcome is inverse hyperbolic sine (IHS) of cause-specific mortality rate per 100,000 age-specific population, using 1990 population as a constant denominator. Panels are split by indicated age ranges. All graphs show our event study, where each coefficient in event study represents an interaction between treatment (within 100 miles of a regulated plant) and year, with 1994 as the omitted year range. Regressions span 1990-2005 and include county fixed effects, weather controls, and decile-binned p-score indicators interacted with year fixed effects. We trim the sample on 1% p-scores. We cluster standard errors at the county level.

Figure 6
 Event Study of Potential Migration and Economic Measures, 1990-2005



Notes: Each panel shows a different regression outcome, used as tests of potential confounders to our private estimates. Panels A and B show the IHS of total wage employment and total personal income, respectively. Panel C shows net migration for all ages (inflows minus outflows, as measured by IRS data on exemptions) divided by 1990 baseline population. Panel D shows IHS of the mortality rate for all deaths attributed to external cases for the 35-64 age range. All graphs show our event study, where each coefficient in event study represents an interaction between treatment (within 100 miles of a regulated plant) and year, with 1994 as the omitted year range. All regressions weight by 1990 age 35-64 population as in our main mortality regressions. Regressions span 1990-2005 and include county fixed effects, weather controls, and decile-binned p-score indicators interacted with year fixed effects. We trim the sample on 1% p-scores. We cluster standard errors at the county level.

Figure 7
Treatment Effects by Varied Distances



Notes: Regressions span 1990-2005 and include county fixed effects, year fixed effects, and weather controls, following column 2 of Table 2. We derive each coefficient and 95% confidence interval from a separate regression. The sample of included counties varies by indicated distance bin. In all cases, we include counties beyond 200 miles from a treatment plant as controls. For treatment counties, we include only those in the indicated distance ranges. For example, the coefficient at point 25 uses all counties within 25 miles of a Phase I plant as treatment counties, all counties beyond 200 miles as control counties, and omits all counties between 25 and 200 miles. We cluster standard errors at the county level.

Table 1
Means by Treatment and Control Counties in 1990

	Control	Treatment
Cause of Death		
Deaths per 100,000	495.94	478.37
Cardio per 100,000	190.70	177.87
Respiratory per 100,000	29.89	28.83
External per 100,000	64.91	53.25
Wage Employment, Income, and Population (REIS & SEER)		
Total Personal Income (\$1,000s)	16481.78	22183.06
Total Wage Employment	24382.53	31353.48
Population	54411.63	69551.64
Counties	796	1281

Notes: Overall internal deaths excludes deaths by AIDS due to changes in treatment technology introduced in 1995. ICD9 (up to 1997) and ICD10 (1998 and beyond) codes used to assign cause of death. Wage employment and personal income data are from the Bureau of Economic Analysis Regional Economic Accounts. Age-specific population is from SEER data.

Table 2
Difference in Adult Mortality Over Time

	(1)	(2)	(3)
Treat X 1990	0.0212** (0.009)	0.0014 (0.009)	-0.0023 (0.011)
Treat X 1991	0.0210*** (0.008)	0.0166** (0.008)	0.0013 (0.010)
Treat X 1992	0.0091 (0.007)	-0.011 (0.008)	0.0014 (0.009)
Treat X 1993	-0.0001 (0.007)	-0.0094 (0.008)	-0.0058 (0.010)
Treat X 1994	-	-	-
Treat X 1995	-0.0168** (0.008)	-0.0168** (0.008)	-0.0125 (0.010)
Treat X 1996	-0.0205** (0.008)	-0.0264** (0.010)	-0.0141 (0.010)
Treat X 1997	-0.0332*** (0.008)	-0.0446*** (0.010)	-0.0294*** (0.011)
Treat X 1998	-0.0326*** (0.010)	-0.0393*** (0.011)	-0.0217* (0.012)
Treat X 1999	-0.0302*** (0.011)	-0.0403*** (0.014)	-0.0262** (0.011)
Treat X 2000	-0.0478*** (0.012)	-0.0369*** (0.011)	-0.0286** (0.013)
Treat X 2001	-0.0471*** (0.013)	-0.0502*** (0.014)	-0.0237* (0.013)
Treat X 2002	-0.0513*** (0.015)	-0.0649*** (0.018)	-0.0420*** (0.015)
Treat X 2003	-0.0600*** (0.016)	-0.0559*** (0.017)	-0.0362** (0.014)
Treat X 2004	-0.0716*** (0.017)	-0.0692*** (0.017)	-0.0377** (0.016)
Treat X 2005	-0.0715*** (0.019)	-0.0781*** (0.022)	-0.0512*** (0.017)
Counties	2,077	2,077	2,077
Observations	33,232	33,232	33,232
County FEs	X	X	X
Year Fes	X	X	X
Weather		X	X
p-Score X Year			X

Standard errors in parentheses
* p<0.10, ** p<0.05, *** p<0.01

Notes: We cluster standard errors at the county level. Overall internal deaths excludes deaths by AIDS due to changes in treatment technology introduced in 1995. ICD9 (up to 1997) and ICD10 (1998 and beyond) codes used to assign cause of death. Coefficients show the regression-adjusted estimated difference between treatment and control counties, as compared to their estimated difference in 1994, the year before the ARP begins. p-score X Year indicates we include indicators for deciles of estimated propensity score interacted with year fixed effects.

Table 3
Difference in Adult Mortality Over Time: Variations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Age-Specific Pop.	SEER Age Share	Log of Deaths	Interacted Covariates	Do Not Trim	Trim More	Omit Plant Counties
Treat X 1990	-0.0004 (0.010)	-0.0043 (0.011)	-0.0025 (0.011)	0.0023 (0.009)	0.0049 (0.011)	-0.0083 (0.011)	-0.0039 (0.011)
Treat X 1991	-0.0003 (0.010)	0.0005 (0.010)	0.0012 (0.010)	0.0047 (0.008)	0.0026 (0.010)	-0.0048 (0.011)	0.0019 (0.010)
Treat X 1992	0.0063 (0.009)	0.002 (0.009)	0.0012 (0.009)	-0.0041 (0.008)	0.0076 (0.010)	-0.0039 (0.010)	0.0014 (0.010)
Treat X 1993	-0.0051 (0.010)	-0.0052 (0.010)	-0.0058 (0.010)	-0.007 (0.008)	-0.0024 (0.010)	-0.0069 (0.010)	-0.0077 (0.010)
Treat X 1994	-	-	-	-	-	-	-
Treat X 1995	-0.0113 (0.010)	-0.0101 (0.010)	-0.0126 (0.010)	-0.0122 (0.008)	-0.0126 (0.010)	-0.0176* (0.010)	-0.0155 (0.010)
Treat X 1996	-0.0099 (0.010)	-0.0118 (0.010)	-0.0142 (0.010)	-0.013 (0.008)	-0.0116 (0.010)	-0.0189* (0.010)	-0.0162 (0.010)
Treat X 1997	-0.0248** (0.011)	-0.0262** (0.011)	-0.0295*** (0.011)	-0.0271*** (0.008)	-0.0317*** (0.012)	-0.0330*** (0.011)	-0.0316*** (0.012)
Treat X 1998	-0.0179* (0.010)	-0.0204* (0.011)	-0.0220* (0.012)	-0.0169** (0.008)	-0.0293** (0.012)	-0.0223* (0.012)	-0.0245** (0.012)
Treat X 1999	-0.0182* (0.010)	-0.0223** (0.011)	-0.0272** (0.011)	-0.0205** (0.009)	-0.0287*** (0.011)	-0.0251** (0.011)	-0.0267** (0.012)
Treat X 2000	-0.0272*** (0.011)	-0.0262** (0.012)	-0.0286** (0.013)	-0.0267*** (0.009)	-0.0357*** (0.013)	-0.0312** (0.013)	-0.0327*** (0.013)
Treat X 2001	-0.0183* (0.011)	-0.0188 (0.012)	-0.0237* (0.013)	-0.0257*** (0.008)	-0.0255** (0.013)	-0.0234* (0.013)	-0.0282** (0.013)
Treat X 2002	-0.0342*** (0.012)	-0.0353*** (0.013)	-0.0427*** (0.015)	-0.0337*** (0.010)	-0.0468*** (0.015)	-0.0389*** (0.015)	-0.0455*** (0.015)
Treat X 2003	-0.0314*** (0.012)	-0.0316** (0.013)	-0.0367*** (0.014)	-0.0309*** (0.009)	-0.0402*** (0.014)	-0.0375*** (0.014)	-0.0419*** (0.014)
Treat X 2004	-0.0332*** (0.013)	-0.0326** (0.015)	-0.0376** (0.016)	-0.0378*** (0.010)	-0.0461*** (0.016)	-0.0389** (0.017)	-0.0429*** (0.016)
Treat X 2005	-0.0446*** (0.014)	-0.0430*** (0.015)	-0.0517*** (0.017)	-0.0407*** (0.010)	-0.0566*** (0.017)	-0.0478*** (0.016)	-0.0564*** (0.017)
Counties	2,077	2,077	2,077	2,077	2,414	1,632	1,989
Observations	33,232	33,232	33,212	33,232	38,624	26,112	31,824

Standard errors in parentheses
* p<0.10, ** p<0.05, *** p<0.01

Notes: We cluster standard errors at the county level. Overall internal deaths excludes deaths by AIDS due to changes in treatment technology introduced in 1995. ICD9 (up to 1997) and ICD10 (1998 and beyond) codes used to assign cause of death. Coefficients show the regression-adjusted estimated difference between treatment and control counties, as compared to their estimated difference in 1994, the year before the ARP begins. Each column shows a different regression, based off Column 3 of Table 2 — column headers indicate regression model adjustments.

Table 4
Annual Avoided Deaths in Treatment Counties


Year	(1) Infants	(2) 1-4	(3) 5-14	(4) 15-24	(5) 25-34	(6) 35-44	(7) 45-54	(8) 55-64	(9) 65-74	(10) 75-84	(11) 85 and up
1995	-829	-237	81	214	82	-68	-451	-1,401	-594	-4,790	-3,174
1996	254	-383	-133	18	89	444	-1,066	-1,340	-187	-5,600	-6,567
1997	-1,257	-166	-3	249	-277	-394	-1,441	-2,499	-2,252	-4,073	-8,655
1998	-2,393	-123	-267	207	-85	-403	-1,028	-1,724	-665	-5,045	-5,859
1999	-1,411	-284	-91	138	-327	31	-1,922	-1,452	-4,238	-6,166	-8,627
2000	-2,050	-165	47	301	-700	-121	-1,425	-2,194	-1,657	-5,552	-5,784
2001	-154	-57	-102	-30	-637	61	-371	-2,413	-3,649	-9,303	-7,332
2002	-919	-263	14	145	-554	-110	-1,563	-4,038	-5,485	-9,134	-7,606
2003	-971	-508	10	178	-763	160	-1,780	-2,693	-4,048	-9,022	-7,698
2004	-431	-291	154	27	-650	636	-1,890	-3,654	-5,429	-8,029	-5,548
2005	-1,777	-518	-110	-227	-147	-720	-2,130	-4,412	-7,995	-10,346	-7,663

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Notes: Table shows estimate of annual avoided mortality in treatment counties. We estimate changes as follows. First, we obtain annual percentage changes in mortality, as compared to 1994. We run the regression equation equivalent of our estimates in Column 3 of Table 2, but do so by indicated age groups. Second, we take county-level internal mortality levels in included treatment counties in 1994. Finally, we sum estimated differences across all treatment counties.

Figure A-1
p-score Distribution



Notes: Propensity scores by increments of 0.1. See Section 3.3 for details.

Table A-1
Additional Coefficients From Event Study Graphs

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Cardio	Resp.	Infants	Age 1-14	Age 15-34	Age 65+	Wage Emp.	Personal Income	Net Migration	External
Treat X 1990	-0.0212 (0.018)	-0.0346 (0.054)	-0.0426 (0.066)	-0.0627 (0.091)	-0.0005 (0.045)	-0.0002 (0.007)	0.0013 (0.006)	-0.01 (0.007)	-0.0035 (0.003)	0.003 (0.031)
Treat X 1991	-0.0003 (0.017)	-0.0025 (0.050)	-0.0209 (0.062)	-0.1367 (0.088)	-0.0152 (0.046)	0.0052 (0.006)	-0.0002 (0.005)	0.0005 (0.005)	-0.0025 (0.002)	0.0027 (0.031)
Treat X 1992	-0.0097 (0.016)	-0.0316 (0.054)	-0.0936 (0.066)	-0.2180** (0.099)	0.0201 (0.047)	-0.0142** (0.006)	-0.0083* (0.004)	-0.0097** (0.005)	-0.0022 (0.002)	-0.0071 (0.033)
Treat X 1993	-0.0084 (0.016)	-0.0443 (0.048)	-0.1398** (0.067)	-0.1829* (0.097)	0.0407 (0.050)	-0.0021 (0.005)	-0.0014 (0.003)	0.0002 (0.003)	0.0002 (0.001)	-0.0172 (0.030)
Treat X 1994	-	-	-	-	-	-	-	-	-	-
Treat X 1995	-0.0326** (0.016)	-0.0583 (0.049)	-0.0786 (0.073)	-0.1118 (0.097)	0.054 (0.052)	-0.0125*** (0.004)	-0.0017 (0.002)	-0.0005 (0.002)	-0.0008 (0.002)	-0.0019 (0.036)
Treat X 1996	-0.0207 (0.016)	-0.0333 (0.052)	0.0241 (0.068)	-0.2371*** (0.090)	0.0281 (0.045)	-0.0190*** (0.006)	-0.0051 (0.004)	-0.0053 (0.004)	-0.0016 (0.002)	-0.0366 (0.031)
Treat X 1997	-0.0331* (0.017)	-0.0738 (0.046)	-0.1191* (0.066)	-0.0589 (0.090)	0.0252 (0.051)	-0.0243*** (0.006)	-0.0065 (0.006)	-0.0045 (0.006)	-0.0025 (0.003)	0.0015 (0.033)
Treat X 1998	-0.0325* (0.017)	-0.028 (0.051)	-0.2269*** (0.070)	-0.1969** (0.100)	0.0364 (0.051)	-0.0200*** (0.007)	-0.0018 (0.007)	0.0031 (0.008)	-0.0025 (0.002)	-0.0027 (0.033)
Treat X 1999	-0.0303* (0.018)	-0.0185 (0.045)	-0.1338* (0.069)	-0.1612* (0.089)	0.0124 (0.048)	-0.0287*** (0.008)	-0.0028 (0.009)	0.0001 (0.010)	0.0009 (0.002)	-0.0086 (0.029)
Treat X 2000	-0.0369** (0.018)	-0.0939* (0.053)	-0.1944*** (0.073)	-0.0693 (0.095)	-0.0258 (0.051)	-0.0231** (0.009)	0.0052 (0.010)	0.0093 (0.010)	0.0007 (0.003)	0.0039 (0.031)
Treat X 2001	-0.0402** (0.018)	-0.0197 (0.047)	-0.0146 (0.073)	-0.0994 (0.095)	-0.0446 (0.049)	-0.0334*** (0.010)	-0.0039 (0.011)	0.004 (0.011)	-0.001 (0.003)	0.0151 (0.032)
Treat X 2002	-0.0431** (0.020)	-0.0204 (0.052)	-0.0871 (0.075)	-0.1088 (0.092)	-0.0353 (0.050)	-0.0379*** (0.011)	-0.007 (0.012)	0.0009 (0.013)	0.0009 (0.003)	0.0072 (0.034)
Treat X 2003	-0.0580*** (0.019)	-0.1197** (0.047)	-0.0921 (0.083)	-0.1660* (0.100)	-0.0348 (0.049)	-0.0346*** (0.012)	-0.0009 (0.013)	0.0047 (0.014)	0.0038 (0.004)	-0.0079 (0.036)
Treat X 2004	-0.0450* (0.025)	-0.0944* (0.050)	-0.0409 (0.082)	-0.0649 (0.103)	-0.0431 (0.047)	-0.0320** (0.013)	-0.0009 (0.014)	-0.0057 (0.018)	0.0045 (0.004)	-0.0065 (0.036)
Treat X 2005	-0.0710*** (0.023)	-0.0726 (0.048)	-0.1685** (0.078)	-0.2565** (0.105)	-0.0236 (0.049)	-0.0428*** (0.015)	-0.0023 (0.017)	-0.0233 (0.021)	-0.0059 (0.006)	-0.0761** (0.037)
Counties	2,077	2,077	2,077	2,077	2,077	2,077	2,077	2,077	2,077	2,077
Observations	33,232	33,232	33,232	33,232	33,232	33,232	33,232	33,232	33,231	33,232

Standard errors in parentheses
* p<0.10, ** p<0.05, *** p<0.01

Notes: All regressions follow column 3 of Table 2, with varied outcomes or samples specified in the table header, and correspond to event study graphs in Figures 4, 5, and 6.