The Long-run Effect of Air Pollution on Survival*

Tatyana Deryugina and Julian Reif
University of Illinois and NBER

July 2023

WORKING PAPER

Abstract

Many environmental hazards produce health effects that take years to arise, but quasi-experimental studies typically measure outcomes and treatment over short time periods. We develop a new approach to overcome this challenge and use it to gauge the effect of exposure to sulfur dioxide on US life expectancy. Using changes in wind direction as an instrument for daily air pollution levels, we first characterize the short-run mortality effects of acute exposure during the time period 1972–1988. Exposure causes two distinct mortality patterns: a short-run mortality displacement effect, and a persistent accelerated aging effect. We then incorporate our estimates into a health production model to quantify the lifelong effects of chronic air pollution exposure for a cohort born in 1972. Model calculations of the effect of chronic exposure on life expectancy are 7–8 times larger than the effect implied by simple extrapolation of our short-run empirical estimates. Ninety percent of the survival benefits accrue after the first fifty years of life, implying that most of the 1970 Clean Air Act's health benefits have yet to emerge for cohorts born after its passage.

JEL Classification: I18, Q52, Q53
Keywords: Air pollution, chronic exposure, mortality, survival

*We thank Josh Gottlieb, Peter Hull, Nicole Riemer, Hannes Schwandt, Matthew West, and Nikos Zirogiannis for helpful feedback. Prakrati Thakur, Lexi Chen, and Xinhui Sun provided excellent research assistance. We are grateful to seminar participants at Arizona State University, Baltic Economic Association, Cornell University, DePaul University, Exeter University, GSSI School of Advanced Studies, Iowa State University, ITAM, IZA Workshop on Environment and Labor Markets, Jinan University, Johns Hopkins University, London School of Economics, MHEC, Nova SBE, Oxford University, Santa Clara University, Southern Economic Association meetings, Southern Illinois University, Stanford University, Tinbergen Institute, Toulouse School of Economics, University of Birmingham, UC Berkeley, University of Chicago, University of Duisburg-Essen, University of Geneva, UIC, University of Indiana, University of Mannheim, University of Navarra, Wake Forest University, and Wharton School for feedback.

1 Introduction

Environmental hazards such as air pollution, extreme temperatures, and water pollution are important causes of human morbidity and mortality. For example, The Lancet Commission on Pollution and Health estimates that air pollution caused 6.5 million premature deaths in 2015, amounting to about 12 percent of all deaths worldwide (Landrigan et al., 2018). Such assessments are generally based on observational studies, which are prone to estimation bias.[1] Although quasi-experimental studies can address this bias, they typically measure health outcomes and treatment exposure over short time periods that span one year or less, thereby overlooking effects that may emerge years or decades later. This limitation is challenging to overcome because quasi-experimental variation in chronic (long-run) treatment exposure is difficult to find, data that track individuals over long time periods are rare, and endogenous responses such as migration complicate the interpretation of estimates. However, optimal health and environmental policies require an accurate estimate of the lifelong health consequences of permanent changes in exposure.

We propose a novel approach to address this limitation when the outcome is mortality. The approach combines well-identified short-run estimates with a versatile model of survival that can accommodate a number of different mortality dynamics (Lleras-Muney and Moreau, 2022). Once calibrated to match our empirical estimates, the model can fully characterize the short- and long-run survival effects of both acute and chronic changes in treatment exposure. The calibrated model can be validated by comparing its short-run predictions to empirical estimates of outcomes that are measured over different time windows or using different ages than those used for calibration.

We use our approach to estimate the short- and long-run effects of exposure to air pollution on US population mortality. We assemble a new, comprehensive dataset that combines 17 years of administrative death records covering the period 1972–1988 with daily data on

[1] See Englert (2004) and Dominici, Greenstone and Sunstein (2014) for a review and discussion of observational studies of air pollution.

air pollution and weather. We focus on the mortality effects of sulfur dioxide (SO_2), the most widely measured air pollutant during this time period, and incorporate other air pollutants such as particulate matter into a secondary analysis that uses a subset of our sample. In the first part of our paper, we investigate the causal effect of acute (1-day) air pollution exposure on county-level mortality by instrumenting for observed changes in SO_2 with changes in local wind direction. We estimate that a 1-unit ($\approx$10 percent) increase in SO_2 raises 1-day mortality by 0.08 deaths per million (0.33 percent).

The cumulative mortality effect of this 1-day exposure more than doubles when we extend the outcome window from one day to one month (28 days), demonstrating that air pollution continues to have lethal effects weeks after exposure has ended. While statistical power limits our quasi-experimental estimates to a one-month outcome window, the trend suggests the cumulative mortality effect would continue increasing beyond this window, albeit at a decreasing rate. These results are consistent with findings from the medical literature suggesting that air pollution causes "accelerated aging" by, for example, hardening arteries and increasing the risk of heart disease (Rajagopalan and Landrigan, 2021).

We document striking differences in mortality dynamics across different causes of death. Our 1-day mortality estimate is driven roughly equally by deaths related to three groups of causes: cardiovascular disease, cancer, and "other diseases," a residual category that includes chronic lower respiratory illness and diabetes. Lengthening the outcome window to one month, however, causes the estimates for cardiovascular and other diseases to increase by a factor of 3–4, while the estimate for cancer falls by 70 percent and becomes statistically insignificant. This fall in the magnitude of the cancer-related mortality estimate over time implies that the cancer-related deaths caused by acute exposure were deaths of frail individuals who had short (less than one month) counterfactual life expectancies—a phenomenon often referred to as "mortality displacement." Altogether, we conclude that acute exposure to air pollution produces two distinct mortality patterns: mortality displacement in a subpopulation of frail individuals, where the cumulative mortality effect quickly dissipates, and

accelerated aging in a subpopulation of healthier individuals, where the cumulative effect grows slowly over time. On net, the accelerated aging effect dominates.

Our short-run analysis has two important limitations common to quasi-experimental studies of air pollution: treatment is limited to acute (1-day) exposure, and the outcome window is limited to one month following exposure. In the second part of our paper, we quantify the effect of chronic (lifetime) exposure on long-run survival by adapting a dynamic production model of health to our daily mortality setting (Lleras-Muney and Moreau, 2022). While there are many ways to model survival, this model is particularly well-suited to our needs because it can accommodate both the accelerated aging and mortality displacement patterns present in our setting. We first calibrate the model's baseline parameters using a 1972 US life table. We then use our 1-day cancer mortality estimates to calibrate the effect of pollution exposure on the model parameter governing mortality displacement and use the 1-day non-cancer mortality estimates to calibrate the effect of exposure on the aging parameter.

To validate the model, we compare its age-specific mortality predictions in the month following acute exposure to IV estimates not used for calibration. The overwhelming majority of model predictions lie inside the 95% confidence intervals of the IV estimates. These predictions depend meaningfully on the estimated fraction of deaths that are due to mortality displacement: alternatively assuming that acute exposure produces either 0% or 100% mortality displacement yields predictions that lie far outside the 95% confidence intervals of the IV estimates. These results demonstrate that while all-cause mortality estimates are insufficient for drawing reliable conclusions about longer-term survival, researchers can overcome this problem by appropriately incorporating information about cause of death into a survival model.

We then use our model to quantify the effect of a chronic (permanent), 1-ppb decrease in SO_2 exposure on life expectancy. The model predicts that, all else equal, such a decrease improves life expectancy at birth by 1.2–1.3 years, which is 7–8 times larger than a naive

estimate that extrapolates our acute exposure IV estimates of the monthly mortality effect to the whole life cycle. Although the decrease in chronic exposure begins at birth, ninety percent of the improvements in life expectancy occur after age 50 and over three-quarters occur after age 65, suggesting that the vast majority of the survival benefits of the drastic reductions in air pollution in the United States over the past fifty years have yet to emerge for cohorts born after the passage of the 1970 Clean Air Act. Because the model holds behavior fixed, its forecasts can be interpreted as the gross benefits associated with pollution reduction, uncontaminated by longer-run behavioral responses such as migration (Graff Zivin and Neidell, 2012; Currie et al., 2014).

The main contribution of our study is the development and application of a new framework for estimating the long-run survival effects of chronic exposure to environmental hazards. The standard approach focuses on acute exposure and quantifies long-run survival effects using population life tables (e.g., Deschênes and Greenstone, 2011), and recent work has improved the accuracy of this method by incorporating individual-level predictions of counterfactual life expectancy (Deryugina et al., 2019). However, that approach remains prone to bias if unobserved characteristics are correlated with both life expectancy and the probability of dying from exposure, and it cannot quantify the effects of chronic exposure. Our study takes a different approach, combining short-run estimates of acute exposure with a dynamic model of health production. This method leverages demographic knowledge about life-cycle mortality patterns to project long-run survival effects and is similar in spirit to Athey, Chetty and Imbens (2020), who combine experimental and observational data to infer the long-run effects of a short-run experimental treatment. The approach complements efforts to estimate the long-run survival effects of pollution directly. Direct estimation requires identifying quasi-experimental variation in long-run exposure and accounting for behavioral responses, such as migration. Ignoring these responses leads to an underestimate of the health costs of air pollution (Graff Zivin and Neidell, 2012; Currie et al., 2014), and is a possibility that cannot be ruled out among the few studies measuring mortality over mul-

tiple years (Chen et al., 2013; Ebenstein et al., 2017; Barreca, Neidell and Sanders, 2021).[2] By contrast, our approach employs quasi-experimental variation at the daily level, which alleviates concerns about many forms of avoidance behavior.

Previous studies of temporary pollution spikes such as London's Great Smog of 1952 have found that mortality effects can linger for weeks after the pollution event ends (Logan, 1953). Our study shows that this lingering effect persists even at moderate levels of pollution, and to our knowledge we are the first to show that pollution exposure causes two distinct mortality patterns: mortality displacement, where cumulative mortality quickly falls to zero, and accelerated aging, where the mortality effect grows with time. Our findings underscore that great care must be taken when inferring longer-run mortality effects from short-run estimates. Although lengthening the outcome window to multiple years helps address this challenge, studies with short outcome windows of one year or less make up the vast majority of papers employing quasi-experimental variation in air pollution.[3] Our methodology of incorporating detailed short-run mortality estimates into a calibrated survival model provides a way to form long-run forecasts that appropriately accounts for both mortality displacement and accelerated aging effects.

Our short-run analysis also makes important contributions to the literature on the health effects of acute exposure to air pollution. A key advantage of our empirical approach is that it harnesses variation across a large geographic area (the continental United States) and over a long time period (nearly two decades) for the majority of the US population. Our paper is thus the largest quasi-experimental study of acute pollution exposure on mortality, encompassing 18 million deaths and enabling us to decompose mortality effects with precision.[4]

[2]For example, Barreca, Neidell and Sanders (2021) utilize repeated cross-section mortality data at the county-year level, and are thus unable to control for any migration response to changes in pollution levels.

[3]See, for example, Currie and Neidell (2005); Knittel, Miller and Sanders (2016); Schlenker and Walker (2016); Deschênes, Greenstone and Shapiro (2017); Hollingsworth, Konisky and Zirogiannis (2021); Hollingsworth and Rudik (2021); and Heo, Ito and Kotamarthi (2023). A handful of papers study the effect of early-life air pollution exposure on later-life outcomes (e.g., Isen, Rossin-Slater and Walker, 2017; Voorheis, 2017; Colmer and Voorheis, 2020), including one study that considers mortality before age 55 (Arenberg and Neller, 2023). These studies do not consider chronic exposure, however.

[4]Deryugina et al. (2019) use a similar empirical approach, but focus on fine particulate matter ($PM_{2.5}$) rather than SO_2, consider a shorter and more recent time period (1999–2013), and are limited to a sample

Finally, ours is one of the first quasi-experimental studies to estimate the mortality effect of acute exposure to SO_2, a pollutant that has declined substantially in the US over the past fifty years but which remains elevated in lower-income countries such as China and India.

The rest of the paper is organized as follows. Section 2 provides background on air pollution and describes our data. Section 3 describes our short-run empirical strategy. Section 4 presents estimates of the short-run mortality effects of acute air pollution exposure. Section 5 presents the dynamic production of health, calibrates it, and quantifies the long-run survival effects of chronic exposure. Section 6 concludes.

2 Background and data

2.1 Air pollution

Sulfur dioxide can harm human health through two main channels. First, direct exposure to SO_2 in controlled clinical trials impairs respiratory function, especially in people with asthma (Agency for Toxic Substances and Disease Registry, 1998). Animal experiments have also demonstrated that SO_2 inhalation can cause brain damage (Sang et al., 2010; Yao et al., 2015) and contribute to cardiac and mitochondrial dysfunction (Qin et al., 2016). Second, SO_2 transforms rapidly into sulfate, on the order of several percent per hour (Luria et al., 2001). Sulfates are a primary component of fine particulate matter ($PM_{2.5}$), a catch-all term for particles whose diameter is 2.5 micrometers (μm) or less. $PM_{2.5}$ is thought to be particularly harmful to health because of its ability to cross the blood-alveolar and blood-brain barriers. Prior quasi-experimental research has found causal links between short-run exposure to $PM_{2.5}$ and a number of adverse health outcomes, such as short-run elderly mortality and hospitalizations (Deryugina et al., 2019).

While research on the exact pathophysiological mechanisms underlying these health effects continues, medical studies have documented significant associations between air pollu-

population over age 65.

tion and hypertension, diabetes, coronary artery calcification, and the progression of chronic kidney disease, all of which are risk factors for cardiovascular disease (Rajagopalan and Landrigan, 2021). Air pollution has also been linked to the initiation, promotion, and progression phases of lung cancer (Turner et al., 2020; Hill et al., 2023). Once initiated, lung cancer typically grows for over 10 years before it is diagnosed (Nadler and Zurbenko, 2014). Thus, while the quasi-experimental literature has consistently found significant adverse effects of exposure on short-run health, long-run health effects are likely to be even larger.

We measure air pollution using the EPA's Air Quality System database, which provides hourly data at the pollution monitor level for criteria pollutants regulated by the EPA. The extent of spatial and temporal coverage varies by pollutant. Our analysis focuses on sulfur dioxide (SO_2), the most widely monitored air pollutant during our time frame. In robustness checks, we also examine four other air pollutants that have been monitored since the 1970s or the 1980s: nitrogen dioxide (NO_2), total suspended particulates (TSP), ozone (O_3), and carbon monoxide (CO). TSP comprises all particulates with diameters less than 100 μm, thus including $PM_{2.5}$. Since $PM_{2.5}$ was not consistently monitored until the late 1990s, we cannot include it directly in our analysis.

Figure A.2 displays the population-weighted concentrations and the number of monitored counties over time for each pollutant. Except for ozone, the population-weighted mean for all pollutants declines substantially during our sample period. CO data are readily available since the mid-1970's and maintain consistent coverage of approximately 225 counties per year, while ozone data are unavailable prior to 1980. Data on SO_2 and NO_2 are available for a larger number of counties than CO, although this spatial coverage declines beginning in 1976. Each year during our sample period, at least 400 counties monitor sulfur dioxide concentrations, and about 50 percent of US individuals live in a county that monitors SO_2.

Panel A of Table 1 shows county-level summary statistics for daily ambient pollution concentrations during 1972–1988. The average SO_2 concentration during our sample period is 9.0 parts per billion, with a standard deviation of 12.6. The average levels of nitrogen

dioxide and ozone are higher, at about 21 and 26 parts per billion, respectively. The most prevalent pollutant is carbon monoxide, with an average concentration of 1.64 parts per million (1,640 parts per billion). We are more than twice as likely to observe SO_2 levels than any of the other four pollutants, which is one of the reasons we focus our analysis on SO_2.

2.2 Mortality

We obtain death counts from the National Vital Statistics. These data are based on death certificate records and include information on the cause of death and the county of occurrence. We focus our analysis on the years 1972–1988 because those years include information on the exact date of death.[5] We calculate death rates by dividing death counts by annual population estimates provided by the Surveillance, Epidemiology, and End Results (SEER) Program.

Figure A.3 reports death rates by age group and cause of death during our sample time period 1972–1988. The infant mortality rate steadily declines over time, and nearly equals the average death rate of the population by the end of our sample period. Panel B of Table 1 summarizes daily mortality rates for various subgroups over this time period. The all-age daily death rate is about 25 per million. The rate is higher for infants (33 deaths per million), and much higher for those over age 85 (443 deaths per million).

We classify causes of death into four categories: cardiovascular, cancer, external, and other. Panel B of Table 1 reports that cardiovascular disease is the leading cause of death during our time period, accounting for nearly half of all deaths in our sample (12 daily deaths per million). Cancer deaths make up slightly more than twenty percent of overall mortality (5 deaths per million). External causes of deaths are responsible for about eight percent of all deaths and include car accidents, poisonings, suicides, and other causes not originating in the body. We group the remaining twenty percent of deaths into the "other" category. The two largest components of this category are chronic lower respiratory illnesses and diabetes.

[5]The exact date of death is unavailable prior to 1972, and is available after 1988 only in the Research Data Center of the National Center of Health Statistics.

2.3 Wind and weather

Our empirical strategy is motivated by the observation that wind currents can carry air pollution over long distances. For example, Figure A.4 shows the fraction of $PM_{2.5}$ and sulfates—for which SO_2 is a precursor—that can be attributed to local versus regional sources for thirteen large US cities, as calculated by Environmental Protection Agency (2004). Regional contributions to air pollution substantially exceed local contributions in nearly every case.

It is theoretically possible to explicitly model the wind transport of pollutants and then use the model's prediction as an instrument for changes in local pollution levels. But doing so is computationally intensive at the daily level and requires comprehensive data on emissions, which are largely not available during our study period. Instead, we follow Deryugina et al. (2019) and instrument for changes in SO_2 using changes in wind direction. Our key identifying assumption is that, conditional on other climatic variables and comprehensive fixed effects, wind direction affects mortality only through its effects on air pollution.

We obtain wind speed and wind direction data from a 6-hour reanalysis dataset published by the Japan Meteorological Agency (JMA).[6] These data are available starting in 1958. They consist of vector pairs, one for the East-West wind direction (u-component) and one for the North-South wind direction (v-component), reported in 6-hour intervals on a grid with a resolution of 1.25 degrees (≈86 miles). We first interpolate between the grid points to calculate the 6-hour u- and v-components at the centroid of each county. We then average the u- and v- components within a county-day to match the frequency of our mortality data. Finally, we use trigonometry to convert the average u- and v- components into daily wind direction and wind speed.

We obtain daily temperature and precipitation data from Schlenker and Roberts (2009). Combining gridded monthly climate data from the PRISM Climate Group with daily data from weather stations, Schlenker and Roberts (2009) construct a gridded weather dataset

[6] Available from http://rda.ucar.edu/datasets/ds628.0/.

at the daily level. The final dataset spans the years 1972–1988 and includes total daily precipitation and daily maximum and daily minimum temperatures for each point on a 2.5-by-2.5 mile grid covering the contiguous US.[7]

3 Empirical strategy

Our first objective is to estimate the causal effect of acute (1-day) exposure to sulfur dioxide on short-run mortality. We model this relationship using the following regression:

$$Y_{cd}^{k} = \beta^{k}\text{SO2}_{cd} + X_{cd}^{k\prime}\delta + \alpha_{cm} + \alpha_{my} + \varepsilon_{cd} \tag{1}$$

where Y_{cd}^{k} is cumulative mortality rate in county c in the k days following exposure on day d (including same-day mortality). The parameter of interest is β^{k}, the coefficient on daily SO_2 levels. Similar to Deryugina et al. (2019), the controls X_{cd}^{k} include contemporaneous and $k-1$ leads of our weather variables (described below) in order to ensure that β^{k} is not capturing the effects of weather conditions during the outcome window. To minimize concerns about autocorrelation, we also control for two leads and two lags of the instruments.

Equation (1) includes county-by-month (α_{cm}) and month-by-year (α_{my}) fixed effects. The county-by-month fixed effects control for geographic and seasonal differences in mortality and air pollution. The month-by-year fixed effects control for common time-varying shocks, such as those induced by environmental policy changes during our study period. Standard errors are clustered by county, and the regression is weighted by the relevant county-year population.

Our main specification controls for daily maximum temperature, precipitation, and wind speed. We control for maximum temperature using a set of indicators based on county-specific percentiles in our sample. The literature on temperature and mortality generally

[7] See http://www.prism.oregonstate.edu/ for the original PRISM dataset and http://www.wolfram-schlenker.com/ for the daily data.

finds that it is extreme temperatures that matter for the mortality rate (e.g. Barreca et al., 2016) and that the average climate of a county determines which temperatures are extreme in this sense (e.g. Heutel, Miller and Molitor, 2021). We therefore generate nine indicators for maximum temperatures falling into intervals whose minimum is defined by the following county-specific percentile cutoffs: 0, 1, 5, 10, 25, 75, 90, 95, and 99.

Guided by the principle of controlling for extremes, we control for daily precipitation by including four indicators for whether precipitation is below the 75th percentile for that county (which in most counties corresponds to very little or no precipitation), between the 75th and 95th percentiles, between the 95th and 99th percentiles, or above the 99th percentile. We control for daily average wind speed with six indicators whose minimum is defined by the following county-specific percentile cutoffs: 0, 25, 75, 90, 95, and 99. Finally, we also control for the set of all possible interactions of these atmospheric controls, yielding 203 different temperature-precipitation-wind-speed combinations.[8]

OLS estimates of Equation (1) are susceptible to bias because exposure to SO_2 may not be randomly assigned and is measured with error. We therefore instrument for SO_2 using daily wind direction in the county, allowing the effect of wind on SO_2 to vary by geographic group g. The regression specification for our first stage is:

$$\text{SO2}_{cd} = \sum_{g=1}^{50} f^g(\theta_{cd}) + X_{cd}^{k\,\prime}\delta + \alpha_{cm} + \alpha_{my} + \varepsilon_{cd} \tag{2}$$

where:

$$f^g(\theta_{cd}) = \gamma_g^1 1[G_c = g] \times \sin(\theta_{cd}) + \gamma_g^2 1[G_c = g] \times \sin(\theta_{cd}/2)$$

The indicator function $1[G_c = g]$ is equal to 1 if county c is a member of group g and 0 otherwise. The variable θ_{cd} is the local wind direction, measured in radians. The excluded instruments are the 100 regressors formed by the interaction of group indicators, $1[G_c = g]$ with $\sin(\theta_{cd})$ and $\sin(\theta_{cd})/2$. As we demonstrate in Section 4.3, our results are robust to

[8]The number of combinations is theoretically higher, but we do not observe all of them in our data.

alternative ways of parameterizing $f^g(\theta_{cd})$.

Equation (2) requires the relationship between wind direction and air pollution to be constant within the geographic group g, which encompasses multiple counties. We define these groups using a k-means clustering algorithm that classifies all SO_2 pollution monitors into 50 spatial groups based on monitor location.[9]

3.1 Discussion

Air pollution sources can be local or distant. For example, air pollution exposure varies with the amount of vehicle exhaust emitted by nearby automobiles and with the amount of smoke emitted by fires and power plants located many hundreds of miles away. We designed the first stage, Equation (2), to exploit variation in pollution emitted by distant sources and to ignore variation emitted by local (within-county) sources. While air pollution emitted by distant sources will have an approximately uniform effect on the pollution levels of counties far from those sources, pollution from local sources within the county will disperse unevenly, affecting some parts of the county and not affecting other parts. Because pollution is sparsely monitored at the county level, it is thus important to avoid employing local variation in pollution transport when estimating the effects of average county-level pollution exposure on health. Failing to do so may result in significant measurement error, which will produce bias in estimated effects of air pollution.[10]

Figure 1 illustrates the variation we use to estimate the causal effects of acute sulfur dioxide exposure, using the Greater Philadelphia and Southern California areas as examples. The black dots on the maps show the locations of the SO_2 monitors in these two

[9]The inputs into the k-means clustering algorithm are monitor latitude and longitude and the desired number of groups (50). If monitors in the same county are assigned to more than one group, we assign the larger integer group number to the county, which is effectively random assignment. The locations of in-sample SO_2 monitors and monitor groups are displayed in Figure A.1. On average, each monitor group encompasses 57 monitors with SO_2 readings and 12 counties.

[10]Consider a county with one pollution monitor that is placed due west of a major power plant. Suppose that this power plant is the only source of pollution and is located in the center of the county. The monitor will register high pollution readings when the wind blows from the east, and low levels when it blows from the west. However, wind direction in this example actually has no effect on average pollution levels in the county.

geographic groups. The plots on the right show the group-specific relationships between daily average wind direction and SO_2.[11] Each specification controls for maximum temperature, precipitation, wind speed, and county-by-month and month-by-year fixed effects in the same way we control for them in our main specification.

Figure 1 reveals a strong first-stage relationship between wind direction and SO_2 levels. In the Greater Philadelphia area, pollution levels are highest when the wind blows from the west-southwest direction, and lowest when the wind is blowing from the east-southeast direction, where the Atlantic Ocean lies. This pattern differs in the Southern California area, where SO_2 levels are highest when the wind is blowing from the east, a densely populated area, and lowest when the wind is blowing from the south-southwest direction, where the Pacific Ocean lies. Figure 1 shows that a change in wind direction can increase SO_2 levels by 3–4 ppb, equal to 30–40 percent of the national mean during this time period (Table 1).

Figure A.5 shows the geographic distribution of the strength of the first stage, as measured by the difference in SO_2 levels between the most and least polluting wind directions.[12] The areas with the largest variation (4+ ppb) are located primarily in the Midwest and the Northeast, although there are fairly strong compliers (variation of 2–4 ppb) throughout the country. The weakest compliers have variation of <1 ppb and thus contribute relatively little to our main estimates.

In Table A.1, we test for differences among these complier areas by regressing county characteristics obtained from the Regional Economic Information System (REIS) dataset on the size of the first stage. The unit of observation is a county-year, and we control for year fixed effects. The size of the first stage is not significantly correlated with population or the percent of residents over age 65. Counties with stronger first stages have a smaller share of residents who are Black, have higher per-capita income, receive higher per-capita transfers, and have a lower employment rate, but the magnitudes of these coefficients are in most cases

[11] Figure A.11 shows corresponding plots for all 50 geographic groups.

[12] For each monitor group g, we calculate $\widehat{\gamma_g^1}\sin(\theta) + \widehat{\gamma_g^2}\sin(\theta/2)$ for $\theta \in [0, 2\pi)$ and take the difference between the maximum and minimum values.

economically small (e.g., a $90 difference in per-capita income for each 1 ppb difference in the first stage). Overall, the complier group appears to be fairly representative of the US population.

Our empirical approach permits us to instrument for multiple pollutants simultaneously because the relationship between wind direction and pollution levels within and across regions is pollution-specific.[13] In a later analysis, we investigate the sensitivity of our main estimate to controlling for the four other pollutants measured during our sample time period: nitrogen dioxide, carbon monoxide, ozone, and total suspended particulates. Including these additional controls causes our sample size to fall by 90 percent or more because these other pollutants are monitored much less frequently than SO_2 during this time period (Table 1). We therefore do not include them in our primary specification.

4 Short-run empirical results

4.1 Mortality by age and cause

Table 2 presents OLS and IV estimates of Equation (1). Column (1) reports that a 1-ppb increase in daily SO_2 is associated with a same-day mortality increase of 0.008 deaths per million, about ten times smaller than the corresponding IV estimate of 0.08 deaths per million reported in Column (2). This downward bias for OLS is common in quasi-experimental studies of air pollution, and is often hypothesized to be at least partly due to measurement error in pollution exposure (Deryugina et al., 2019; Alexander and Schwandt, 2022).

Figure 2 shows the IV estimates of the effects of a 1-day, 1-ppb increase in SO_2 levels on mortality in the month following exposure. The first blue point depicts the estimate from Column (2) of Table 2. If short-term mortality displacement were the predominant driver of this 1-day mortality effect, we would expect the estimated mortality effects to decline

[13] For example, automobile emissions are a major contributor to ozone, but not to SO_2.

over time, potentially all the way to zero. Instead, the estimated effect increases steadily to 0.10 deaths per million in the week following exposure and to 0.18 deaths per million in the month following exposure.

Figure 3 shows how the 1-day mortality effect varies by age group. Panel (a) shows the absolute magnitude of the effect (in deaths per million), while panel (b) reports it as a percent of the average 1-day mortality for that age group. We fail to detect significant mortality increases for the two youngest age groups (covering ages 0–19). For older ages, our estimates are statistically significant and range from 0.017 deaths per million for 20–44-year-olds to 2.3 deaths per million for 85+ year-olds (Figure 3a). When expressed in relative terms, estimates vary little with age. The smallest relative effect (0.16 percent of daily mortality) is found among 55-59 year olds, while the largest relative effect (0.51 percent of daily mortality) is found among 85+ year olds. However, we cannot reject that the relative effects are equal to each other for most age group pairs.

In relative terms, our estimates for older age groups are similar to corresponding estimates reported in Deryugina et al. (2019), who investigate the effect of acute exposure to fine particulate matter on 3-day mortality. To the best of our knowledge, no prior quasi-experimental estimates exist for the effect of air pollution on short-run mortality for ages 1–64. For infant mortality, the two most comparable studies are Currie and Neidell (2005) and Knittel, Miller and Sanders (2016), who estimate the effect of PM10 (coarse particulate matter) on weekly infant mortality. Our statistically insignificant infant mortality estimate aligns with estimates from Currie and Neidell (2005) but is smaller than those from Knittel, Miller and Sanders (2016).

Table A.4 reports estimates by age for different outcome windows. For ages 60 and over, the monthly estimates are at least twice as large as the 1-day estimates, suggesting that any short-run mortality displacement among these groups of individuals are more than offset by delayed effects of acute exposure. For those between ages 20 and 59, however, the monthly estimates are much smaller than the 1-day estimates and are statistically insignifi-

cant, suggesting that acute pollution exposure among young and middle-aged adults has an immediate mortality impact only on those individuals who are very frail. By contrast, older adults who are killed by acute air pollution exposure have longer counterfactual lifespans.

Figure 4 shows how our estimates vary by cause of death as reported on death certificates. The increase in 1-day mortality is split roughly equally among cardiovascular disease, cancer, and "other" diseases. As we consider longer time horizons, however, the estimated effect on cancer-related mortality falls, implying that these individuals were already quite ill and would have died anyway in the days or weeks that followed exposure. After one month, the cancer estimate becomes small and statistically insignificant, which means that we cannot reject the null hypothesis that all cancer-related deaths would have occurred within one month even absent the pollution shock. By contrast, deaths from cardiovascular and other diseases increase with the time horizon—more than doubling over one month—implying that acute SO_2 exposure continues to have lethal effects even after exposure has ended.

We find a small but statistically significant increase in same-day external deaths (0.005 deaths per million), but this estimates becomes statistically insignificant when measured over longer time horizons. While we lack detailed data to probe this result further, it could be due to negative effects of pollution on cognitive function (Crüts et al., 2008; Fonken et al., 2011; Bishop, Ketcham and Kuminoff, 2023); for example, recent work has suggested this pollution-induced cognitive effect causes vehicle fatalities (Burton and Roach, 2023).

Figure A.12 shows estimates of the effect of SO_2 over time for 30 different causes of death. All subcategories of cardiovascular disease except for "other" and hypertension show significant same-day and monthly mortality increases. We also find both simultaneous and lagged effects for deaths from chronic obstructive pulmonary disease (COPD), pneumonia, and influenza, all of which are plausibly related to air pollution exposure. By contrast, we find no significant monthly effects for most conditions that have not been linked to air pollution exposure, such as stomach ulcers, cholelithiasis, appendicitis, and chronic liver

diseases.[14] These patterns further support our identification strategy.

4.2 Other air pollutants

So far we have interpreted our estimates as the causal effects of SO_2 on mortality. It is possible, however, that other harmful air pollutants are co-transported with SO_2. A second, related issue is that SO_2 transforms rapidly into sulfate at a rate of several percent per hour (see Section 2.1), raising the possibility that it is exposure to this component of $PM_{2.5}$ rather than its precursor, SO_2, that drives our mortality estimate.

Because different air pollutants are produced by a variety of different sources in different geographic locations, are carried differently by the wind, and exhibit different atmospheric chemistry patterns, our empirical strategy allows us to instrument separately for multiple pollutants in a single regression. The main challenge is that other pollutants are not consistently measured during our time period, and including all of them in our regression significantly reduces our sample size. We therefore create two different subsamples. The first one is the smallest and includes measures of TSP, NO_2, ozone, and CO in addition to SO_2. The second is larger but includes only SO_2 and TSP (our best proxy for $PM_{2.5}$).

Table 3 reports estimates from these two subsamples, using 1-day all-cause mortality as the outcome. Column (1) in Panel A reports that a 1-ppb increase in SO_2 increases 1-day mortality by 0.084 deaths per million in this sample. This coefficient's 95% confidence interval, $[0.0614, 0.107]$, includes all the other estimates except for two presented in Columns (2) and (5), which fall just outside of the interval. The estimated mortality effects of nitrogen dioxide and carbon monoxide are never significant. The coefficient on ozone is negative and marginally significant, possibly because ozone is negatively correlated with other, unmonitored pollutants that matter for mortality.[15] The estimated effect of SO_2, however, is

[14] Interestingly, we find a lagged effect on deaths from meningitis. Although not obviously pollution-related, there is suggestive evidence that air pollution can negatively affect the local immunity of the pharynx, leading to increased susceptibility to meningitis (e.g., Michele et al., 2006; Jusot et al., 2017; Shears et al., 2020).

[15] Ground-level ozone forms when oxides of nitrogen NO_2 react with volatile organic compounds (VOCs) on hot summer days. Many different chemicals are classified as VOCs, and car exhaust is a major source of

insensitive to the inclusion of ozone, suggesting that these unobserved pollutants are not biasing our SO_2 estimates. The inclusion of TSP possibly affects the SO_2 coefficient, reducing it by 20-30 percent, consistent with the notion that the SO_2 measured by pollution monitors is accompanied by sulfates, which are not measured directly but included in measures of TSP.

Panel B shows the effect of controlling for TSP in a larger sample that only conditions on observing both SO_2 and TSP. Comparing Column (2) to Column (1) shows that the coefficient on SO_2 is 55 percent lower after controlling for TSP.[16] This difference is statistically significant, suggesting that some of the SO_2 mortality effect we estimate is driven by particulate matter that was either co-transported with or formed by SO_2.

To investigate this possibility more directly, we use simulations from the Intervention Model for Air Pollution (InMAP), which models how emissions of SO_2, $PM_{2.5}$, and NO_2 are transported across the US.[17] Importantly, InMAP simulates the conversion of SO_2 and NO_2 into particulate matter, allowing us to gauge both the co-transport of SO_2 and $PM_{2.5}$ and the importance of SO_2 conversion. We use the 1990 National Emissions Inventory to determine emissions, the earliest year for which appropriate data are available.[18] We focus on emissions from coal-fired power plants, the largest source of SO_2 emissions in the NEI, although we also report results from all-source emissions. Given emissions information, InMAP outputs a spatial distribution of: (1) ambient SO_2; (2) primary $PM_{2.5}$ (fine particulate matter that is directly emitted by the source); and (3) secondary pollutants (e.g., sulfate created as a result of atmospheric conversion of SO_2, and $PM_{2.5}$ derived from NO_2 transport).

---

them. NO_x emissions come from a variety of vehicles and industrial sources. Higher ozone levels may mean there are fewer precursor chemicals in the ambient air, although testing this explicitly is beyond the scope of this paper.

[16] For completeness, Table A.6 shows results when restricting the sample to observations that have pollution readings for SO_2, NO_2, O_3, and CO, which yields almost four times the number of observations as Panel A of Table 3. The ozone coefficients cease to be statistically significant, and the estimated effects of SO_2 remain stable across the different sets of included pollutants.

[17] Available from https://github.com/spatialmodel/inmap/releases/tag/v1.9.6. Evaluation data for the simulations are from Tessum et al. (2019).

[18] Available from https://gaftp.epa.gov/air/nei/nei_criteria_summaries/1990criteriasummaryfiles/.

Using the results of the InMAP simulation, we calculate the ratio of total $PM_{2.5}$ for each part per billion of ambient SO_2.[19] Using coal-fired power plants emissions only, the ratio is approximately 2.3, indicating that each transported ppb of SO_2 is accompanied by 2.3 $\mu g/m^3$ of $PM_{2.5}$. The vast majority (>99%) of this $PM_{2.5}$ consists of secondary pollutants. The all-source ratio is slightly higher (approximately 3.0). Because pollution sources that are not coal-fired power plants emit a lot more $PM_{2.5}$ relative to SO_2, the proportion of primary to secondary $PM_{2.5}$ is substantially higher: about 40% of this $PM_{2.5}$ is primary particulate matter. These ratios suggest that a natural way to bound our projections of long-run effects is to assume that 100% of the mortality is due to fine particulate matter and divide by 3 to get an estimate of the mortality effects of one $\mu g/m^3$ of $PM_{2.5}$. The fact that most of the $PM_{2.5}$ accompanying SO_2 is secondary underscores the policy relevance of our estimates, as reducing SO_2 emissions necessarily reduces relevant secondary pollutants as well.

4.3 Robustness

Our key identifying assumption is that changes in wind direction are unrelated to mortality except through their effects on pollution levels. This identifying assumption is violated if wind direction is correlated with unobserved weather patterns that cause mortality. While this identifying assumption is impossible to test directly, we can probe it indirectly by assessing the sensitivity of our estimates to different ways of controlling for temperature, precipitation, and wind speed.

Columns (2)–(5) of Table 2 shows the estimated effect of SO_2 on same-day all-age all-cause mortality for different sets of weather controls. Column (2) reports our baseline estimate. In Column (3), we also include bins of minimum temperature—with thresholds defined as for maximum temperature—in our weather interaction indicators. In Column (4), we control for all possible interactions of minimum temperatures (specified as indicators for minimum temperatures falling into 3-degree Celsius bins, with outer bins defined by temperatures

[19] InMAP reports all pollutants in $\mu g/m^3$. At standard temperature (15 degrees Celsius), one ppb of SO_2 corresponds to about 2.62 $\mu g/m^3$.

below -15 Celsius or above 30 Celsius); maximum temperatures (specified in the same way); ten deciles of wind speed; and ten deciles of precipitation. This flexible definition of weather follows Deryugina et al. (2019) and gives us 28,899 possible weather conditions ($17 \times 17 \times 10 \times 10 - 1$), although only about one-third of these combinations exists in our data. Column (5) reports estimates for a specification with no weather controls at all. Regardless of the specification, our estimates remain similar.

In Table A.7, we vary the fixed effects. Column (1) reproduces our baseline estimate, while columns (2)-(6) present specifications with alternative fixed effects, such as county, year, and month (column (2)), county and state-by-month-by-year (column (3)), and county-year and state-month fixed effects (column (4)). None of the variations has a meaningful effect on our results.[20] Thus, we are not concerned that omitted variables along these dimensions are biasing our IV estimates.

While acute exposure produces a growing effect on all-cause mortality, Figure 4 shows short-term mortality displacement for deaths where cancer is the underlying cause of death. This disparity in trends raises concerns about composition bias, as cancer-related deaths that occur weeks after exposure could potentially be misattributed to other causes like cardiovascular disease. To address this possibility, we estimate the effect of acute exposure on all deaths where cancer was listed as the underlying or secondary (contributing) causes of death. Those estimates, shown in Figure A.6, are very similar to our main estimates and do not change our conclusions.

The remaining robustness checks are relegated to the Online Appendix and demonstrate the insensitivity of our results to varying the number of included instrument leads and lags (Table A.8); to alternative ways of specifying the instruments (Table A.9); and to using LIML instead of 2SLS (Table A.10). Finally, we generate random wind directions to see if the way we construct our instruments is susceptible to spurious correlation. We leave the other controls as in Equation (1) and use these placebo wind directions as instruments for

[20] We have also estimated specifications that include only county and year fixed effects and a few other variations. These produce very similar results and are not shown for the sake of brevity.

SO_2. The resulting first-stage F-statistics are very low, demonstrating that our strong first stage is not spurious (Table A.11).

5 Long-run survival

5.1 Framework

Our framework for quantifying the long-run survival effects of chronic exposure to SO_2 is based on a dynamic production model of health originally developed by Lleras-Muney and Moreau (2022). Let H_{it} denote the health capital of individual $i \in \{1, ..., N\}$ at time $t \in \{0, ..., T\}$. At birth, an individual is endowed with a stock of initial health, H^*_{i0}, which is drawn randomly from a normal distribution. This health stock evolves over the individual's lifetime according to the following formula:

$$H_{it} = H_{i,t-1} - d(t) + I + \varepsilon_{it} \tag{3}$$

where:

$$\begin{aligned} H_{i0} = H^*_{i0} &\sim N(\mu_H, \sigma_H), \\ d(t) &= \delta t^{\alpha}, \\ \varepsilon_{it} &\sim N(0, \sigma_\varepsilon) \end{aligned}$$

The health stock depreciates at a rate, $d(t)$, which increases with the age of the individual. It is replenished at a constant rate, I, which captures time-invariant factors such as early-life parental investment, and varies with an iid health shock, ε_{it}. Death occurs when the individual's health stock falls below a critical threshold, $\underline{H}$, and is denoted by the indicator

variable D_{it}, where:

$$D_{i0} = 1\left[H_{i0} < \underline{\mathrm{H}}\right],$$
$$D_{it} = 1\left[H_{it} < \underline{\mathrm{H}} \,\middle|\, D_{i,t-1} = 0\right], \; t > 0$$

The model is fully characterized by seven parameters: $\{\alpha, \delta, I, \mu_H, \sigma_H, \sigma_e, \underline{\mathrm{H}}\}$.[21]

Cohort mortality and survival can be calculated to an arbitrary degree of precision by simulating the model given by Equation (3) for a sufficiently large number of individuals, N. The mortality rate at time t, M_t, is equal to the number of individuals in the simulation who died in period t divided by the number of individuals alive at the beginning of period t. Survival can then be calculated from mortality:

$$S_1 = 1 - M_0,$$
$$S_t = S_{t-1}\left(1 - M_{t-1}\right), \; t > 1$$

This parsimonious model is well-suited to our long-run survival analysis. Our objective is to forecast long-run survival, and Lleras-Muney and Moreau (2022) demonstrate that their model's fit across a diverse set of population survival curves is on par with the best models employed by demographers. In addition, the model can separately capture the mortality displacement and accelerated aging effects we observed in our empirical analysis (Figure 4). As we show below, these two effects have very different implications for long-run survival.

The health effects of pollution exposure can be modeled as temporary or permanent changes to one or more model parameters. We focus on changes to the death threshold, $\underline{\mathrm{H}}$, which produces mortality displacement, and the depreciation rate parameters, α and δ, which govern the aging process. Depreciation, $d(t) = \delta t^{\alpha}$, is a power function; changes to

[21] The model can be extended to incorporate external death factors such as car accidents by including two additional parameters specifying the age of onset and the severity of these external mortality causes (Lleras-Muney and Moreau, 2022). However, this extension is unnecessary for our analysis, which focuses on deaths from direct biological (internal) causes.

depreciation parameters will result in larger mortality effects among older individuals than younger ones, consistent with the patterns from our empirical analysis (Figure 3). While exposure could in principle change the level of investment, I, this effect would result in high and persistent mortality increases for both old and young. Because the mortality effects we observe for younger individuals appear to consist solely of short-run displacement, we rule out I as a possibility.[22]

A temporary elevation in the death threshold leads to increased mortality among frail individuals who are closest to death, but has no effect on the health of those still alive. When the threshold reverts to its initial value, mortality rates decline due to the lack of surviving individuals with health capital values near the newly lowered threshold. Consequently, a temporary rise in the death threshold leads to short-run mortality displacement without any longer-lasting effects.

In contrast, depreciation affects the health capital of all individuals—both healthy and ill—leading to persistent changes in mortality. For instance, consider a temporary increase in the depreciation rate, $d(t)$. This event causes an immediate rise in mortality rates among frail individuals whose health capital values fall below the death threshold as a result of increased depreciation. However, because the health capital of the entire population has been reduced, mortality rates remain elevated even after $d(t)$ reverts to its original path.

Our main structural assumption is that the effect of pollution exposure on model parameters depends only on current exposure. This assumption implies that exposure alters the death threshold and depreciation rate parameters by the same amount for the old and the young, and that those changes are independent of exposure history. Thus, observed differences in the mortality effects of exposure are driven by differences in pre-existing stocks of health capital and by the functional form of $d(t)$. A key feature of this assumption is that it allows us to calibrate the effect of exposure on model parameters using IV estimates from

[22]Exposure could also alter the variance of the iid health shock, ε_{it}. We rule out this possibility because it implausibly implies that an increase in pollution exposure would produce health benefits for half of the population.

any age group. It also yields the testable implication that calibrated parameter values from one age group should yield accurate mortality predictions for any other age group.

Our structural assumption is based on research suggesting that, while the health consequences of pollution exposure may vary from person to person, the biological mechanisms are similar at a molecular level. For example, ambient air pollution exposure has been found to increase oxidative stress and have a number of other pathophysiological effects (see Brook et al. (2010) for a review). In a double-blind experiment, Li et al. (2017) exposed healthy young adults to either filtered air or real-world $PM_{2.5}$ concentrations. Higher $PM_{2.5}$ exposure led to increased levels of stress hormones, inflammatory markers, and blood pressure, supporting our assumption of common biological processes. Viewed through the lens of the depreciation parameter, a uniform increase in δ across the age distribution would reduce the health stock more for older individuals than for younger individuals, and would largely raise contemporaneous and near-term mortality only among those individuals who are frail.

Equation (3) ignores optimization behavior, which is consistent with our use of daily-level IV estimates to calibrate the effect of air pollution exposure on model parameters: behavioral adjustments such as buying an air purifier or relocating are unlikely responses to daily changes in air pollution levels. The absence of endogenous responses to exposure means that our survival forecasts should be interpreted as holding behavior fixed, which is the relevant comparison for quantifying the total benefits associated with pollution reduction (Graff Zivin and Neidell, 2013; Currie et al., 2014). In a scenario involving optimal behavioral change, the realized health benefit would be lower than the model's forecast due to risk compensation.

Our long-run survival analysis proceeds as follows. First, we calibrate the baseline model parameters using a period life table from the beginning of our sample time period (1972). Second, we solve for the change in parameter values required to match our IV estimates of the effect of acute pollution exposure on 1-day mortality. We assess the performance of the model by comparing its short-run (up to one month) predictions to IV estimates not used for

the calibration. Finally, we use the model to predict the long-run effects of chronic changes in pollution exposure on survival.

5.2 Calibration

5.2.1 Baseline parameters

The dynamic production model of health given by Equation (3) depends on seven parameters: $\{\alpha, \delta, I, \sigma_e, \mu_H, \sigma_H, \underline{\mathrm{H}}\}$. To achieve identification (i.e., to ensure a unique solution), we follow Lleras-Muney and Moreau (2022) and normalize two parameters: $\underline{\mathrm{H}} = 0$ and $\sigma_H = 1$. We calibrate the five remaining parameters using simulated method of moments. Specifically, we use the Nelder-Meader method to solve numerically for the parameter values that minimize the squared distance between the model's predicted age-specific survival and US population survival in 1972, the first year of our sample period.

Our daily-level IV estimates average about 0.1 deaths per million for a 1-unit change in SO_2 (e.g., Table 2). To measure changes in mortality at that level of granularity, the model must include many millions of individuals. However, calibrating the baseline parameters is computationally infeasible at the daily level when N is very large.[23] We therefore proceed in three steps.

First, we calibrate the baseline model using annual survival data and $N = 1,000,000$ agents. Second, we recalibrate the baseline model using daily survival data and $N = 100,000$ agents, employing (scaled) estimates from the annual calibration as our starting values.[24] Third, we simulate the recalibrated baseline model for a population of $N = 20,000,000$ indi-

[23] If the maximum lifespan is 110 years, then the number of periods is $T = 110 \times 365 = 40,150$ days. Simulating a population of $N = 10,000,000$ individuals thus produces 401.5 billion health capital values (3.2 terabytes, if each value is an 8-byte number). Simulating a single population of this size on a large server requires several hours of runtime, and the Nelder-Mead method requires simulating the population hundreds of times to converge to a solution.

[24] We obtain the starting values by dividing the annual estimates of I and δ by 365, and of σ_e by $\sqrt{365}$. We do not rescale μ_H or α. While these starting values provide a good guess for the solution, the model's (nonlinear) dependence on t causes the guess to be suboptimal, hence the need for recalibration. In the special case where $\alpha = 1$, one can show that if a year is partitioned into P periods, then the optimal value for δ in the partitioned model approaches $1/P$ times the optimal value from the annual model.

viduals, using parameter estimates from the second (daily-level) calibration. This third step provides the high-resolution model estimates of population health capital that we combine with our IV estimates to calibrate the effect of exposure on model parameters.

Figure A.7 illustrates the results of our baseline calibration. The solid blue line shows the survival curve for the US 1972 life table. The red dashed line reports the survival curve produced by our calibrated health model. There are some small deviations between the model's prediction and the observed data in infancy, but otherwise the two curves track each other closely and produce life expectancy estimates that differ by only 0.1 years, indicating a successful model fit. The model's parameter values are reported in Column (2) of Table A.13.

5.2.2 Pollution exposure parameters

Our IV estimates from Section 4 identify the effect of a 1-day, 1-ppb increase in SO_2 exposure on mortality in the following month. In the context of our health production model, we assume that exposure can raise the value of the death threshold parameter, $\underline{H}$, and can increase depreciation by raising the values of δ or α. A 1-day increase in the death threshold will produce short-run mortality displacement that increases current mortality but then reduces near-future mortality by about the same amount, while an increase in depreciation will increase both current and future mortality. Changes in δ and α produce near-identical mortality effects in the short run, so we do not distinguish between them until we turn to our long-run forecasts.[25] For expositional purposes, here we consider changes to δ only. Employing α instead of δ in our validation exercises produces indistinguishable results.

In line with the patterns shown in Figure 4, we assume that cancer-related deaths reflect mortality displacement while other deaths are due to accelerated aging (depreciation). Let $\hat{\beta}^k_{a,c}$ denote the IV estimate of the effect of acute exposure on cumulative mortality for age

[25] The marginal effects of δ and α on depreciation, $d(t) = \delta t^{\alpha}$, vary over time so in theory one could identify which parameter is a better fit for pollution exposure by using variation in mortality by age group. In practice, a 1-day exposure event provides far too little variation for us to separately identify these parameters.

group a from cause of death c in the k days following exposure. Consider a specific daily age, t, that lies inside the age interval spanned by age group a. We first use the high-resolution model estimates described above to solve numerically for a new death threshold, $\widetilde{\underline{H}}_a$, such that elevating the baseline threshold to $\widetilde{\underline{H}}_a$ for a single day at age t yields a mortality increase equal to the 1-day estimate for cancer-related deaths, $\hat{\beta}^1_{a,cancer}$. We then solve for the value of $\widetilde{\delta}_a$ that further increases mortality at age t to match the all-cause (total) estimate, $\hat{\beta}^1_{a,all}$.

Calibrating the parameters $\{\widetilde{\underline{H}}_a, \widetilde{\delta}_a\}$ can be done using IV estimates for any age group a. To maximize efficiency, we perform the calibration using several age groups, and then use the simple average across ages when predicting long-run survival. Because we need death rate magnitudes to be large enough to match the granularity of our model, we only include age groups over 65.[26] The specific IV estimates used for these calibrations are reported in Table A.12.

We use the approximate midpoint of each age bin when solving for the one-day change $\{\widetilde{\underline{H}}_a, \widetilde{\delta}_a\}$. Because health capital depends on an iid health shock, using only a single day for this calibration produces noisy solutions. To improve precision, we solve for the 1-day changes using 50 different days around the midpoint and take the average. For example, for the 65–69 age group, we solve for the changes using ages 68y1d, 68y2d, ..., 68y50d.[27] Appendix A.2 provides additional details. The final result is a set of parameters for each age group, $\{\widetilde{\underline{H}}_a, \widetilde{\delta}_a\}$, capturing the effect of an increase in exposure of 1 ppb. The parameter $\widetilde{\underline{H}}_a$ captures the mortality displacement effect, and the parameter $\widetilde{\delta}_a$ captures the accelerating aging effect.

To account for econometric uncertainty in the IV estimates, we use a resampling-based methodology. We randomly draw an estimate of the effect of acute pollution exposure on

[26] Our model includes includes $N = 20$ million individuals, so a daily death rate of 1 per million corresponds to at most 20 deaths. As explained in the appendix, we improve our precision by repeating our calibration using 50 different points; however, that is insufficient for calibrating to the 1-day death rates for the younger age groups, all of which are less than 0.1 deaths per million.

[27] The optimal strategy would employ all $365 \times 5 = 1825$ days in the 5-year age bin, giving more weight to the ages near the midpoint. However, doing so is computationally burdensome and unnecessary for achieving sufficient accuracy.

1-day mortality from a normal distribution with a mean and standard deviation set equal to the mean and standard error of $\widehat{\beta}^1_{a,c}$ in Equation (1), and then calibrate the change in model parameters to match the mortality change draw. We repeat this exercise 100 times and report the 5th and 95th percentiles of the resulting distribution of model parameter estimates.

Our age-specific estimates of changes in the mortality threshold and one of the depreciation parameters, $\{\widetilde{\underline{H}}_a, \widetilde{\delta}_a\}$, expressed as deviations from the baseline parameter values, are shown in Figure A.8. For completeness, we show estimates for $\widetilde{\alpha}_a$ as well. Estimates are generally similar across ages, with most values falling inside the 90 percent confidence intervals of other estimates, which is consistent with our assumption that the effects of exposure on model parameters are constant over the life cycle.

5.3 Validation

We validate the model by comparing its predictions of the mortality effect of acute exposure to IV estimates that were not used in the calibration. Because calibration for age group a employs only the 1-day IV estimate for that age group, our first validation test focuses on model predictions for longer-run (>1 day) outcomes. Figure 5 provides an example using the 65–69 age group. The solid blue line reports the IV estimates, and the green dot-dashed line reports "own-age" predictions, which come from a model that was calibrated using $\hat{\beta}^1_{65,cancer}$ and $\hat{\beta}^1_{65,all}$. By construction, the own-age model prediction matches the IV estimate for the same-day effect. To assess model performance, we compare the own-age predictions of longer-run mortality changes to the corresponding IV estimates. All of these predictions lie within the 95% confidence intervals of the IV estimates.

Our second, more demanding, validation test involves computing mortality predictions for a given age group using the average of the calibrated parameters from other age groups. For example, instead of using the calibrated parameters $\{\widetilde{\underline{H}}_{65}, \widetilde{\delta}_{65}\}$ as we did for the own-age predictions, we predict mortality for the 65–69 age group using the average of the calibrated

values from the 70–74, 75–79, 80–84, and 85+ age groups: $\{\frac{1}{4}\sum_{a>65}\widetilde{\underline{H}}_a, \frac{1}{4}\sum_{a>65}\widetilde{\delta}_a\}$. Here, the 1-day mortality prediction also serves as a validation test, since $\hat{\beta}^1_{65}$ was not used to form this "leave-one-out" prediction. The thick red dashed line in Figure 5 shows that all of these leave-one-out predictions lie inside the 95% confidence intervals of the IV estimates. One compelling feature of this validation exercise is that it produces good predictions despite its reliance on IV estimates that are nearly one order of magnitude larger in size than the 65–69 predictions (Figure 3a).

For purposes of comparison, we also plot predictions from a model specification that assumes none of the 1-day exposure deaths are due to changes in the death threshold ("no displacement") and one that assumes all of the deaths are due to changes in the death threshold ("all displacement"). In the all-displacement scenario, depicted by the black dashed line in Figure 5, the cumulative mortality effect falls rapidly to zero, reflecting complete short-run mortality displacement. In the no-displacement scenario, depicted by the orange dashed line, the mortality predictions increase much more rapidly than either the own-age or leave-one-out predictions. Notably, the predictions from these two extreme specifications lie well outside the 95% confidence intervals of our IV estimates, demonstrating that the latter are precise enough to rule out a range of model predictions. The wide gap between the all-displacement and no-displacement predictions underscores the importance of accurately identifying the fraction of deaths that are attributable to mortality displacement rather than accelerated aging.

Figure A.10 shows corresponding leave-one-out predictions for the other age groups. Overall, 26/30 (87%) of the model predictions lie inside the 95% confidence intervals of the IV estimates. The close concordance between these predictions and the IV estimates suggests that our dynamic production model of health accurately captures important determinants of mortality, at least for 1-day shocks with a one-month follow-up window. It also supports our assumption that the effect of acute exposure on model parameters is the same across different ages.

5.4 Long-run predictions

Finally, we use the methodology outlined in Section 5.2.2 to quantify the effects on life expectancy of a permanent, 1-unit decrease in SO_2 exposure. For each age group a, we first solve for the parameter values $\{\widetilde{\underline{H}}_a, \widetilde{\delta}_a\}$ that yield the 1-day mortality change implied by the IV estimate for that age group. We then predict survival using the average value for the five age groups 65 and over: $\{\frac{1}{5}\sum_{a\geq 65}\widetilde{\underline{H}}_a, \frac{1}{5}\sum_{a\geq 65}\widetilde{\delta}_a\}$. Rather than quantifying the survival effects of a 1-day change in parameter values, however, here we quantify the effect of a permanent, lifetime change beginning at birth ($t = 0$). Because we are quantifying long-run effects, the difference between changing the depreciation parameters δ and α is no longer negligible, so we show both sets of estimates. For purposes of comparison, we also quantify the implied life expectancy effect of extrapolating our age-specific monthly IV estimates to the entire life cycle.[28]

Figure 6 illustrates our results. The baseline life expectancy in our model—which was calibrated using a 1972 period life table—is 71.32 years. The two model-based calculations imply that a permanent, 1-ppb decrease in SO_2 improves life expectancy for this cohort by 1.18–1.32 years (1.65–1.85%).[29] By contrast, extrapolating our IV estimates to the life cycle yields a life expectancy improvement of about 0.16 years (0.22%), which is 7–8 times smaller than the model-based estimates. The IV extrapolation is biased downwards because it is limited by the one-month regression window and does not account for the latent effects of pollution on people's health capital over longer time periods.

Even though decreased SO_2 exposure begins at birth, Figure 6 indicates that survival gains are concentrated among older individuals: more than 90% of the life expectancy im-

[28] These IV estimates are reported in the last row of Table A.4. This naive calculation assumes that daily mortality rates increase at every age by the monthly IV point estimate for the corresponding age group, including point estimates that are statistically insignificant. We assume the estimates in Table A.4 apply to the midpoint of each age bin, and interpolate to calculate the rest of the values.

[29] Table A.14 reports results for the survival effects of chronic exposure changes up to 3 ppb. The relationship between survival improvements and permanent changes in SO_2 is roughly linear, although we caution against generalizing from these results, since they rely heavily on the assumed linear relationship between SO_2 and mortality in Equation (1).

provement occurs after age 50 and more than 75% occurs after the age of 65. Intuitively, this is because additional health capital at younger ages translates into negligible mortality improvements, as overall health is high and most individuals are decades away from dying. As people age and their health capital approaches the death threshold, however, health capital accrued at younger ages becomes increasingly important for determining longevity.

Interpreting the life expectancy improvements requires some care, as the nature of our pollution data means that some unobserved $PM_{2.5}$ accompanies every unit of observed SO_2. Due to lack of monitoring data during our study period, we cannot measure this ratio directly. However, our simulations in Section 4.2 indicated that, on average, 1 ppb of SO_2 is accompanied by 3 $\mu g/m^3$ of $PM_{2.5}$ or less. If we were to make the extreme assumption that all of our estimated mortality effects of SO_2 were in fact operating through $PM_{2.5}$, then our estimates imply that a permanent 1-$\mu g/m^3$ decrease in $PM_{2.5}$ improves life expectancy by 0.39–0.43 years. We note that even under this extreme assumption, our main estimates remain policy relevant because the vast majority of co-transported $PM_{2.5}$ in our sample likely comes from the conversion of SO_2 to $PM_{2.5}$ (see Section 4.2).

6 Conclusion

Accurate estimates of the long-run effect of chronic pollution exposure on health and mortality are vital for making informed policy decisions. Yet, reliable causal estimates remain scarce. We address this gap by proposing a novel two-step approach that combines well-identified short-run estimates of the mortality effect of pollution with a dynamic production model of health. Our approach can simulate a variety of exposure counterfactuals, including the long-run mortality effect of chronic exposure. Although we focus on air pollution, our method can be applied to a number of other health hazards, provided that the researcher can estimate short-run mortality effects and credibly identify the relevant model parameters affected by the hazard.

To obtain well-identified estimates for our model, we assemble a new, comprehensive dataset of daily US mortality and weather, and then instrument for changes in SO_2 levels using changes in wind direction. We show that the short-run mortality effects of acute SO_2 exposure can be decomposed into two distinct phenomena: mortality displacement, where exposure kills frail individuals with short counterfactual life expectancies, and accelerated aging, where mortality continues to increase even after exposure has ceased. In our setting, the aging effect dominates, causing net mortality to rise significantly in the month following a 1-day increase in SO_2.

We integrate these short-run mortality estimates into the model of survival by separately calibrating its parameters to our mortality displacement and accelerated aging estimates. We then calculate that a permanent, 1 ppb decrease in SO_2 exposure improves life expectancy by 1.2–1.3 years. This increase is almost an order of magnitude larger than a naive extrapolation of the monthly IV estimates, demonstrating the importance of accounting for latent health changes caused by pollution exposure.

While our projections of the effect of pollution exposure on life expectancy help inform the benefits of air pollution regulations, these estimates should not be interpreted as predictions of how US life expectancy will change, for two reasons. First, we do not account for other phenomena that matter greatly for life expectancy, such as the opioid epidemic and medical innovation. Second, the model explicitly ignores endogenous responses to air pollution, such as migration. The estimates are thus best interpreted as the benefits of reducing air pollution exposure holding longer-run behavior fixed. To understand whether it is more cost-effective to reduce pollution exposure through private action or through public policy, one could compare the costs of these two options per unit of exposure. Regardless of which option is more cost-effective, our estimates help quantify the benefits of such actions.

References

Agency for Toxic Substances and Disease Registry (1998). Public health statement for sulfur dioxide. Technical report, U.S. Department of Health and Human Services.

Alexander, D. and H. Schwandt (2022). The impact of car pollution on infant and child health: Evidence from emissions cheating. *The Review of Economic Studies 89*(6), 2872–2910.

Arenberg, S. and S. Neller (2023). Ashes to ashes: The lifelong consequences of early-life wildfire exposure.

Athey, S., R. Chetty, and G. Imbens (2020). Combining experimental and observational data to estimate treatment effects on long term outcomes. *arXiv preprint arXiv:2006.09676*.

Barreca, A., K. Clay, O. Deschenes, M. Greenstone, and J. S. Shapiro (2016). Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century. *Journal of Political Economy 124*(1), 105–159.

Barreca, A. I., M. Neidell, and N. J. Sanders (2021). Long-run pollution exposure and adult mortality: Evidence from the acid rain program. *Journal of Public Economics 200*, 104440.

Bishop, K., J. Ketcham, and N. Kuminoff (2023). Hazed and confused: The effect of air pollution on dementia. *The Review of Economic Studies*, forthcoming.

Brook, R. D., S. Rajagopalan, C. A. Pope III, J. R. Brook, A. Bhatnagar, A. V. Diez-Roux, F. Holguin, Y. Hong, R. V. Luepker, M. A. Mittleman, et al. (2010). Particulate matter air pollution and cardiovascular disease: an update to the scientific statement from the american heart association. *Circulation 121*(21), 2331–2378.

Burton, A. and T. Roach (2023). Negative externalities of temporary reductions in cognition: Evidence from particulate matter pollution and fatal car crashes. *Working Paper*.

Chen, Y., A. Ebenstein, M. Greenstone, and H. Li (2013). Evidence on the impact of sustained exposure to air pollution on life expectancy from chinas huai river policy. *Proceedings of the National Academy of Sciences 110*(32), 12936–12941.

Colmer, J. and J. Voorheis (2020). The grandkids aren't alright: the intergenerational effects of prenatal pollution exposure. Working Paper.

Crüts, B., L. van Etten, H. Törnqvist, A. Blomberg, T. Sandström, N. L. Mills, and P. J. Borm (2008). Exposure to diesel exhaust induces changes in eeg in human volunteers. *Particle and Fibre toxicology 5*(1), 4.

Currie, J. and M. Neidell (2005). Air pollution and infant health: What can we learn from california's recent experience? *Quarterly Journal of Economics 120*(3), 1003–1030.

Currie, J., J. G. Zivin, J. Mullins, and M. Neidell (2014). What do we know about short- and long-term effects of early-life exposure to pollution? *Annual Review of Resource Economics 6*(1), 217–247.

Deryugina, T., G. Heutel, N. Miller, D. Molitor, and J. Reif (2019). The mortality and medical costs of air pollution: Evidence from changes in wind direction. *American Economic Review 109*(12), 4178–4219.

Deschênes, O. and M. Greenstone (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us. *American Economic Journal: Applied Economics 3*(4), 152–185.

Deschênes, O., M. Greenstone, and J. S. Shapiro (2017). Defensive investments and the demand for air quality: Evidence from the nox budget program. *American Economic Review 107*(10), 2958–89.

Dominici, F., M. Greenstone, and C. R. Sunstein (2014). Particulate matter matters. *Science 344*(6181), 257–259.

Ebenstein, A., M. Fan, M. Greenstone, G. He, and M. Zhou (2017). New evidence on the impact of sustained exposure to air pollution on life expectancy from chinas huai river policy. *Proceedings of the National Academy of Sciences 114*(39), 10384–10389.

Englert, N. (2004). Fine particles and human health - a review of epidemiological studies. *Toxicology Letters 149*, 235–242.

Environmental Protection Agency (2004). The particle pollution report: Current understanding of air quality and emissions through 2003. Technical report, U.S. Environmental Protection Agency.

Fonken, L. K., X. Xu, Z. M. Weil, G. Chen, Q. Sun, S. Rajagopalan, and R. J. Nelson (2011). Air pollution impairs cognition, provokes depressive-like behaviors and alters hippocampal cytokine expression and morphology. *Molecular Psychiatry 16*(10), 987.

Graff Zivin, J. and M. Neidell (2012). The impact of pollution on worker productivity. *American Economic Review 102*(7), 3652–3673.

Graff Zivin, J. and M. Neidell (2013). Environment, health, and human capital. *Journal of Economic Literature 51*(3), 689–730.

Heo, S. W., K. Ito, and R. Kotamarthi (2023). International spillover effects of air pollution: Evidence from mortality and health data. Technical report, National Bureau of Economic Research.

Heutel, G., N. H. Miller, and D. Molitor (2021). Adaptation and the mortality effects of temperature across us climate regions. *Review of Economics and Statistics 103*(4), 740–753.

Hill, W., E. L. Lim, C. E. Weeden, C. Lee, M. Augustine, K. Chen, F.-C. Kuan, F. Marongiu, E. J. Evans Jr, D. A. Moore, et al. (2023). Lung adenocarcinoma promotion by air pollutants. *Nature 616*(7955), 159–167.

Hollingsworth, A. and I. Rudik (2021). The effect of leaded gasoline on elderly mortality: Evidence from regulatory exemptions. *American Economic Journal: Economic Policy 13*(3), 345–73.

Hollingsworth, A. J., D. M. Konisky, and N. Zirogiannis (2021). The health consequences of excess emissions: Evidence from texas. *Journal of Environmental Economics and Management 108*, 102449.

Isen, A., M. Rossin-Slater, and R. 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.

Jusot, J.-F., D. R. Neill, E. M. Waters, M. Bangert, M. Collins, L. B. Moreno, K. G. Lawan, M. M. Moussa, E. Dearing, D. B. Everett, et al. (2017). Airborne dust and high temperatures are risk factors for invasive bacterial disease. *Journal of Allergy and Clinical Immunology 139*(3), 977–986.

Knittel, C. R., D. L. Miller, and N. J. Sanders (2016). Caution, drivers! children present: Traffic, pollution, and infant health. *Review of Economics and Statistics 98*(2), 350–366.

Landrigan, P. J., R. Fuller, N. J. Acosta, O. Adeyi, R. Arnold, A. B. Baldé, R. Bertollini, S. Bose-O'Reilly, J. I. Boufford, P. N. Breysse, et al. (2018). The lancet commission on pollution and health. *The Lancet 391*(10119), 462–512.

Li, H., J. Cai, R. Chen, Z. Zhao, Z. Ying, L. Wang, J. Chen, K. Hao, P. L. Kinney, H. Chen, et al. (2017). Particulate matter exposure and stress hormone levels: a randomized, double-blind, crossover trial of air purification. *Circulation 136*(7), 618–627.

Lleras-Muney, A. and F. Moreau (2022). A unified model of cohort mortality. *Demography 59*(6), 2109–2134.

Logan, W. (1953). Mortality in the london fog incident, 1952. *Lancet 264*(1), 336–338.

Luria, M., R. E. Imhoff, R. J. Valente, W. J. Parkhurst, and R. L. Tanner (2001). Rates of conversion of sulfur dioxide to sulfate in a scrubbed power plant plume. *Journal of the Air & Waste Management Association 51*(10), 1408–1413.

Michele, M., M. Alberto, S. Liana, and D. Francesco (2006). Do environmental factors influence the occurrence of acute meningitis in industrialized countries? an epidemic of varying aetiology in northern italy. *European Journal of Epidemiology 21*, 465–468.

Nadler, D. L. and I. G. Zurbenko (2014). Estimating cancer latency times using a weibull model. *Advances in Epidemiology 2014*.

Qin, G., M. Wu, J. Wang, Z. Xu, J. Xia, and N. Sang (2016). Sulfur dioxide contributes to the cardiac and mitochondrial dysfunction in rats. *Toxicological Sciences 151*(2), 334–346.

Rajagopalan, S. and P. J. Landrigan (2021). Pollution and the heart. *New England Journal of Medicine 385*(20), 1881–1892.

Sang, N., Y. Yun, H. Li, L. Hou, M. Han, and G. Li (2010). So2 inhalation contributes to the development and progression of ischemic stroke in the brain. *Toxicological Sciences 114*(2), 226–236.

Schlenker, W. and M. J. Roberts (2009). Nonlinear temperature effects indicate severe damages to u.s. crop yields under climate change. *Proceedings of the National Academy of Sciences 106*(37), 15594–15598.

Schlenker, W. and W. R. Walker (2016). Airports, air pollution, and contemporaneous health. *Review of Economic Studies 83*(2), 768–809.

Shears, R. K., L. C. Jacques, G. Naylor, L. Miyashita, S. Khandaker, F. Lebre, E. C. Lavelle, J. Grigg, N. French, D. R. Neill, et al. (2020). Exposure to diesel exhaust particles increases susceptibility to invasive pneumococcal disease. *Journal of Allergy and Clinical Immunology 145*(4), 1272–1284.

Tessum, C., J. Hill, J. Marshall, and D. Paolella (2019, September). Evaluation data for the Intervention Model for Air Pollution (InMAP) version 1.6.1.

Turner, M. C., Z. J. Andersen, A. Baccarelli, W. R. Diver, S. M. Gapstur, C. A. Pope III, D. Prada, J. Samet, G. Thurston, and A. Cohen (2020). Outdoor air pollution and cancer: An overview of the current evidence and public health recommendations. *CA: a Cancer Journal for Clinicians 70*(6), 460–479.

Voorheis, J. (2017). Air quality, human capital formation and the long-term effects of environmental inequality at birth. US Census Center for Economic Studies, CARRA Working Paper Series.

Yao, G., H. Yue, Y. Yun, and N. Sang (2015). Chronic so2 inhalation above environmental standard impairs neuronal behavior and represses glutamate receptor gene expression and memory-related kinase activation via neuroinflammation in rats. *Environmental Research 137*, 85–93.

Figure 1: The relationship between wind direction and SO_2 concentration, Greater Philadelphia and Southern California areas

Greater Philadelphia area

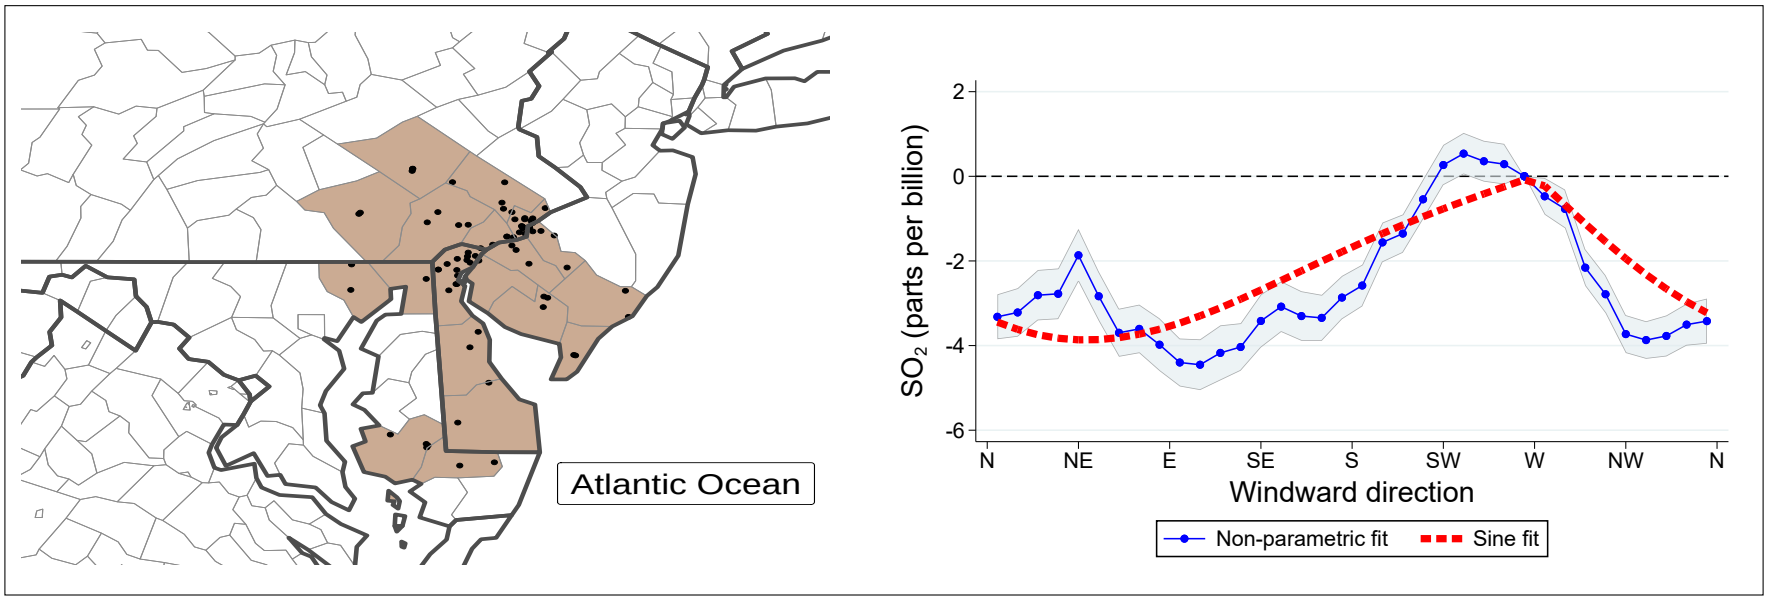

Southern California area

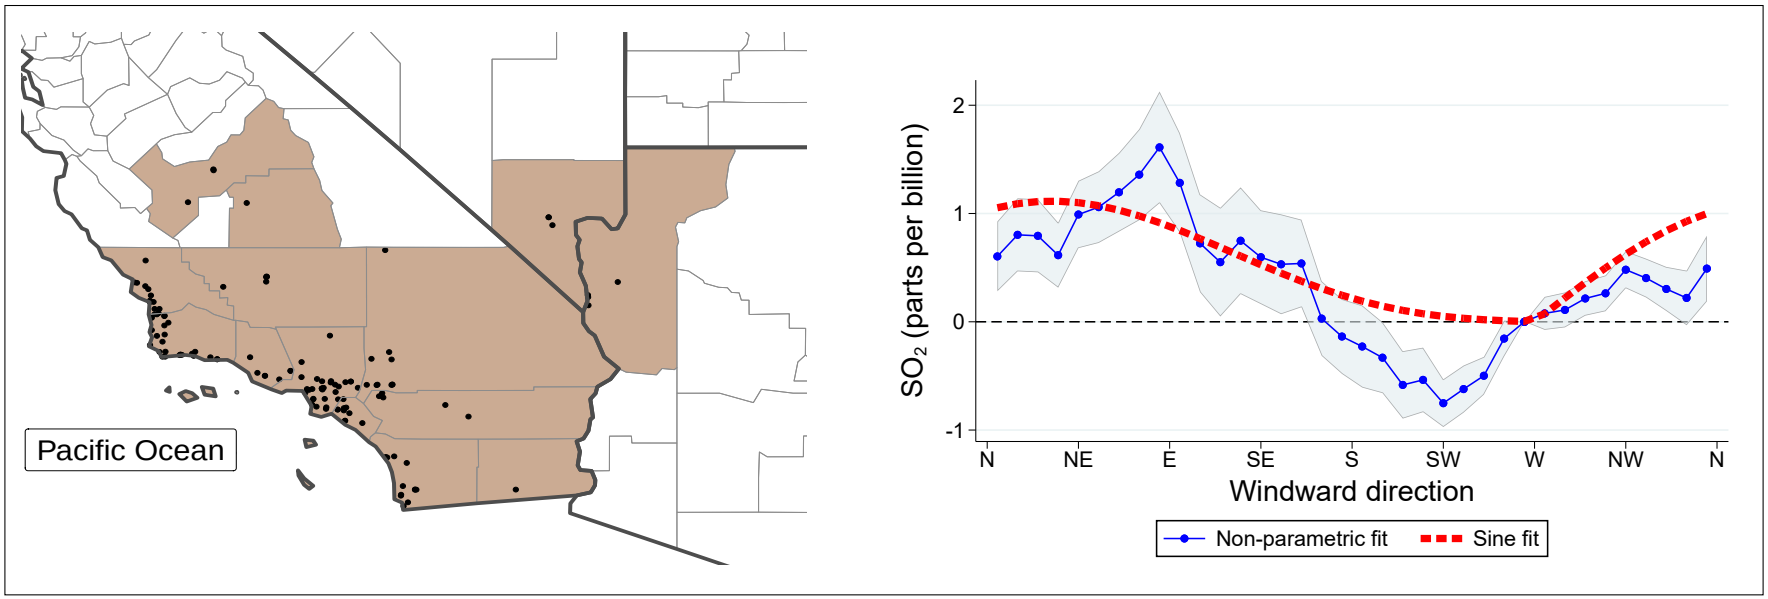

Notes: Sulfur dioxide (SO_2) pollution monitors are depicted as black dots on the two maps. The graphs on the right plot the relationship between SO_2 levels and windward direction in each area. Windward direction describes where the wind is blowing from, with "N" indicating North, "NE" indicating Northeast, etc. The 36 blue points report coefficient estimates from a non-parametric regression of SO_2 on wind direction measured in 10-degree angle bins. The blue shaded area shows the corresponding 95% confidence intervals. The red dashed lines report coefficient estimates from the parametric sine specification given by $f^g(\theta)$ in Equation (2). Regressions include county-by-month and month-by-year fixed effects, and flexible weather controls. Standard errors are robust to heteroskedasticity.

Figure 2: IV estimates of effect of acute (1-day) SO_2 exposure on cumulative mortality

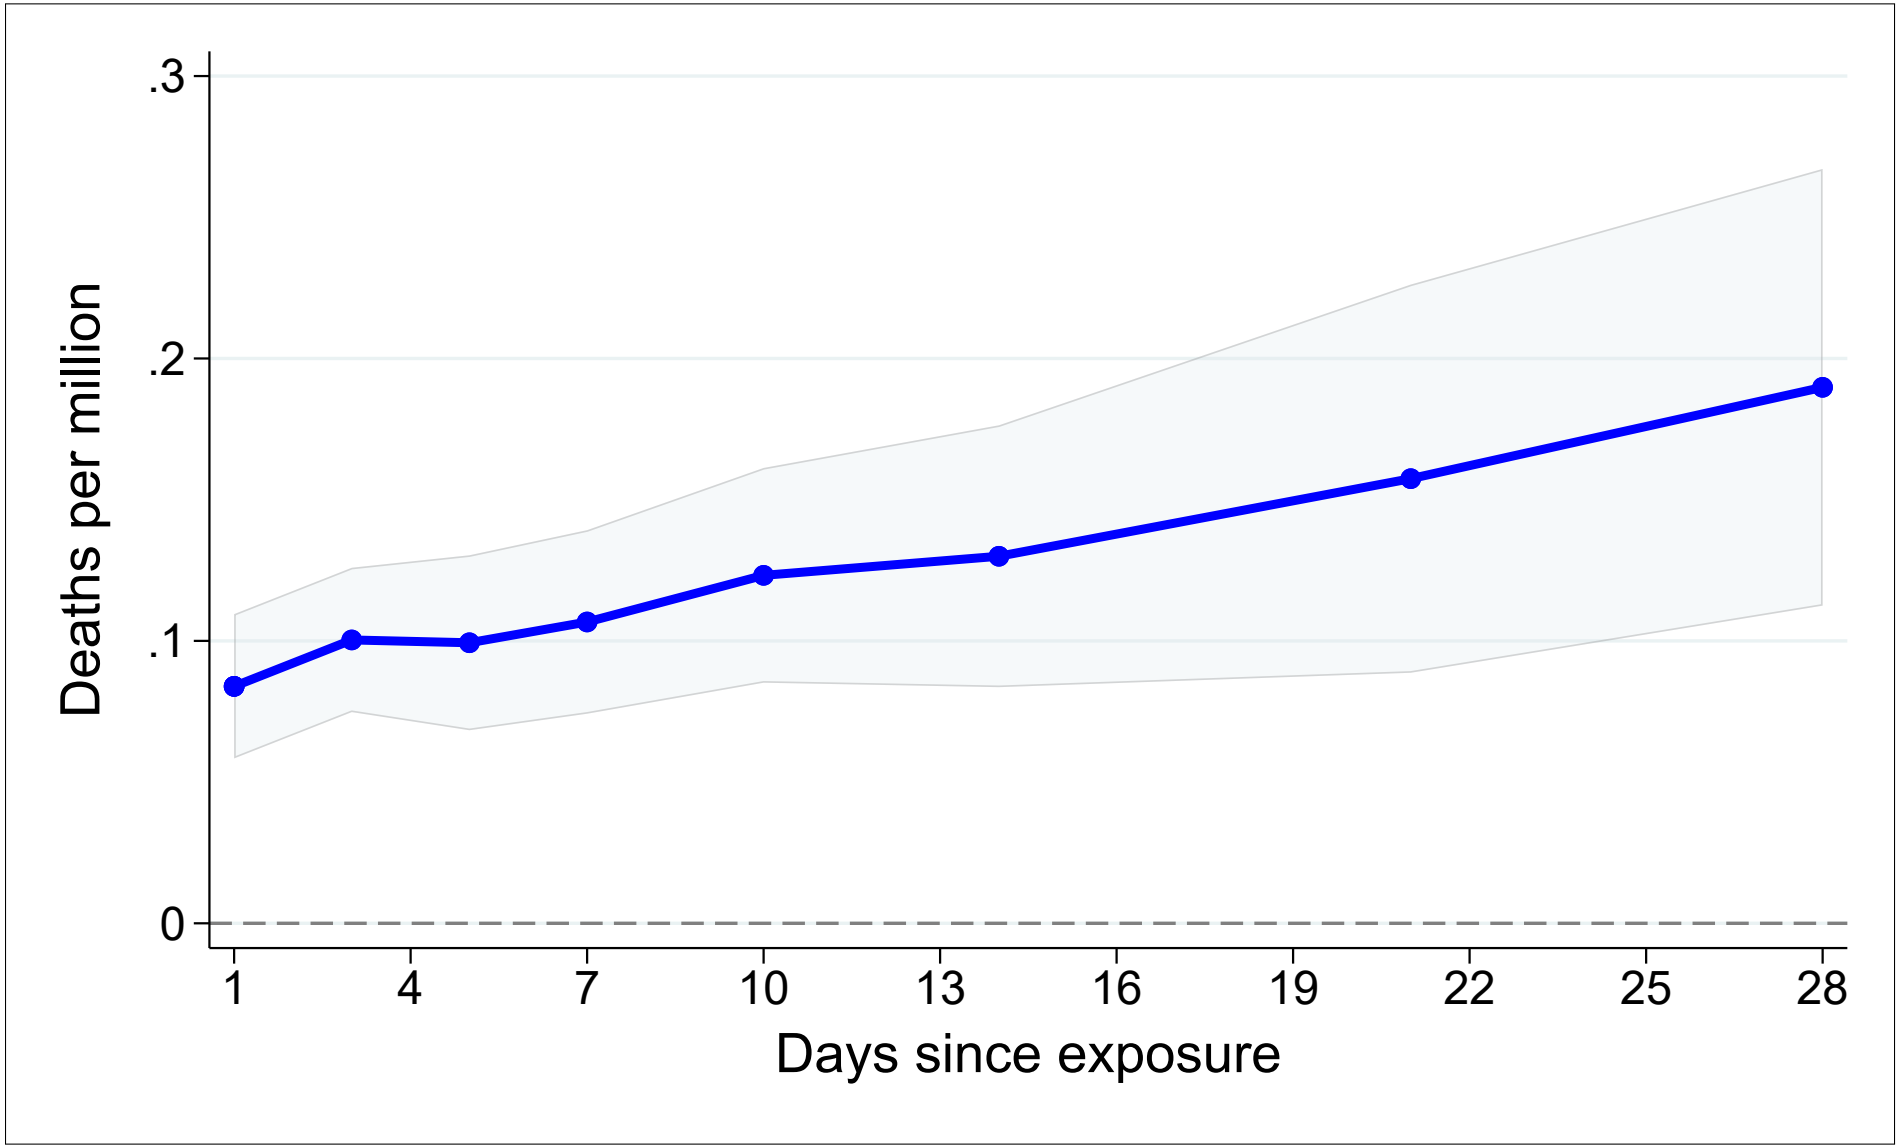

Notes: Each point reports an IV estimate from Equation (1) of the effect of acute (1-day), 1-ppb sulfur dioxide (SO_2) exposure on mortality, where mortality is measured as cumulative deaths per million over a time window ranging from 1 to 28 days, as indicated by the x-axis. Shaded area depicts 95% confidence intervals. Estimates are also reported in Column (2) of Table A.2. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors are clustered by county.

Figure 3: IV estimates of effect of acute (1-day) SO_2 exposure on 1-day mortality, by age group

(a) Absolute increase (deaths per million)

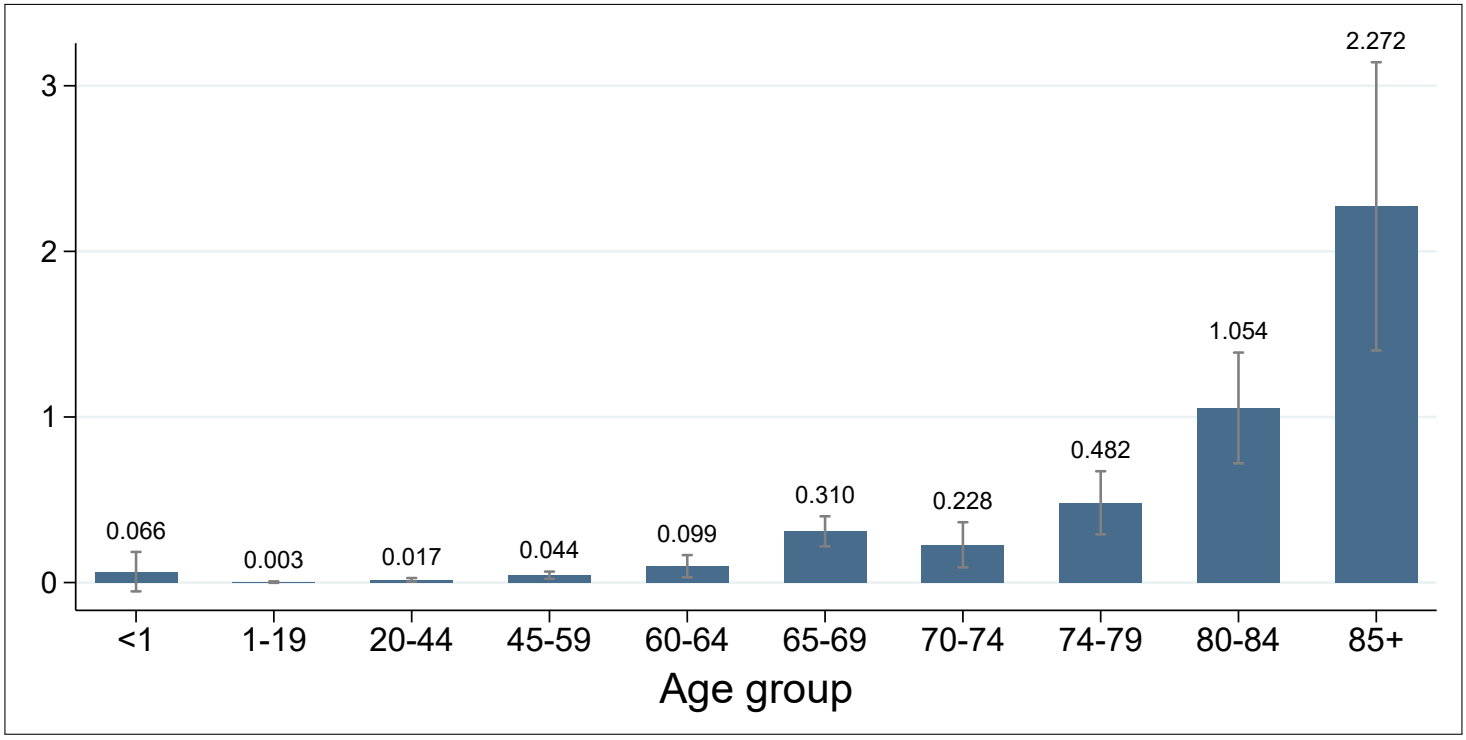

(b) Relative increase (percent of 1-day mortality)

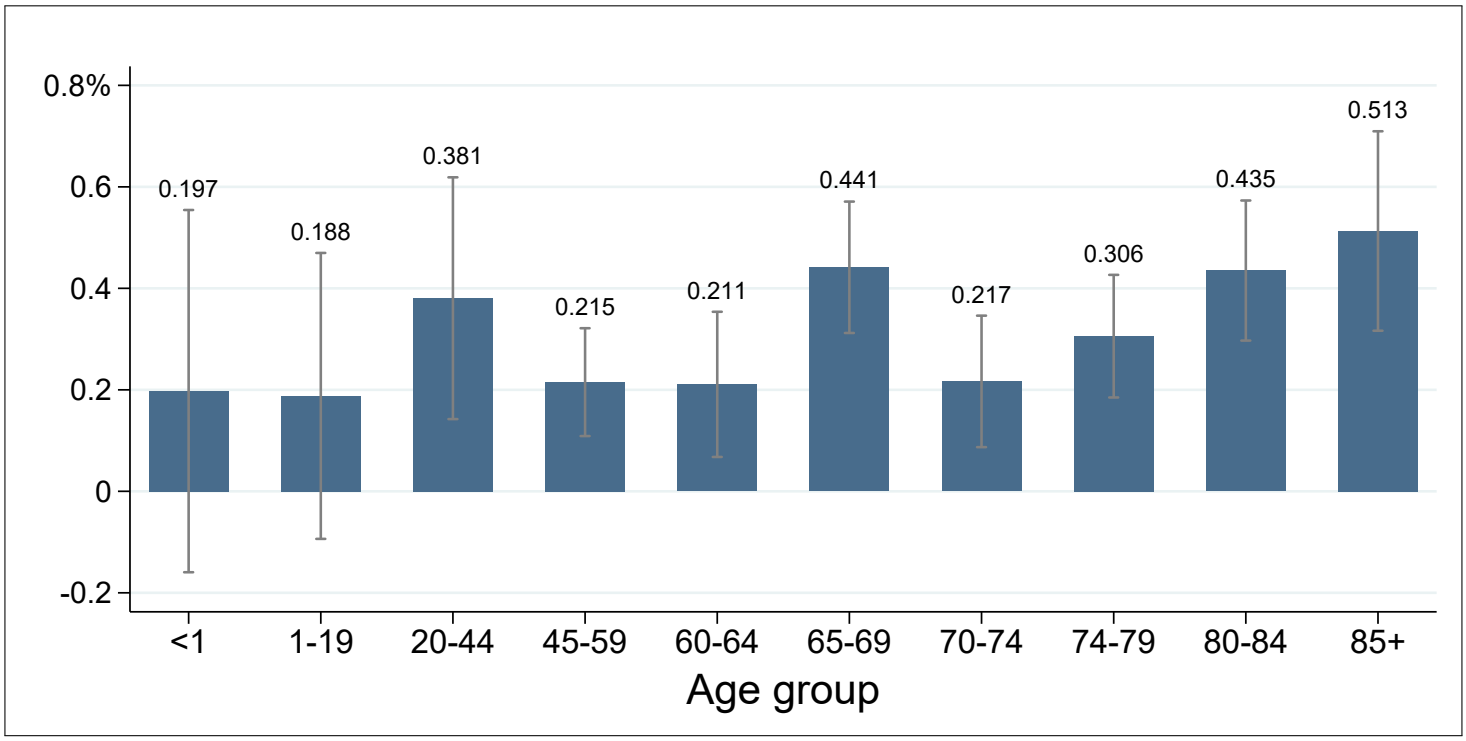

Notes: Each bar represents an IV estimate from Equation (1) of the 1-day mortality rate on a 1-day change in SO_2 for the age group specified. Error bars report 95% confidence intervals. Estimates are also reported in Table A.3. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors are clustered by county.

Figure 4: IV estimates of effect of acute (1-day) SO_2 exposure on cumulative mortality, by cause of death

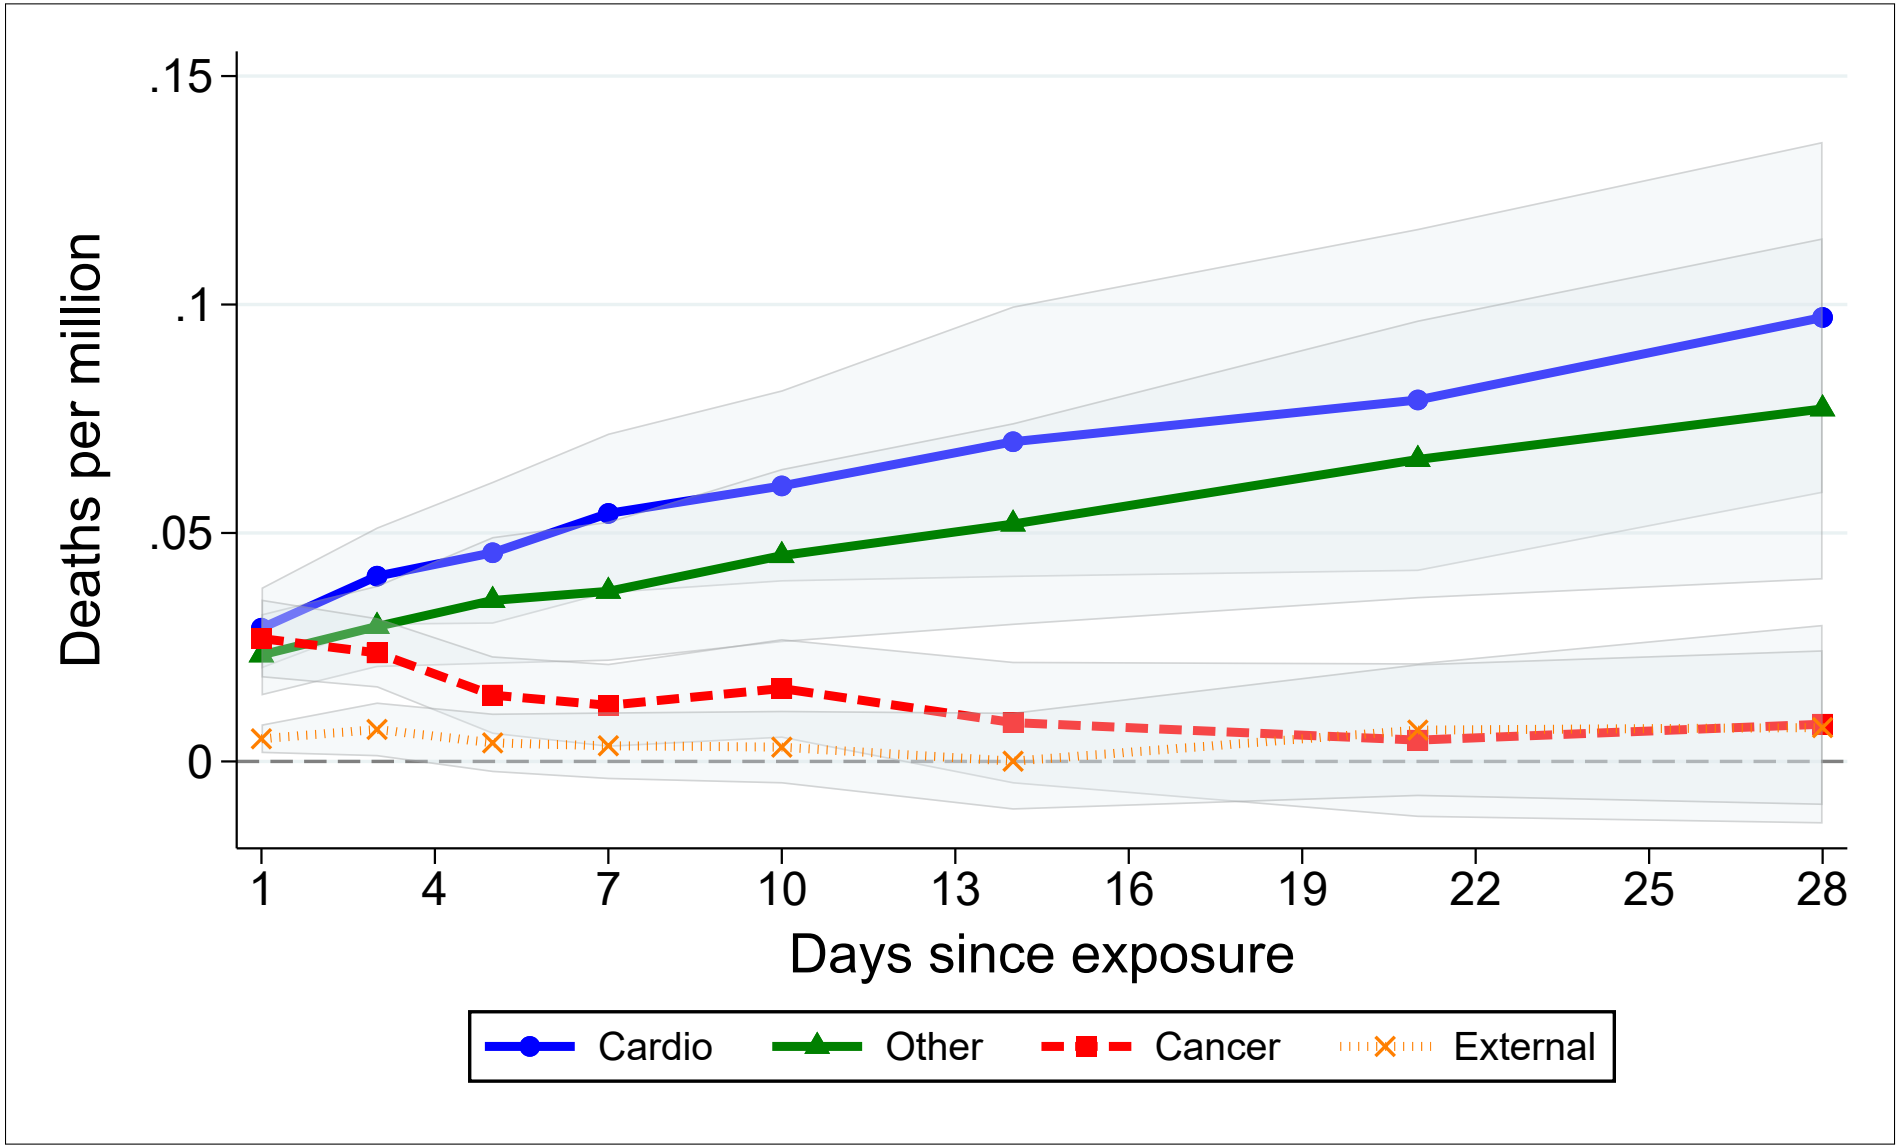

Notes: Each point reports an IV estimate from Equation (1) of the effect of acute (1-day) sulfur dioxide (SO_2) exposure on mortality (deaths per million), for four different causes of death: cardiovascular disease, cancer, other diseases, and external causes of death. Mortality is measured with a time window ranging from 1 to 28 days, as indicated by the x-axis. Shaded areas correspond to 95% confidence intervals. Estimates are also reported in Table A.5. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors are clustered by county.

Figure 5: Model predictions of effect of acute (1-day) SO_2 exposure on one-month survival, ages 65–69

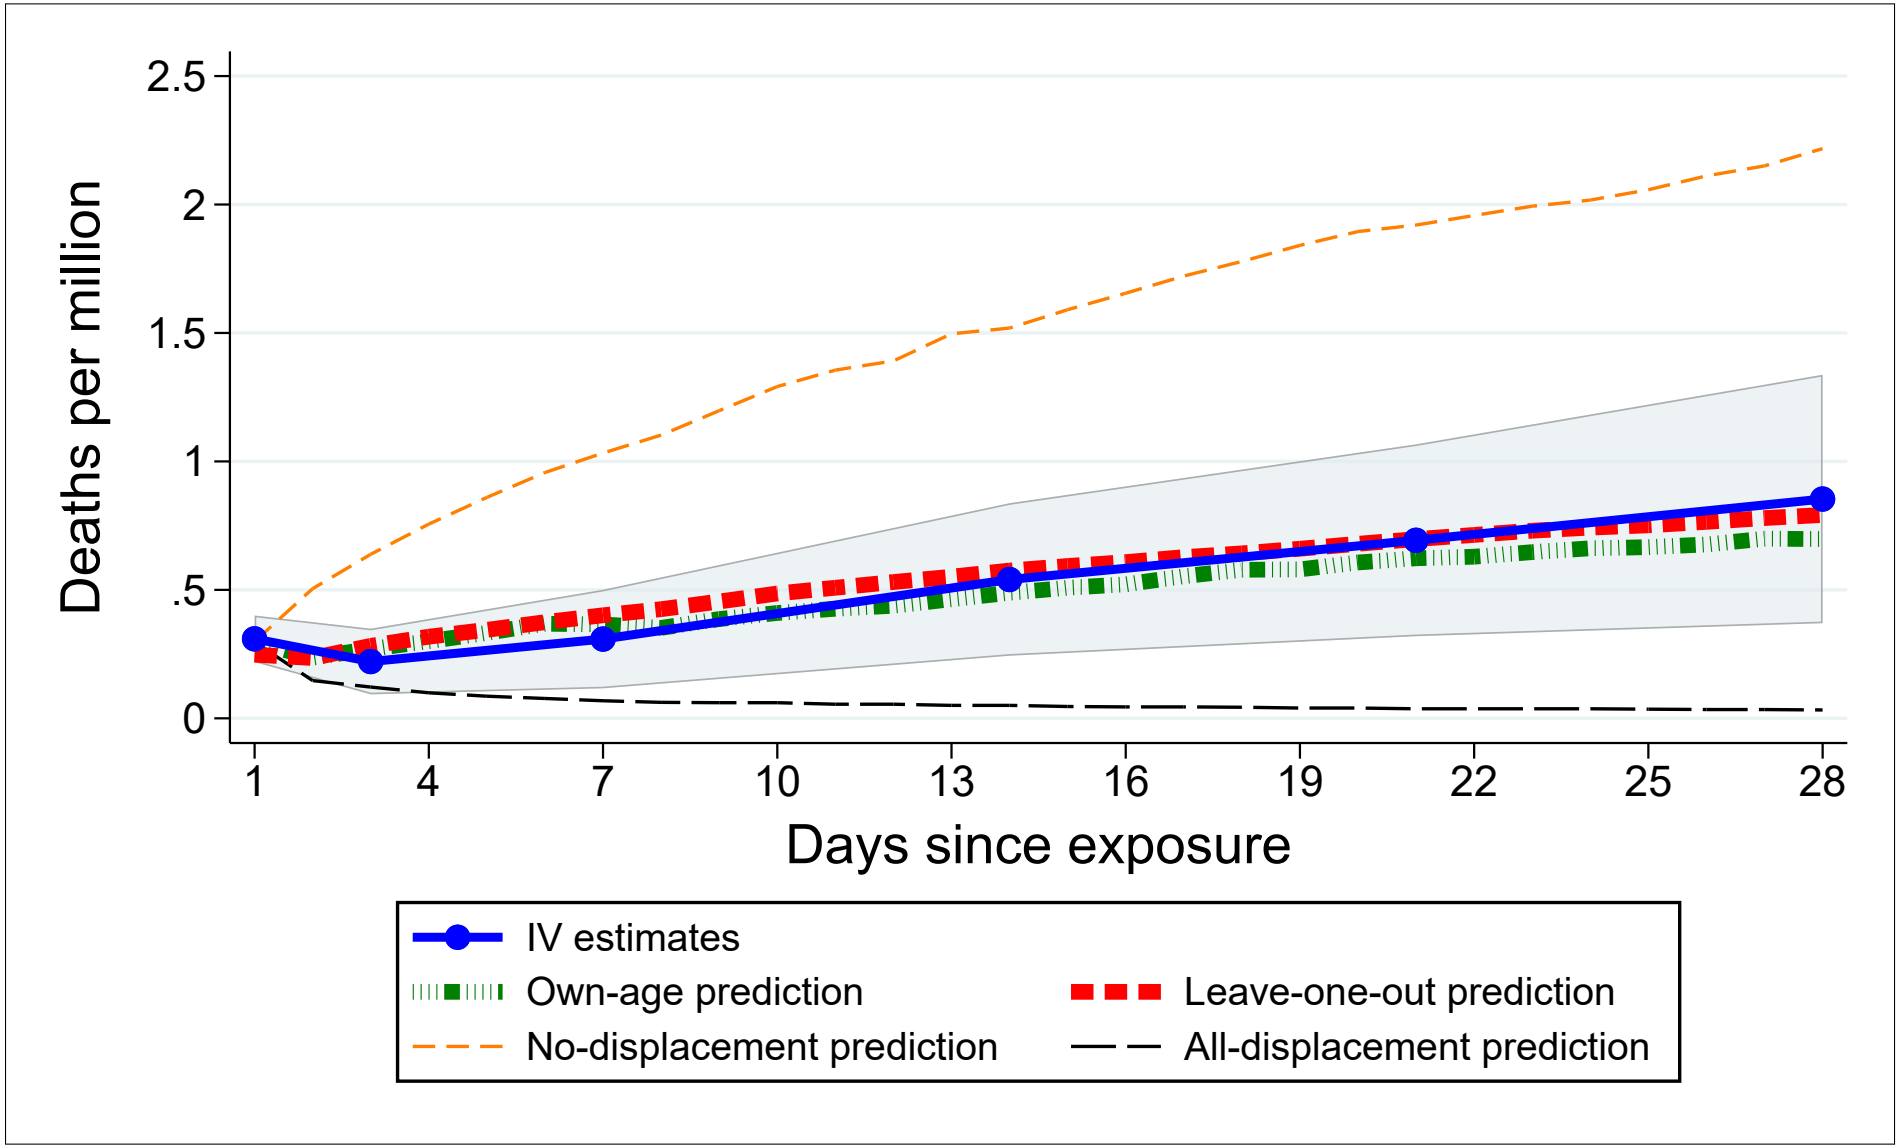

Notes: The solid blue line reports IV estimates from Equation (1) of the effect of acute (1-day) sulfur dioxide (SO_2) exposure on cumulative mortality for the 65–69 age group, with 95% confidence intervals given by the blue shaded area. The thick red and thick green dashed lines reports corresponding predictions from the dynamic production model of health described by equation (3). The "age-specific calibration" model, depicted by the thick red dashed line, is calibrated so that its 1-day prediction matches the 1-day IV estimate (first blue point), and it attributes the cancer-related portion of the 1-day IV estimate to the frailest individuals with the lowest levels of health capital. The "leave-one-out" model, depicted by the thick green dashed line, employs the average of the calibrated values for all of the other older age groups, i.e., ages 70–74, 75–79, 80–84, and 85+. The orange dashed line ("no displacement") reports model predictions under the extreme assumption that none of the 1-day mortality effect is mortality displacement, while the green dashed line ("all displacement") reports predictions under the alternative extreme assumption that all of the 1-day effect is mortality displacement. Figure A.10 shows plots for all age groups 60 and over.

Figure 6: Predicted effect of a permanent 1-unit decrease in SO_2 on survival for cohort born in 1972

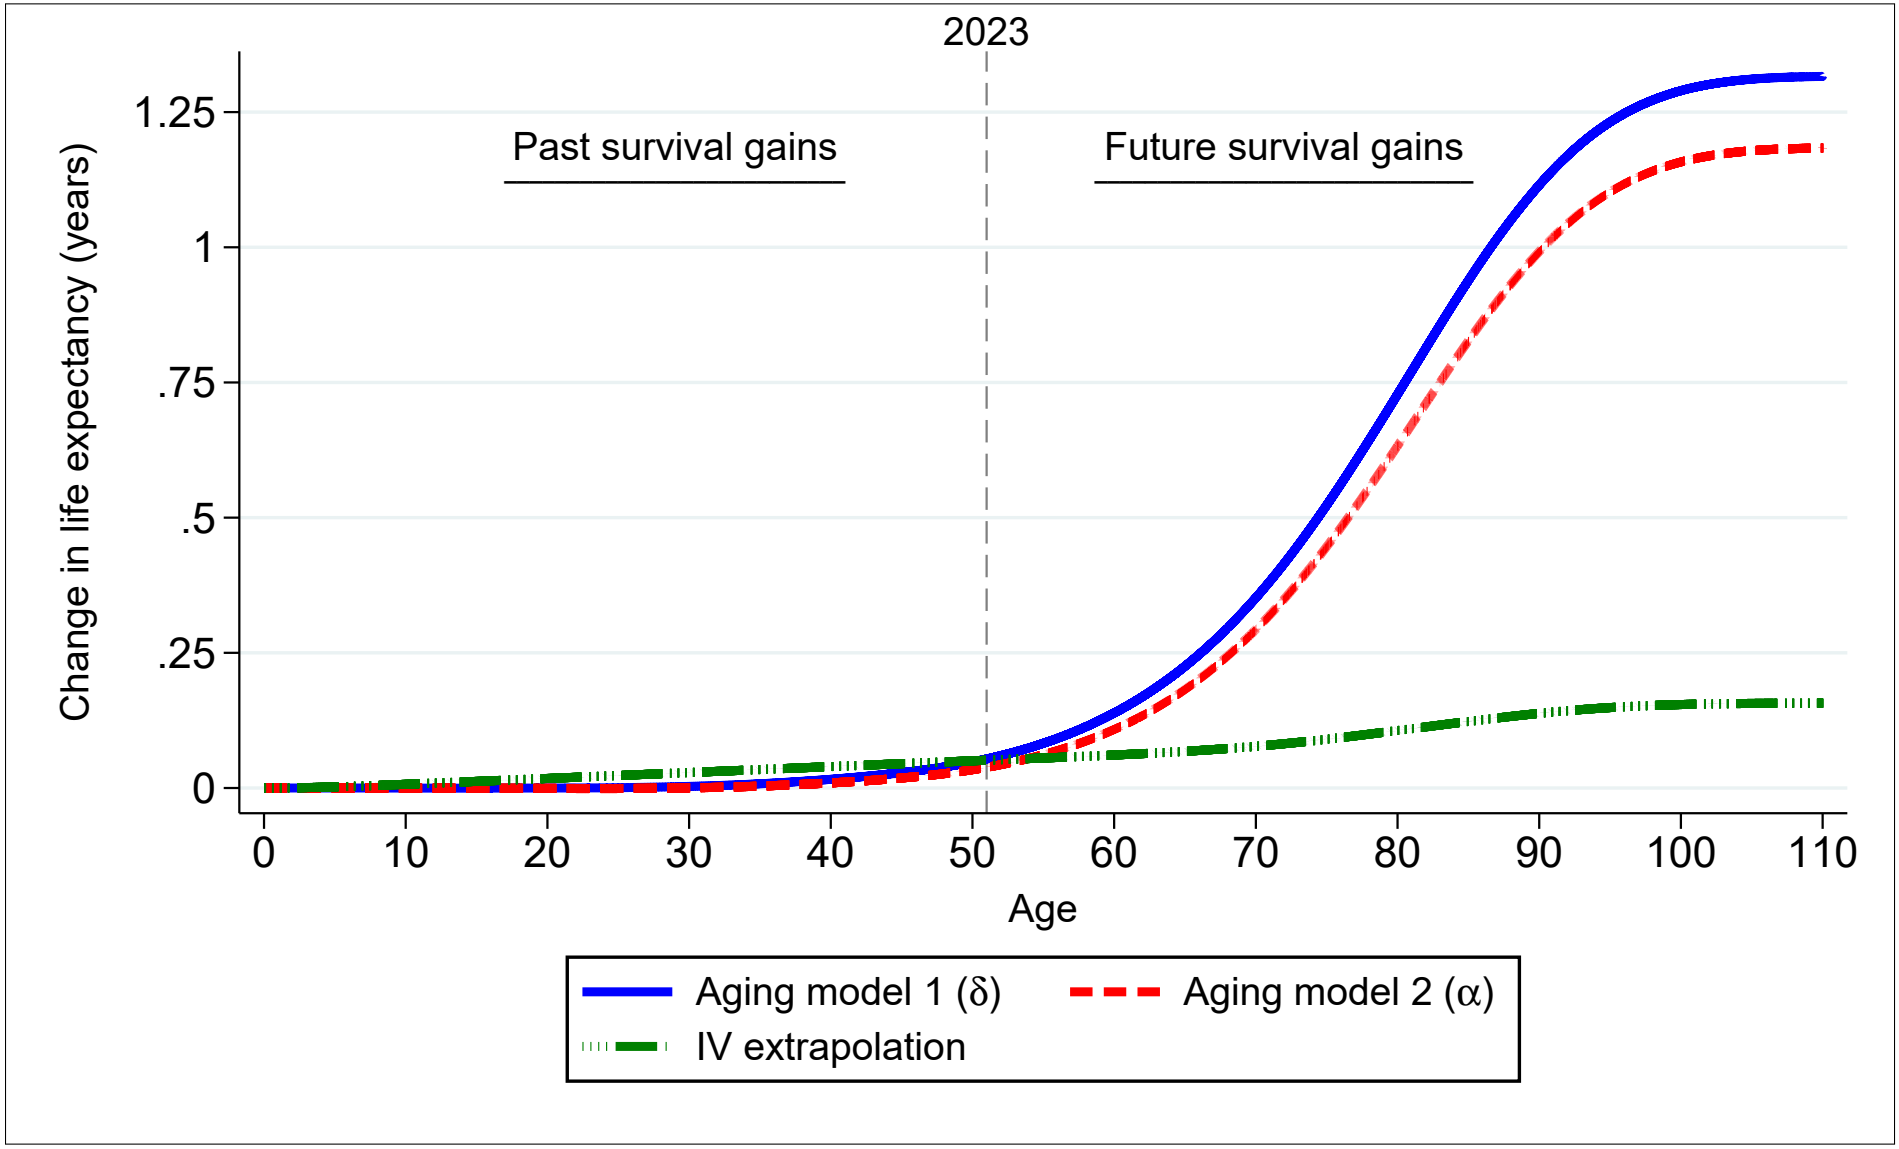

Notes: This figure shows the cumulative effect of a permanent, 1-unit decrease in sulfur dioxide (SO_2) on life expectancy over the life cycle. These projections are produced by the dynamic production model of health (3), which was calibrated using our 1-day IV estimates from Equation (1). The solid blue line (aging model 1) shows projections under the assumption that non-cancer-related pollution deaths are governed by changes in the model parameter α, and the dashed red line (aging model 2) assumes they are governed instead by the parameter δ. The cumulative effect at age 110 equals the total effect on life expectancy for the "Parameter average" reported in Table A.14.

Table 1: Summary statistics, 1972–1988

	(1)	(2)	(3)
	Mean	Std. Dev.	Observations
A. Pollution outcomes			
SO_2, ppb	8.96	12.62	2,032,338
NO_2, ppb	21.25	15.60	792,784
CO, ppm	1.64	1.37	848,067
Ozone, ppb	25.53	13.69	669,261
TSP, $\mu g/m^3$	63.11	40.19	628,932
B. One-day mortality rate outcomes			
All-cause mortality, deaths per million	24.70	24.32	2,032,338
Cardiovascular	12.21	16.04	2,032,338
Cancer	5.15	9.16	2,032,338
Other	5.45	10.02	2,032,338
External	1.89	7.99	2,032,338
All-cause mortality by age group, deaths per million			
Age 1 and under	33.34	166.40	2,032,338
Age 1–19	1.51	11.52	2,032,338
Age 20–44	4.41	18.44	2,032,338
Age 45–59	20.61	48.61	2,032,338
Age 60–64	46.92	129.10	2,032,338
Age 65–69	70.11	170.41	2,032,338
Age 70–74	105.29	241.93	2,032,338
Age 75–79	157.60	360.18	2,032,338
Age 80–84	242.33	575.20	2,032,338
Age 85+	442.67	923.63	2,032,338

Notes: Unit of observation is a county-day. Statistics are unweighted. Sample is restricted to observations where both mortality and sulfur dioxide (SO_2) are non-missing. Mortality is calculated as the number of daily deaths per million individuals. Pollution data are from the Environmental Protection Agency, mortality counts are from the National Vital Statistics, and population estimates are from the Surveillance, Epidemiology, and End Results (SEER) Program.

Table 2: OLS and IV estimates of effect of SO_2 on 1-day mortality

	OLS	IV			
	(1)	(2)	(3)	(4)	(5)
SO_2, ppb	0.0077*	0.084**	0.084**	0.085**	0.098**
	(0.0031)	(0.013)	(0.013)	(0.012)	(0.014)
First-stage F-statistic		42	68	33	32
Mean outcome	25	25	25	25	25
Sample size	2,023,456	2,032,338	2,032,272	2,031,752	2,032,340
Weather controls					
Baseline weather variables	X	X	X	X	
Minimum temperature variables			X	X	
More granular bins				X	

Notes: Dependent variable is number of deaths per million people on the day of exposure. All regressions include county-by-month and month-by-year fixed effects, as well as two leads and two lags of the instruments. Baseline weather variables include bins for maximum temperature, precipitation, and wind speed. The specification in Column (3) additionally includes controls for minimum temperature. The specification in Column (4) increases the number of bins and matches the specification used in Deryugina et al. (2019). Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table 3: IV estimates of effect of SO_2 on 1-day mortality, controlling for other pollutants

	(1)	(2)	(3)	(4)	(5)	(6)
A. All-pollutant sample						
SO_2, ppb	0.084**	0.060**	0.065**	0.066**	0.059**	0.064**
	(0.012)	(0.013)	(0.014)	(0.012)	(0.012)	(0.014)
TSP, $\mu g/m^3$		0.012**	0.014**	0.014**	0.013**	0.015**
		(0.0036)	(0.0037)	(0.0033)	(0.0040)	(0.0035)
NO_2, ppb			–0.014			0.0023
			(0.013)			(0.017)
Ozone, ppb				–0.044*		–0.046*
				(0.021)		(0.022)
CO, ppm					–0.20	–0.24
					(0.17)	(0.20)
First-stage F-statistic	81	21	17	11	20	10
Mean outcome	27	27	27	27	27	27
Sample size	78,946	78,946	78,946	78,946	78,946	78,946
B. SO_2/TSP sample						
SO_2, ppb	0.079**	0.035*				
	(0.014)	(0.015)				
TSP, $\mu g/m^3$		0.019**				
		(0.0045)				
First-stage F-statistic	96	50				
Mean outcome	25	25				
Sample size	627,304	627,304				

Notes: Dependent variable is number of deaths per million people on the day of exposure. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level. Table A.6 presents estimates using a third sample that includes all pollutants except for TSP.

Online Appendix

"The Long-run Effect of Air Pollution on Survival"

Tatyana Deryugina, University of Illinois and NBER
Julian Reif, University of Illinois and NBER

A Supplementary information and analysis

A.1 Data

Monitor-level data for sulfur dioxide (SO_2), total suspended particulates (TSP), nitrogen dioxide (NO_2), ozone, and carbon monoxide (CO) for the years 1972–1988 were obtained by email request from the US Environmental Protection Agency (EPA). Each SO_2 observation provided a sample measure, usually recorded over a period of one hour. We dropped SO_2 observations that had a sample duration greater than 24 hours, or that reported an SO_2 measurement above 1000 parts per billion (ppb) or below –2 ppb.[1] We dropped CO, NO_2, and ozone observations with negative values. All TSP observations had non-negative values. We then aggregated all pollutants to the the monitor-day level, weighting by the time duration of each measure. Finally, data were aggregated to the county-day level by averaging over all monitors within a county.

There were 4,740 active SO_2 pollution monitors during our 1972–1988 sample period, covering 1,041 counties (Figure A.1). Figure A.2 shows trends in air pollution levels over time for different pollutants, and in the fraction of the population living in a county with at least one monitor for that pollutant.

A.2 Model calibration

The dynamic production model of health given by Equation (3) depends on seven parameters: $\{\alpha, \delta, I, \sigma_e, \mu_H, \sigma_H, \underline{\mathrm{H}}\}$. To achieve identification, we normalize two parameters: $\underline{\mathrm{H}}^* = 0$ and $\sigma_H^* = 1$. We calibrate the baseline values for the five remaining parameters using a 1972 period life table, as described in Section 5.2. Let $\{\alpha^*, \delta^*, I^*, \sigma_e^*, \mu_H^*\}$ denote those calibrated values.

We calculate the life expectancy effects of chronic air pollution exposure for two different scenarios. In the first scenario, we assume that pollution exposure only affects the parameters δ and $\underline{\mathrm{H}}$, and denote their post-exposure values as $\{\widetilde{\delta}, \widetilde{\underline{\mathrm{H}}}\}$. In the second scenario, we assume that exposure affects α rather than δ. Below, we describe how we solve for $\{\widetilde{\delta}, \widetilde{\underline{\mathrm{H}}}\}$. It is straightforward to solve for $\widetilde{\alpha}$ instead of $\widetilde{\delta}$ using the methodology described below.

Consider a population of N individuals whose health capital evolves according to Equation (3). Let $\Theta^* = \{\alpha^*, I^*, \mu_H^*, \sigma_H^*, \sigma_e^*\}$ represent the five baseline calibrated parameters that are unaffected by air pollution exposure. Let S be a random-number seed that fixes the initial stock of health capital, H_{i0},

[1] According to the AQS Data Coding Manual version 2.38 (February 2, 2010), the maximum allowable sample value for SO_2 is 1000 ppb. The EPA informed us by email that small negative values can arise due to noise and should be included in sample averages to avoid bias. We chose –2 as the bottom cutoff because it appeared to be the minimum allowable sample value.

and the evolution of the iid shock, ε_{it}, for all individuals. Then we can define $M_t\left(\delta^*, \underline{\mathrm{H}}^* \middle| \Theta^*, N, S\right)$ as the deterministic mortality rate at time $t \geq 0$ for a cohort born in 1972, as computed by the model using the baseline parameters.

Let $\hat{\beta}^k_{a,c}$ denote the IV estimate of the effect of acute exposure on cumulative mortality for age group a from cause of death c in the k days following exposure. Let $[t^0_a, t^1_a]$ define the age interval spanned by age group a, measured in days. Consider a specific age, $t \in [t^0_a, t^1_a]$. To quantify the effect of exposure on model parameters, we first solve for the new mortality threshold, $\widetilde{\underline{\mathrm{H}}}_{at}$, which is defined implicitly by the following equation:

$$\hat{\beta}^k_{a,cancer} = M_t\left(\delta^*, \widetilde{\underline{\mathrm{H}}}_{at} \middle| \Theta^*, N, S\right) - M_t\left(\delta^*, \underline{\mathrm{H}}^* \middle| \Theta^*, N, S\right) \tag{A.1}$$

We then solve for $\widetilde{\delta}$, which is defined implicitly by:

$$\hat{\beta}^k_{a,all} = M_t\left(\widetilde{\delta}_{at}, \widetilde{\underline{\mathrm{H}}}_{at} \middle| \Theta^*, N, S\right) - M_t\left(\delta^*, \widetilde{\underline{\mathrm{H}}}_{at} \middle| \Theta^*, N, S\right) \tag{A.2}$$

Health capital is strictly decreasing in δ for all individuals (see Equation 3), and death occurs when an individual's health stock falls below $\underline{\mathrm{H}}$. Thus, the mortality rate, $M_t(\cdot)$, is monotonically increasing in both δ and $\underline{\mathrm{H}}$. Consequently, the solutions $\{\widetilde{\delta}_{at}, \widetilde{\underline{\mathrm{H}}}_{at}\}$ to Equations (A.1) and (A.2) exist and are unique.[2] The solutions can be computed using standard root-finding algorithms.

We can solve for $\{\widetilde{\delta}_{at}, \widetilde{\underline{\mathrm{H}}}_{at}\}$ for any $t \in [t^0_a, t^1_a]$. We use the approximate integer midpoint of each age bin, $t = round\,[(t^0_a + t^1_a)/2]$.[3] Because health capital depends on an iid health shock, using only a single day for this calibration produces noisy solutions. To improve precision, we solve Equations (A.1) and (A.2) using 50 different days around the midpoint and take the average. For example, for $a = 65$ (the 65–69 age group), we solve for $\{\widetilde{\delta}_{65,t}, \widetilde{\underline{\mathrm{H}}}_{65,t}\}$ using ages $t =$ 68y1d, $t =$ 68y2d, ..., $t =$ 68y50d.[4] Figure A.9 shows the cumulative mortality effects of acute (1-day) exposure to air pollution for 50 separate daily ages from the 65–69 age group, as predicted by the model using the values of $\{\widetilde{\delta}_{65,t}, \widetilde{\underline{\mathrm{H}}}_{65,t}\}$ for each age t. The effect for the first day of each plot is equal to the 1-day IV estimate for the 65–69 age group (see first row of Column (6) in Table A.4). The subsequent estimates show the longer-run effects up to 30 days later. Each value reports mortality in the counterfactual relative to mortality in the baseline case (no exposure). A decrease in the cumulative value indicates mortality displacement. The "own-age prediction" in Figure 5 reports the average of these 50 plots.

[2] Alternatively, one could solve first for $\widetilde{\delta}$ and then for $\widetilde{\underline{\mathrm{H}}}$. That would yield estimates that are numerically very close but not identical to the ones resulting from this method. We solve for $\widetilde{\underline{\mathrm{H}}}$ first because it is more efficient: when solving for $\widetilde{\alpha}$, we do not need to separately solve for $\widetilde{\underline{\mathrm{H}}}$ again as it was already computed when solving for $\widetilde{\delta}$.

[3] For the 85+ age group, we use a midpoint of 90, which is the average age of death in that group during our sample period.

[4] The optimal strategy would employ all $365 \times 5 = 1825$ days in the 5-year age bin, giving more weight to the ages near the midpoint. However, doing so is computationally burdensome and unnecessary for achieving sufficient accuracy.

We compute the age-specific parameter solutions $\{\widetilde{\underline{H}}_a, \widetilde{\delta}_a\}$ as the averages of these 50 solutions:

$$\widetilde{\underline{H}}_a = \frac{1}{50}\sum_t \widetilde{\underline{H}}_{at}$$
$$\widetilde{\delta}_a = \frac{1}{50}\sum_t \widetilde{\delta}_{at}$$

Figure A.8 reports $\{\widetilde{\underline{H}}_a, \widetilde{\delta}_a\}$ for the five age groups over 65, expressed as deviations from the baseline calibrated values (i.e., $\widetilde{\underline{H}}_a - \underline{H}^* = \widetilde{\underline{H}}_a$ and $\widetilde{\delta}_a - \delta^*$). The parameters used to predict the long-run survival effects of air pollution are computed as the average across these five age groups:

$$\widetilde{\underline{H}} = \frac{1}{5}\sum_a \widetilde{\underline{H}}_a$$
$$\widetilde{\delta} = \frac{1}{5}\sum_a \widetilde{\delta}_a$$

To account for econometric uncertainty in the IV estimates, we use a resampling-based methodology. We randomly draw an estimate of the effect of acute pollution exposure on 1-day mortality from a normal distribution with a mean and standard deviation set equal to the mean and standard error of $\widehat{\beta}^1_{a,c}$ in Equation (1), and then calibrate the change in model parameters to match the mortality change draw. We repeat this exercise 100 times and report the 5th and 95th percentiles of the resulting distribution of model parameter estimates. To reduce the computational burden, we use 20 different days rather than 50 when computing $\{\widetilde{\underline{H}}_a, \widetilde{\delta}_a\}$ during the resampling.

Figure A.1: Locations of the 50 geographic groups and SO_2 monitors

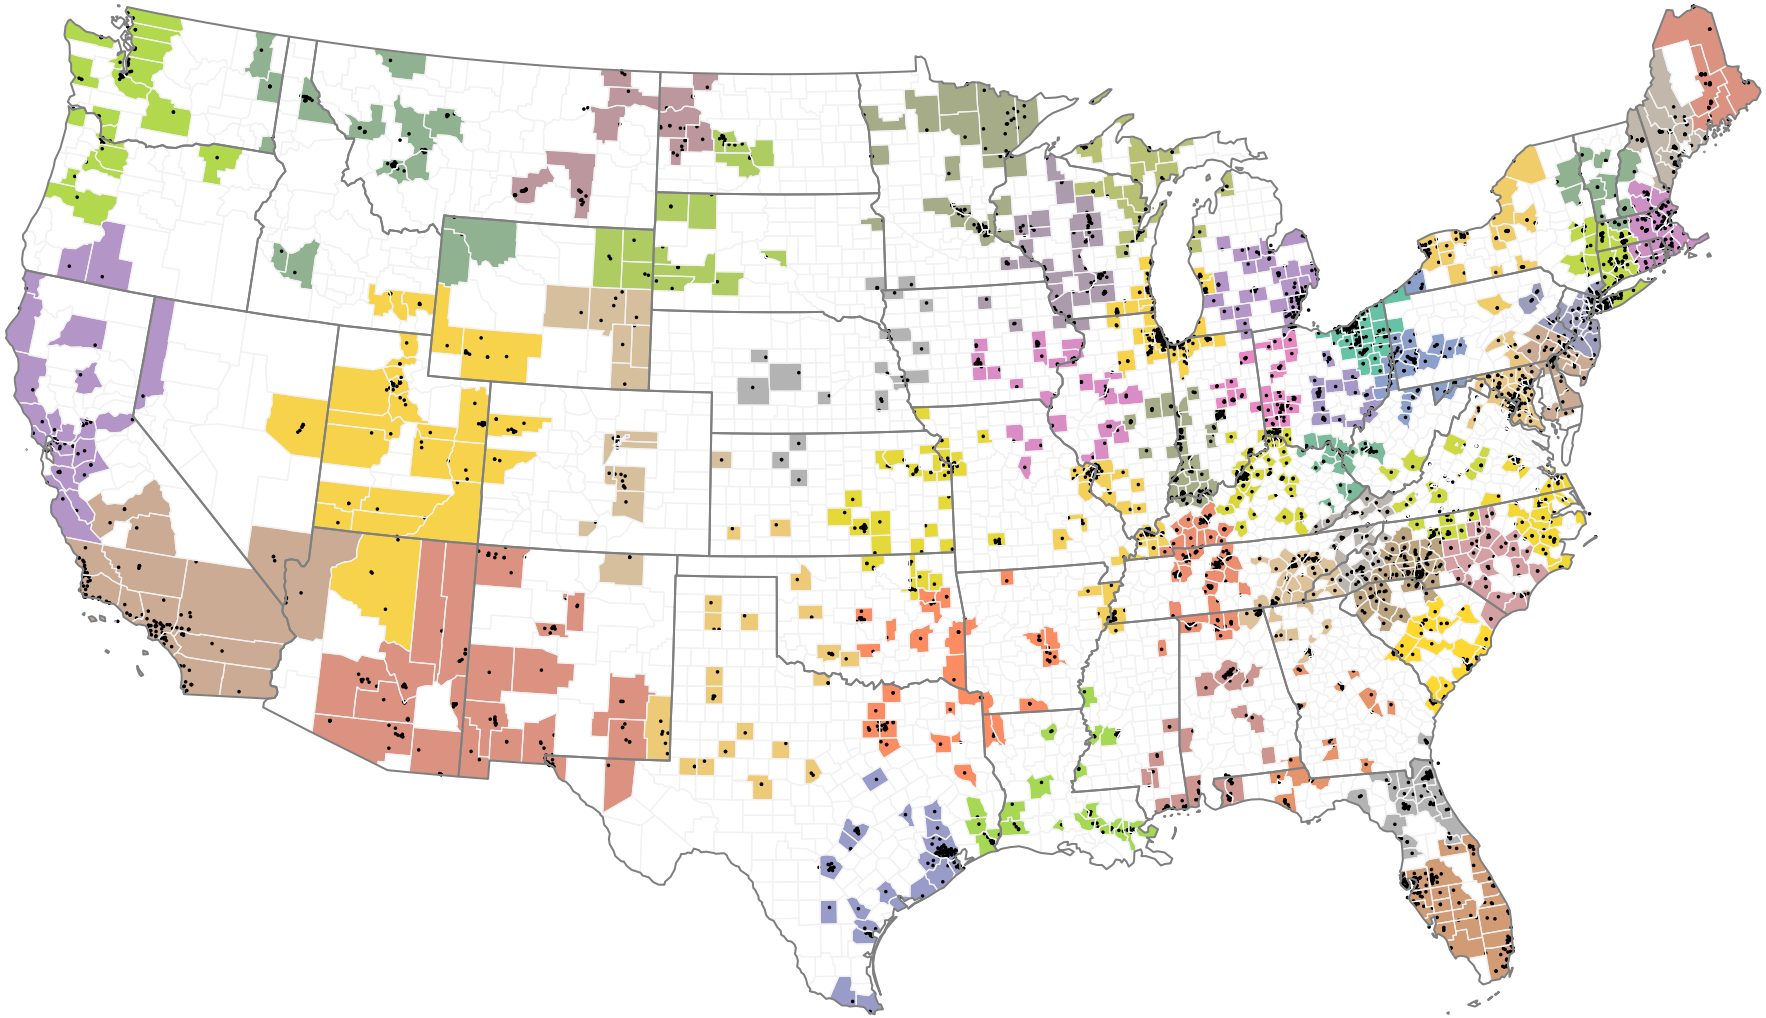

Notes: This map shows the 50 geographic groups included in our main estimation sample. Each group is shaded in a different color. The black dots represent the locations of the SO_2 monitors. Unshaded (white) counties have no SO_2 monitors and are not included in our sample. The Southern California and Greater Philadelphia groups are shown in detail in Figure 1.

Figure A.2: Air pollution means and population coverage levels, 1972–2014

(a) SO_2

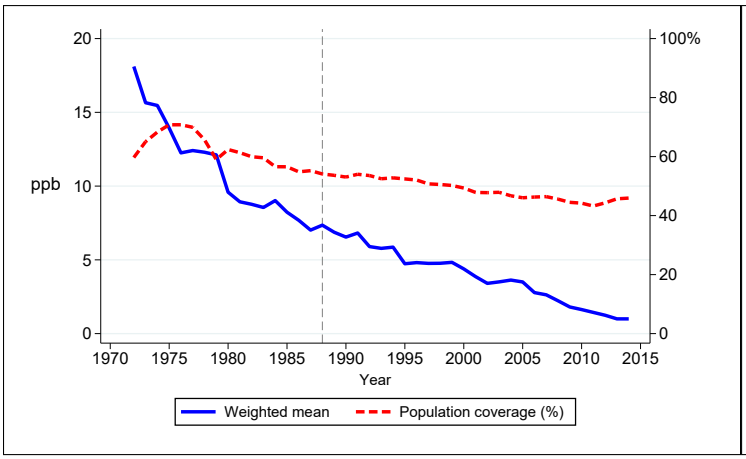

(b) CO

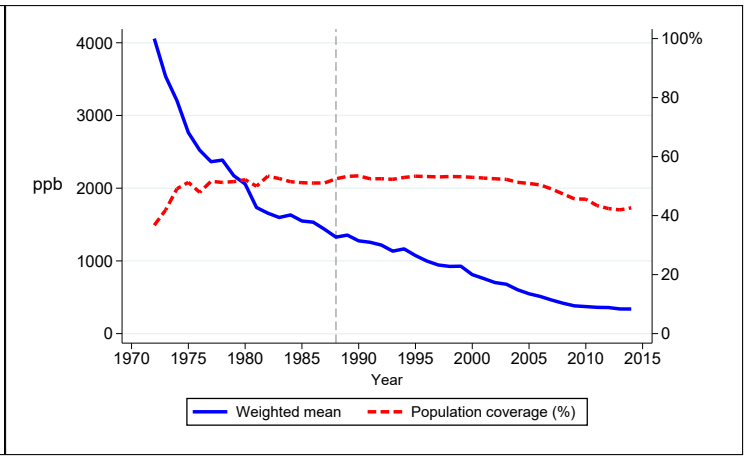

(c) NO_2

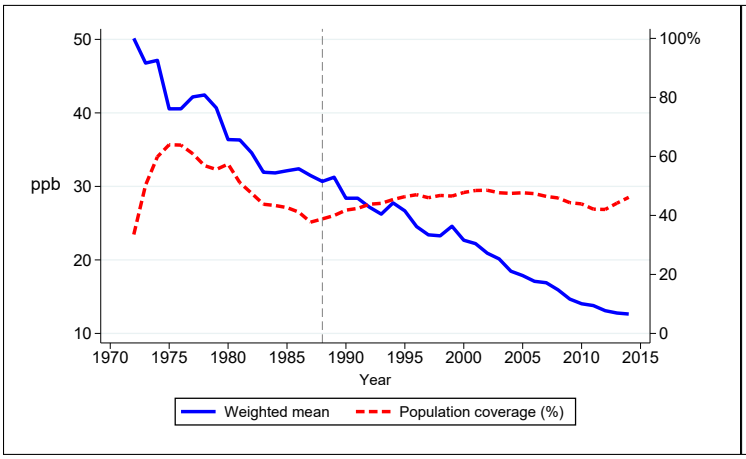

(d) O_3

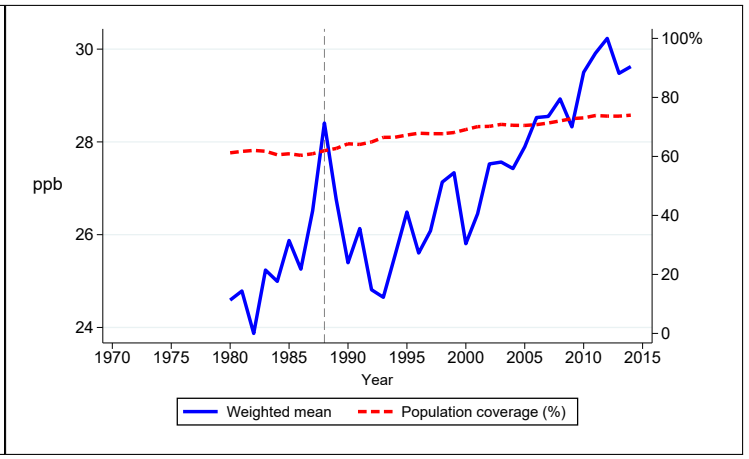

(e) TSP

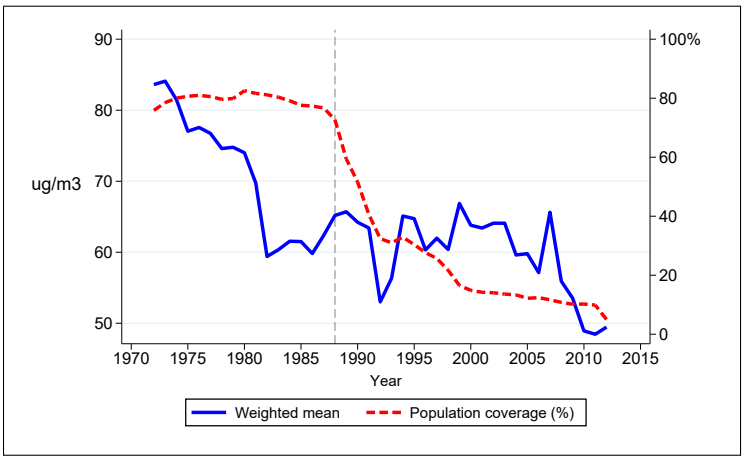

Notes: The solid blue line reports population-weighted pollution levels for all US counties with a monitor for the pollutant that operated for at least one day. The dashed red line reports the percentage of the US population residing in a county with a monitor for the pollutant that operated for at least one day. The dashed vertical line indicates 1988, the end of the sample period used in this study. Data are obtained from the EPA Air Quality database. SO_2, CO, NO_2, and O_3 are measured in parts per billion (ppb). Total suspended particulates (TSP) is measured in micrograms per cubic meter ($\mu g/m^3$).

Figure A.3: Trends in United States mortality rates, 1972–1988

(a) By age group

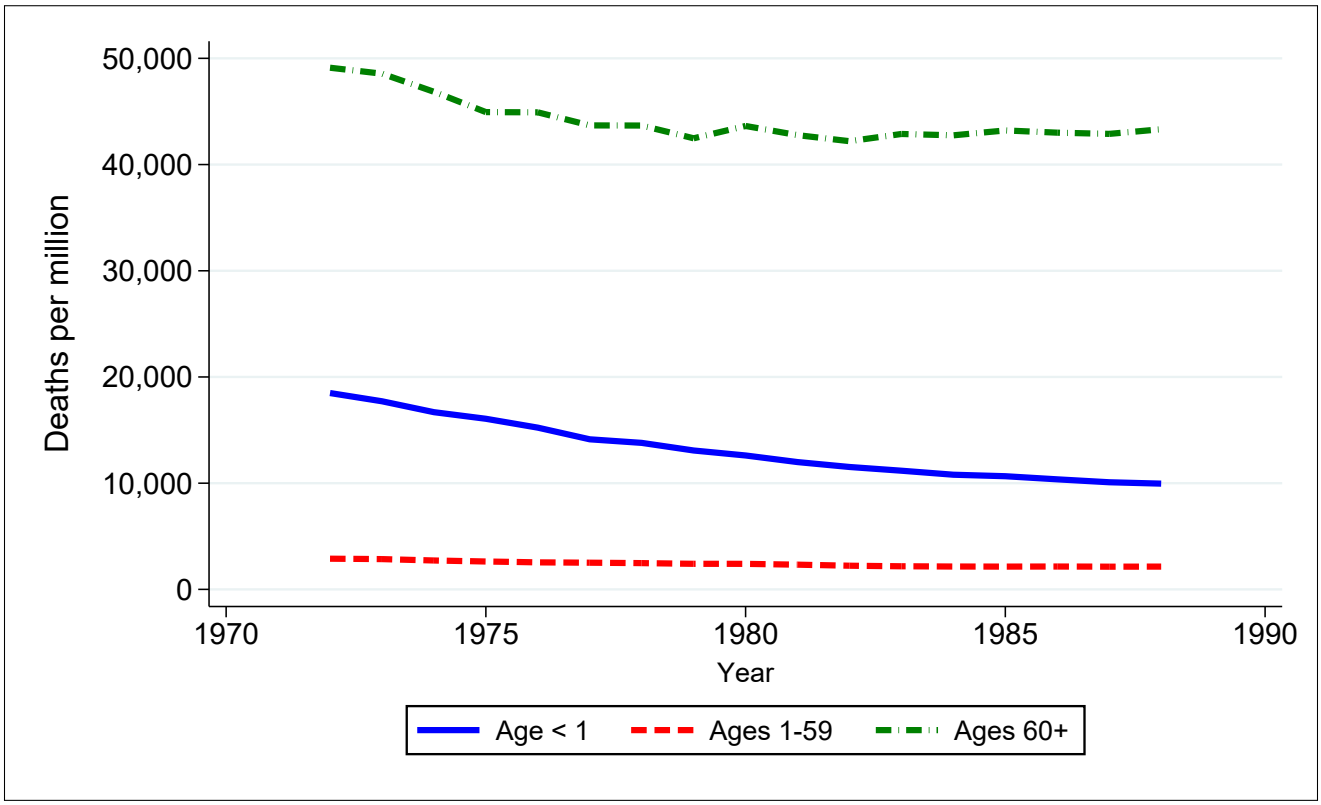

(b) By cause of death

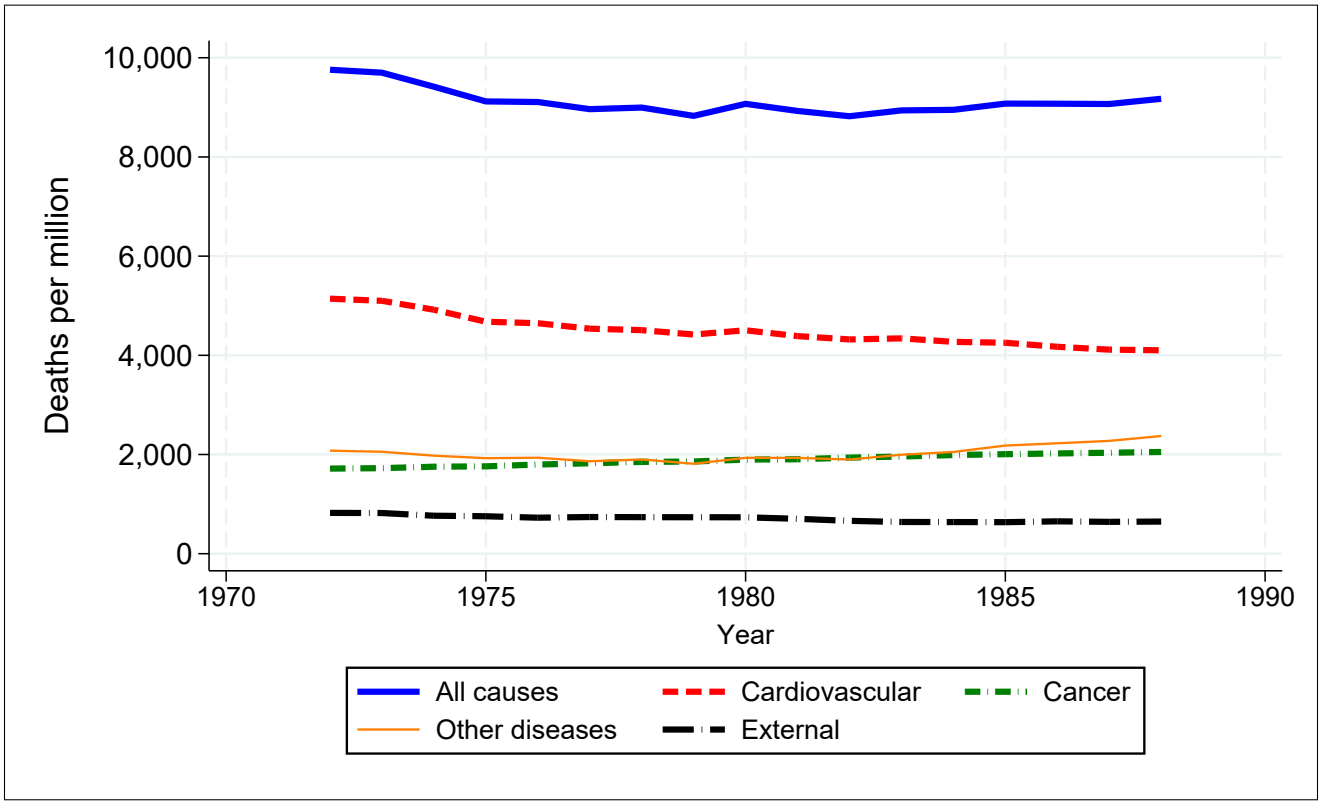

Notes: These two figures report annual mortality rates for the US population. These rates are calculated using mortality data from the National Vital Statistics and population data from SEER. Annual mortality rates are approximately 365 times larger than the daily mortality rates used in the analysis.

Figure A.4: A lot of air pollution originates from distant sources, especially in the East

(a) $PM_{2.5}$ **(b)** Sulfates

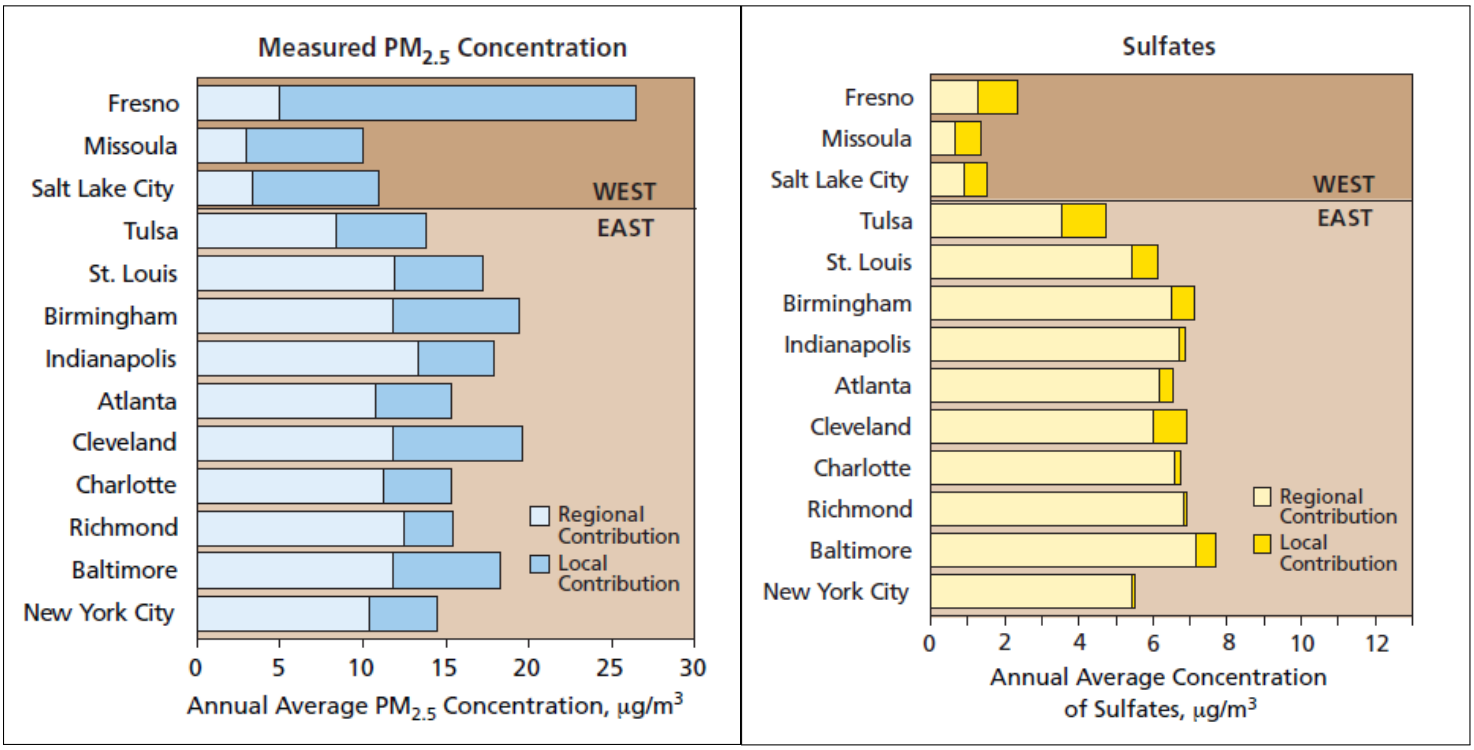

Notes: This figure, reproduced from EPA (2004), shows that a large fraction of measured $PM_{2.5}$ and sulfates does not originate from local sources. Sulfates, which are an important component of $PM_{2.5}$, are formed from the atmospheric transformation of sulfur dioxide.

Figure A.5: Size of first stage by location

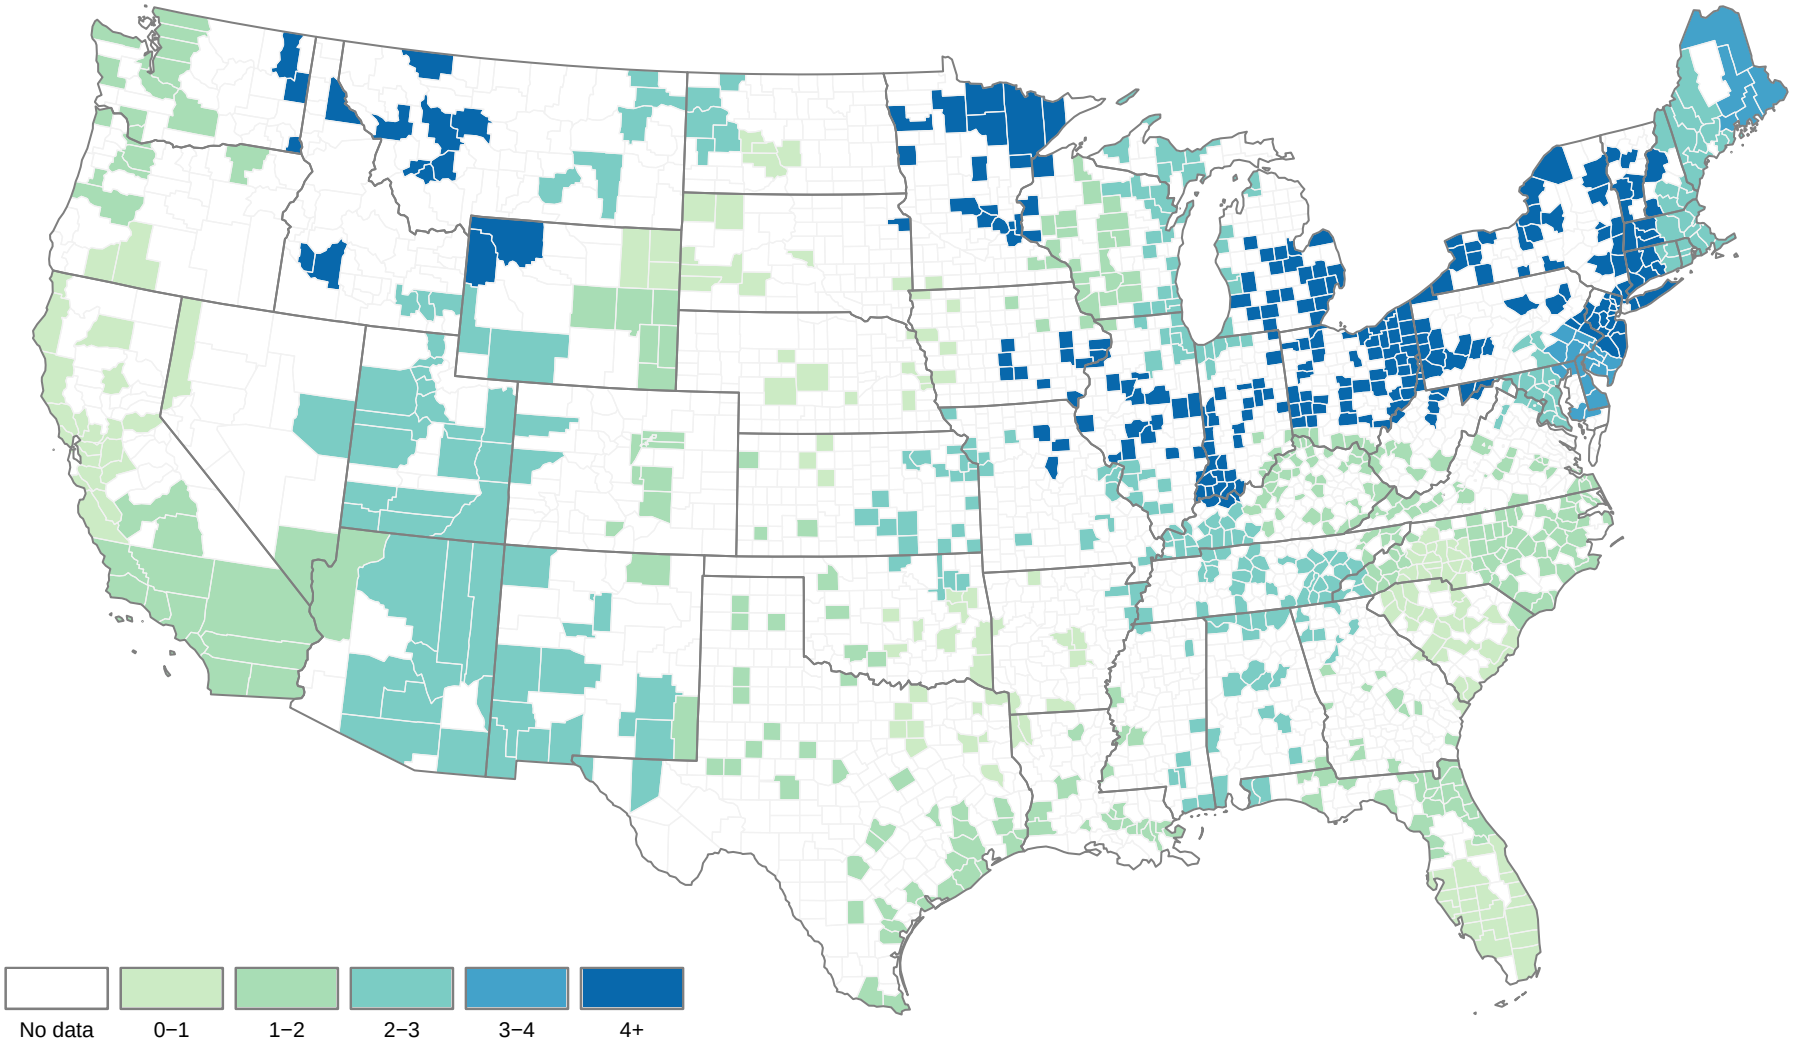

Notes: This map shows the size of the first stage for each of the 50 geographic groups included in our main estimation sample. The size of the first stage is calculated as the difference in SO_2 changes (in parts per billion) between the most and least polluting wind directions, as determined by $\widehat{\gamma_g^1}\sin(\theta) + \widehat{\gamma_g^2}\sin(\theta/2)$, estimated using Equation (2), for $\theta \in [0, 2\pi)$. Unshaded (white) counties have no SO_2 monitors and are not included in our sample.

Figure A.6: IV estimates of effect of acute (1-day) SO_2 exposure on cancer-related mortality

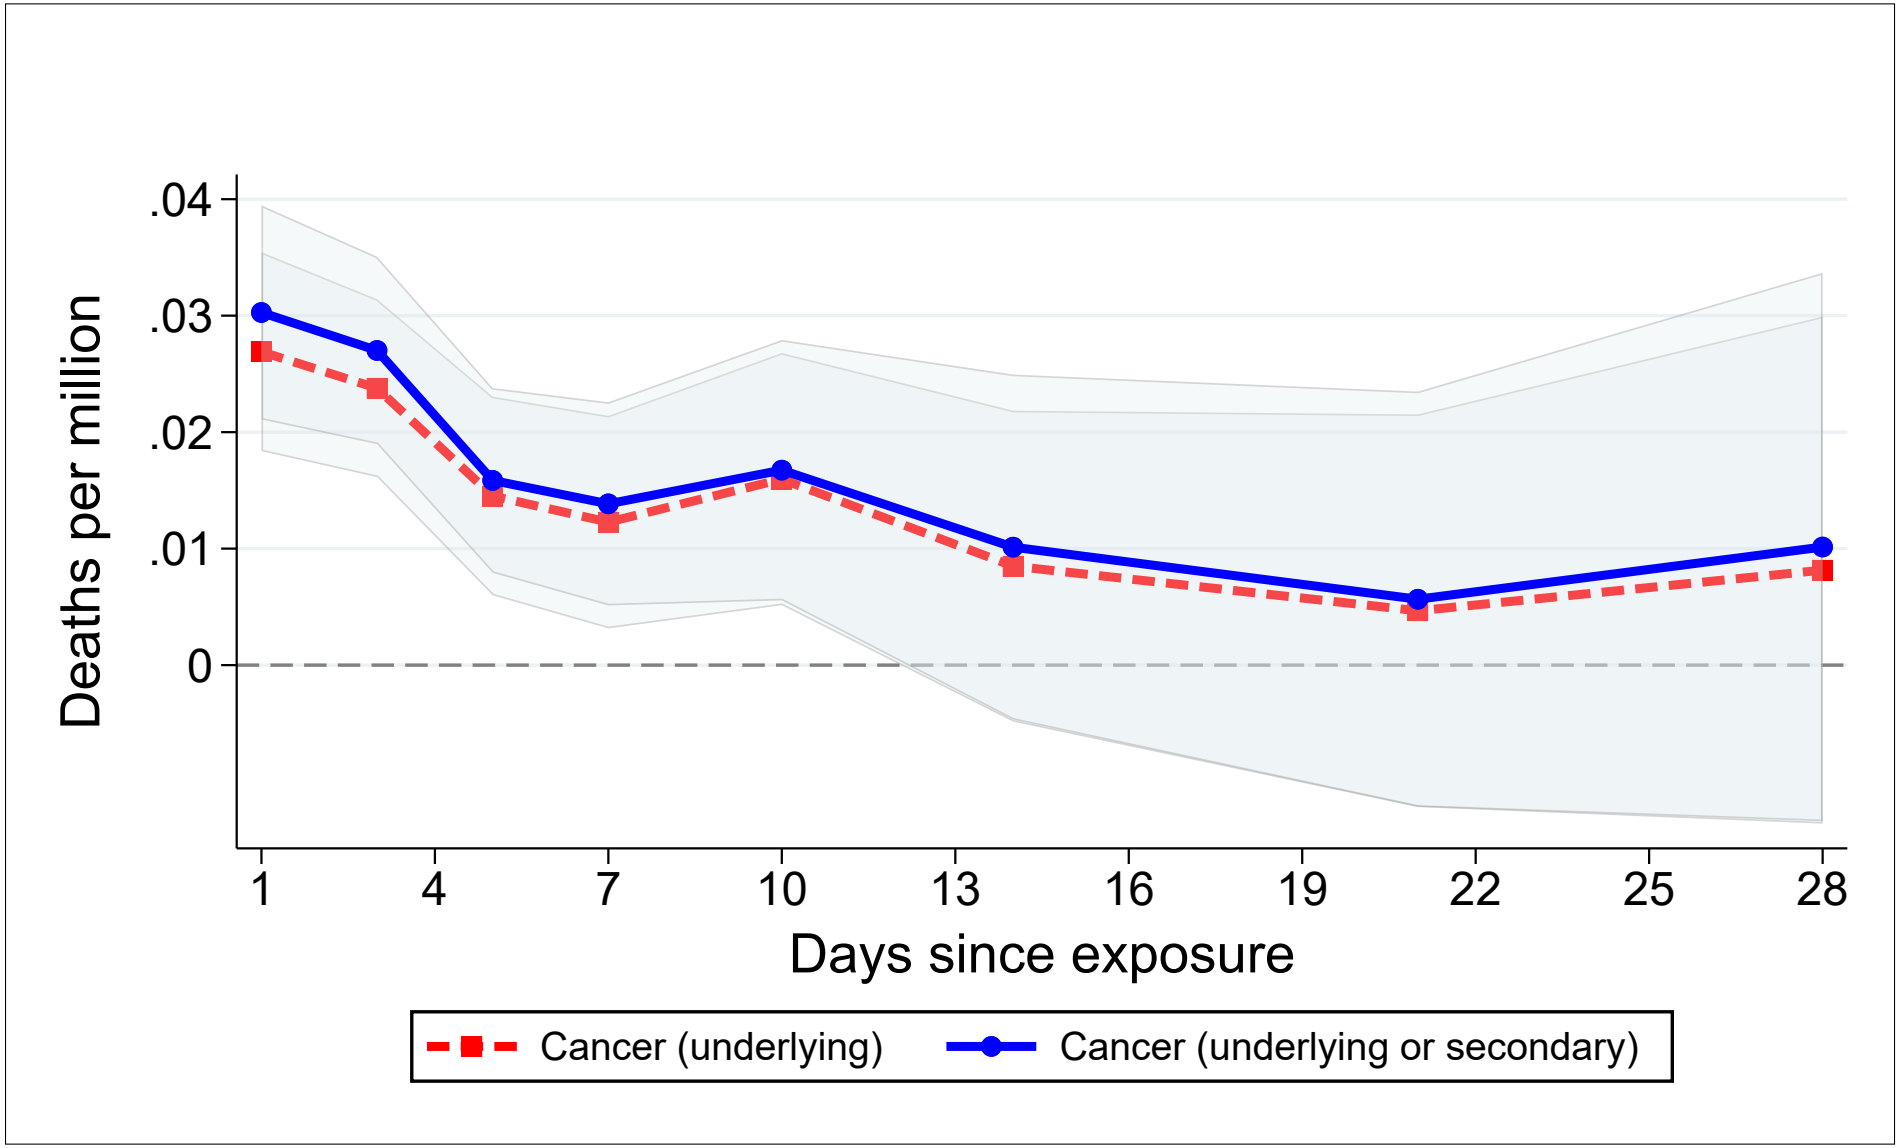

Notes: Each point reports an IV estimate from Equation (1) of the effect of acute (1-day) sulfur dioxide (SO_2) exposure on cancer-related mortality (deaths per million). "Cancer (underlying)", which replicates the cancer estimates shown in Figure 4, includes only deaths where cancer is listed as the underlying cause of death on the death certificate. "Cancer (underlying or secondary)" includes deaths where cancer is listed as the underlying cause of death or as one of the secondary causes. Mortality is measured with a time window ranging from 1 to 28 days, as indicated by the x-axis. Shaded areas correspond to 95% confidence intervals. Estimates are also reported in Table A.5. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors are clustered by county.

Figure A.7: Baseline calibration of the dynamic production model of health

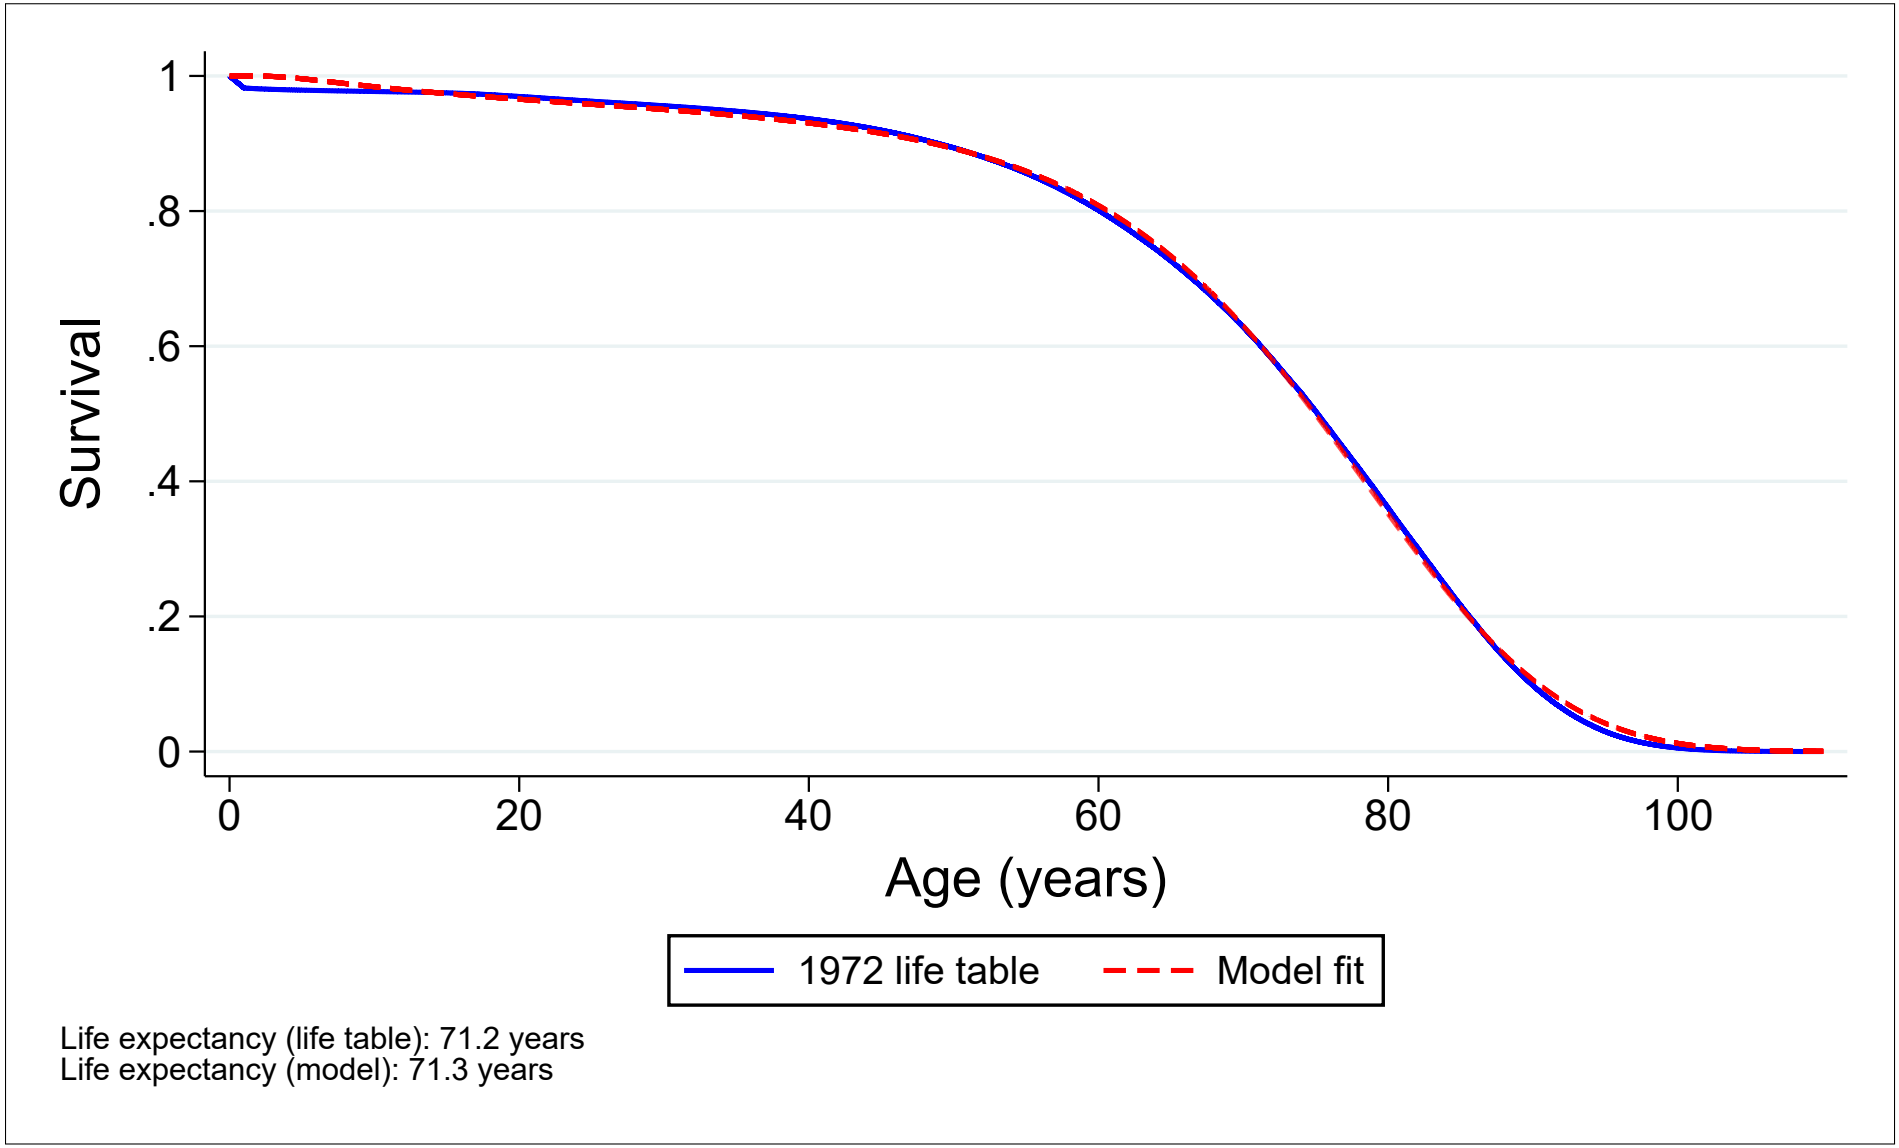

Notes: The solid blue line depicts the survival curve derived from the 1972 period life table for the United States. The dashed red line reports the predicted survival from our dynamic production model of health (3), which was calibrated using these 1972 data. The calibrated model parameters are reported in Column (2) of Table A.13.

Figure A.8: The effect of acute (1-day) SO_2 exposure on model parameters

(a) Increase in α ($\widetilde{\alpha}_a - \alpha^*$)

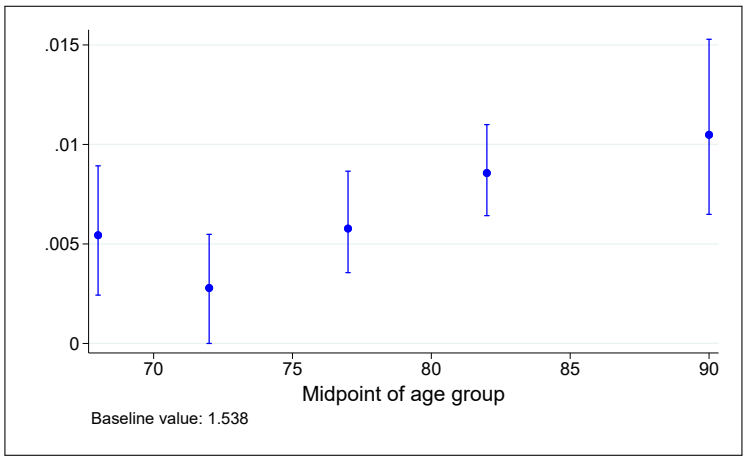

(b) Increase in δ ($\widetilde{\delta}_a - \delta^*$)

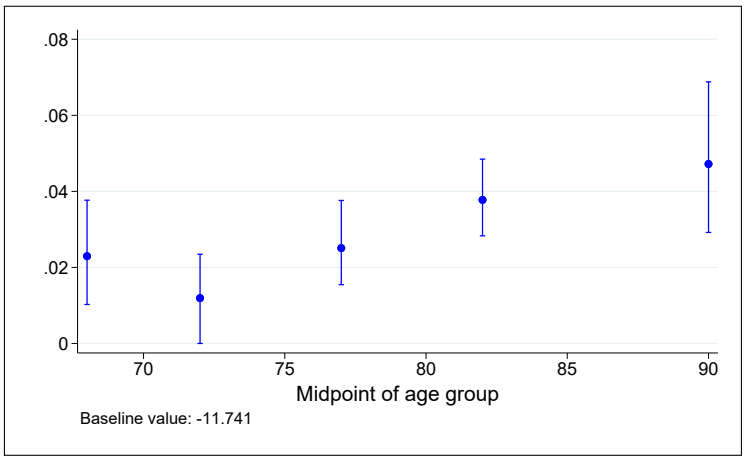

(c) Increase in $\underline{\mathrm{H}}$ ($\widetilde{\underline{\mathrm{H}}}_a$)

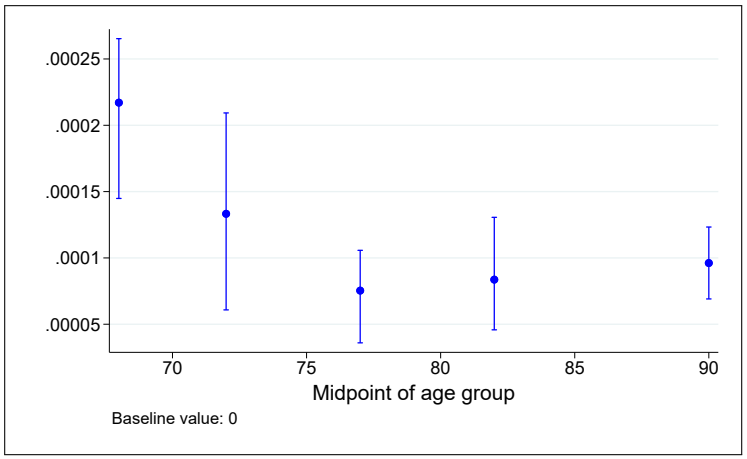

Notes: Panel (a) reports the effect of acute (1-day) pollution exposure on the change in the value of the parameter α from the health model (3), under the assumption that all non-cancer-related deaths operate through changes in α and all other deaths operate through changes in $\underline{\mathrm{H}}$. Panel (b) reports estimates under the alternative assumption that all non-cancer-related deaths operate through changes in the aging parameter δ rather than α. Panel (c) reports the effect on the change in the value of the parameter $\underline{\mathrm{H}}$, which governs mortality displacement and is calibrated using cancer-related deaths. The vertical lines report the 5th and 95th percentiles from a set of 100 bootstrap replications. The bootstrap replications account for econometric uncertainty in the IV estimation of β in (1) and for model uncertainty produced by the error term in the health model (3). Baseline values for all model parameters are reported in Table A.13.

Figure A.9: Predicted effects of acute (1-day) SO_2 exposure on cumulative mortality, for selected ages

Notes: Each plot corresponds to a different age (68 years 1 day, 68 years 2 days, etc.) from the dynamic production model of health (3). Plots show the predicted effect of acute exposure on cumulative mortality (deaths per million) up to 30 days following exposure, using the calibration methodology described in Section 5. The first point of each plot (x = 1 day) is equal to the 1-day IV estimate for the 65–69 age group (Figure 3a). The subsequent values are model predictions. The average of these 50 plots corresponds to the thick red dashed line ("age-specific calibration") shown in Figure 5.

Figure A.10: Comparison of model predictions to IV estimates, by age bin

(a) Ages 65–69

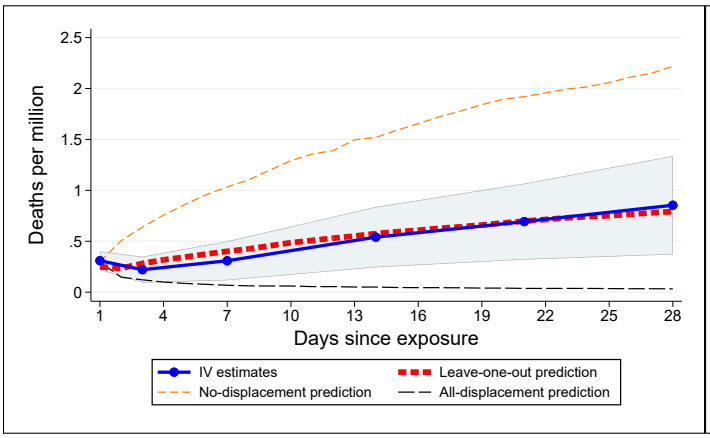

(b) Ages 70–74

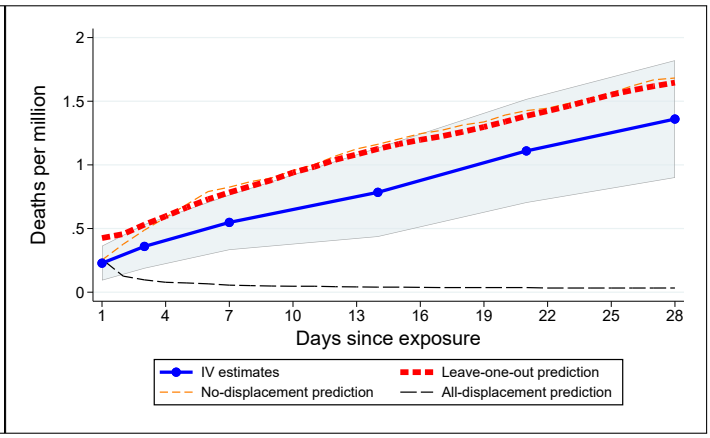

(c) Ages 75–79

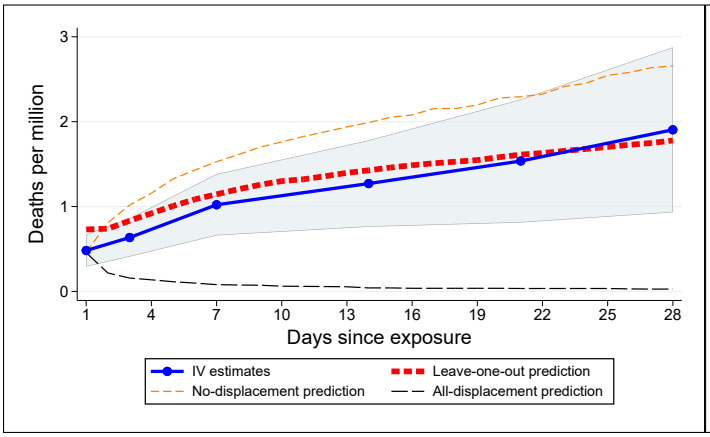

(d) Ages 80–84

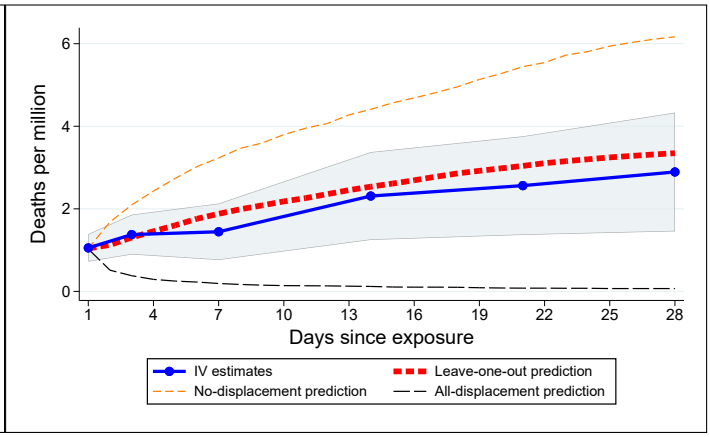

(e) Ages 85+

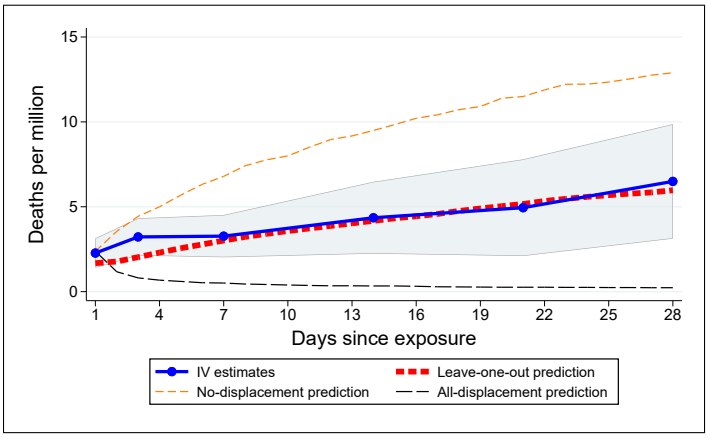

Notes: The red dashed line in panel (b) (ages 65–69), which is reproduced in Figure 5, reports the cumulative average of the 50 plots from Figure A.9.

Figure A.11: The relationship between wind direction and SO_2 concentrations by monitor group

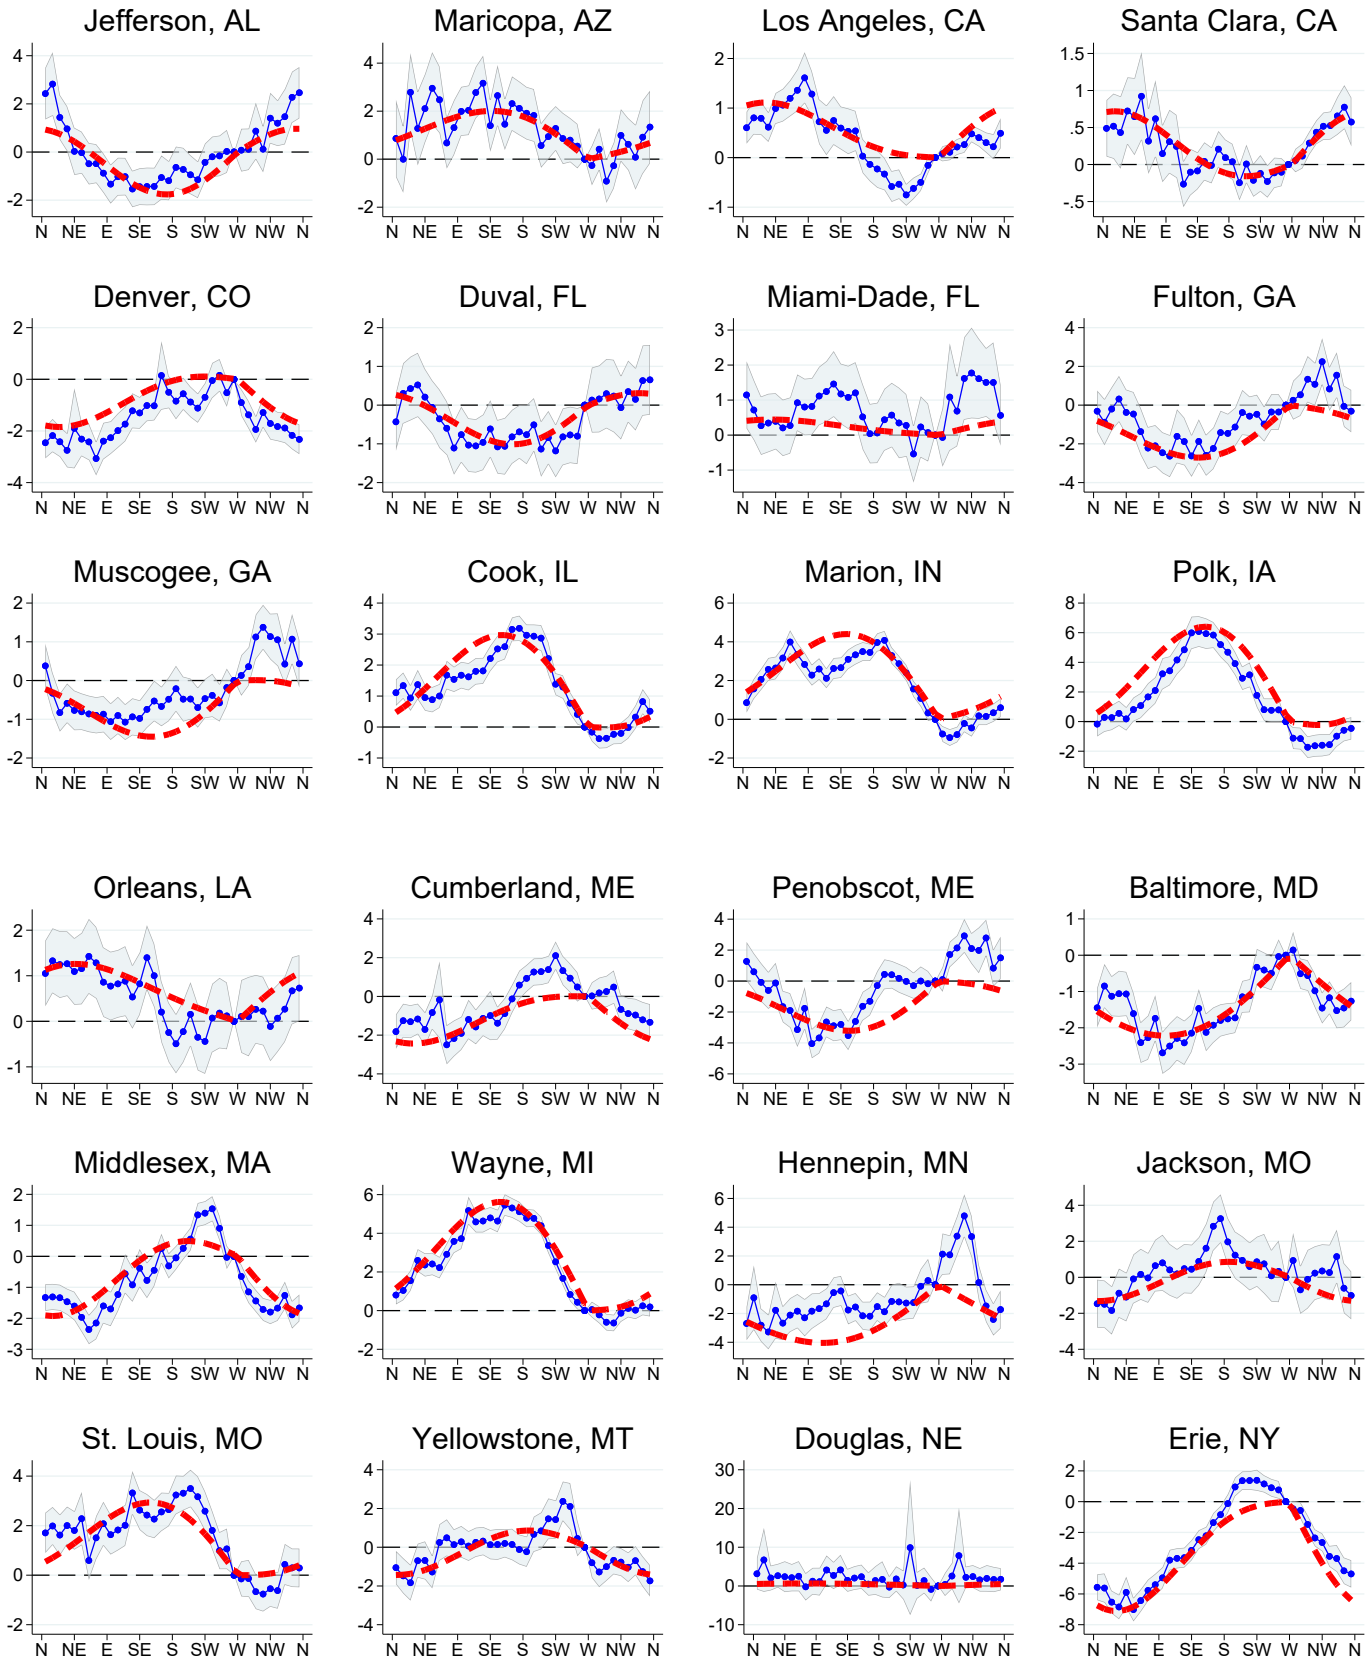

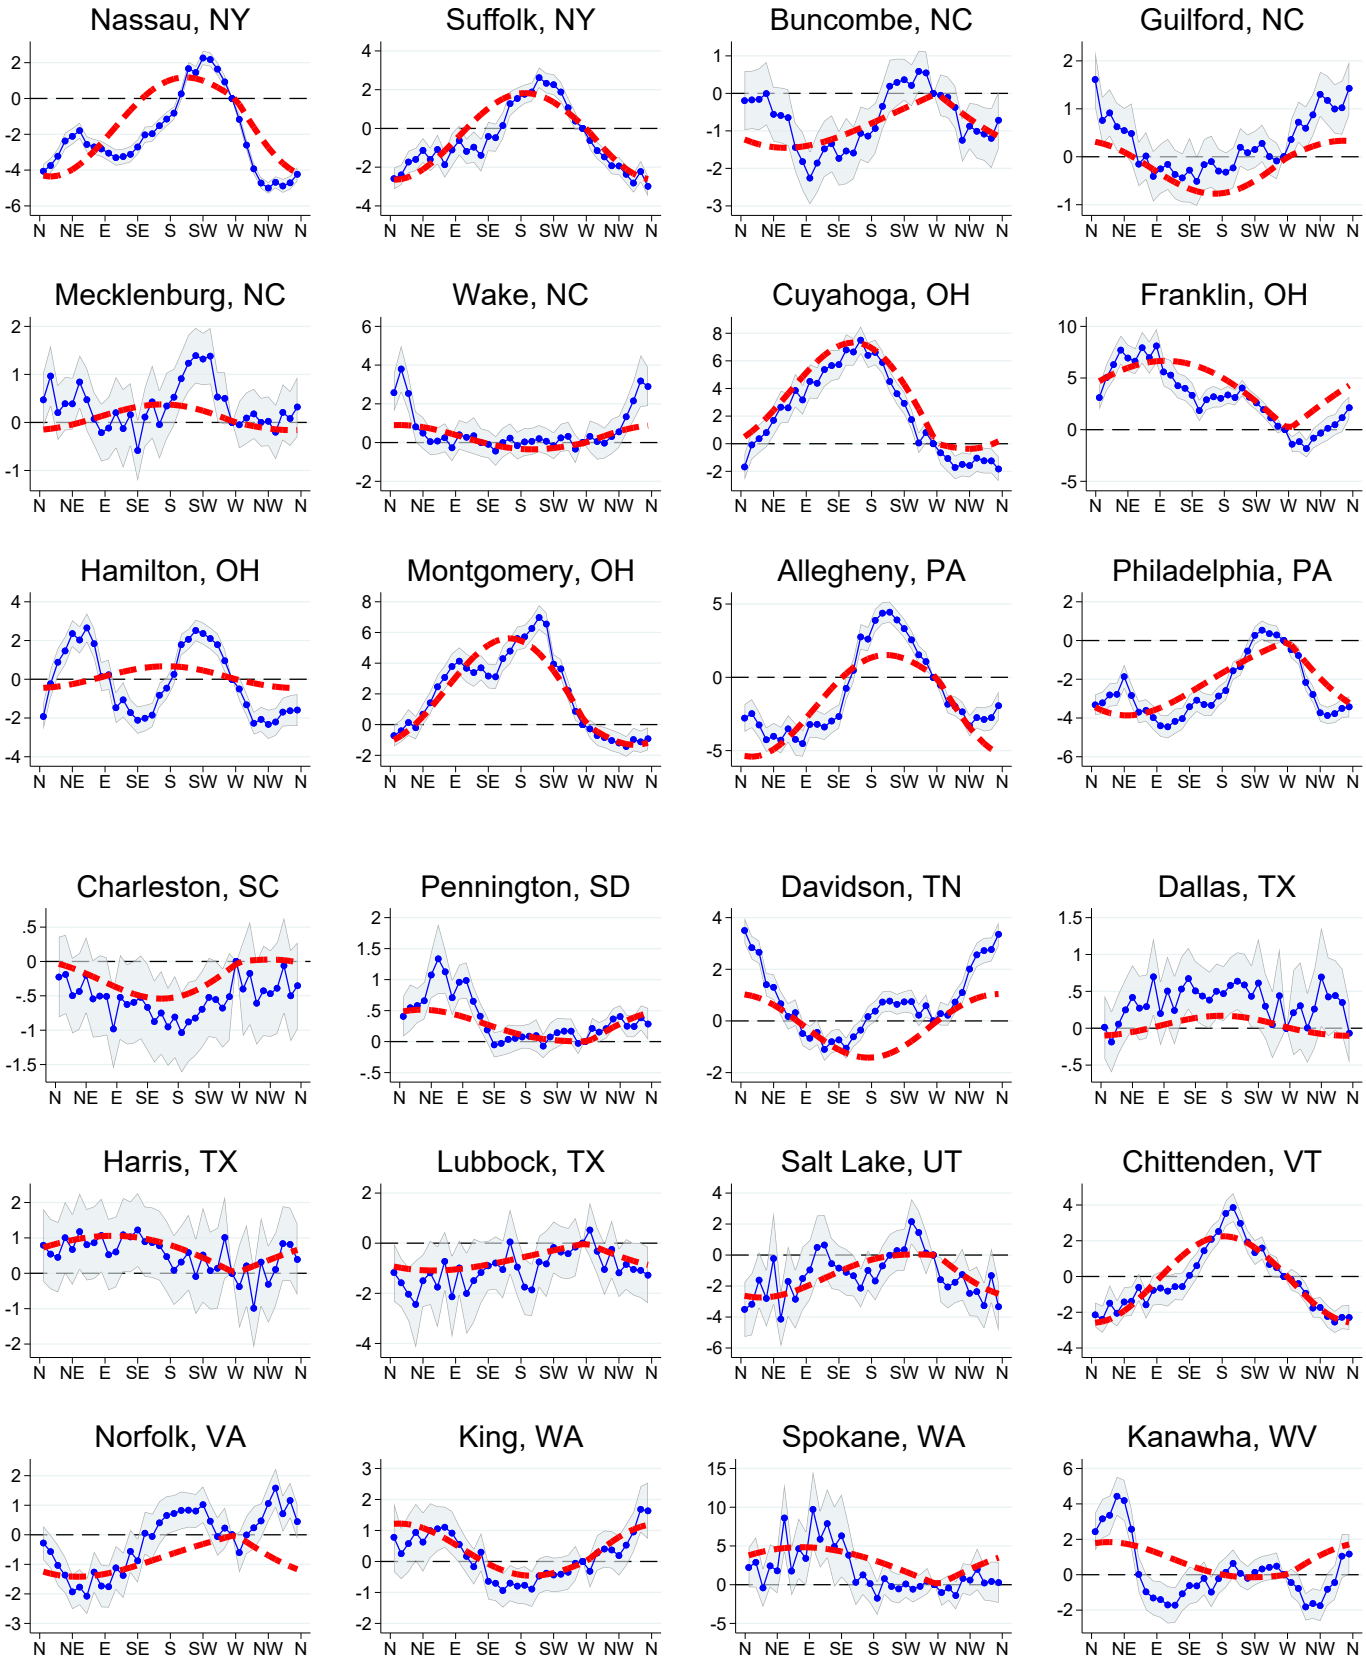
Nassau, NY
Suffolk, NY
Buncombe, NC
Guilford, NC
Mecklenburg, NC
Wake, NC
Cuyahoga, OH
Franklin, OH
Hamilton, OH
Montgomery, OH
Allegheny, PA
Philadelphia, PA
Charleston, SC
Pennington, SD
Davidson, TN
Dallas, TX
Harris, TX
Lubbock, TX
Salt Lake, UT
Chittenden, VT
Norfolk, VA
King, WA
Spokane, WA
Kanawha, WV
N NE E SE S SW W NW N

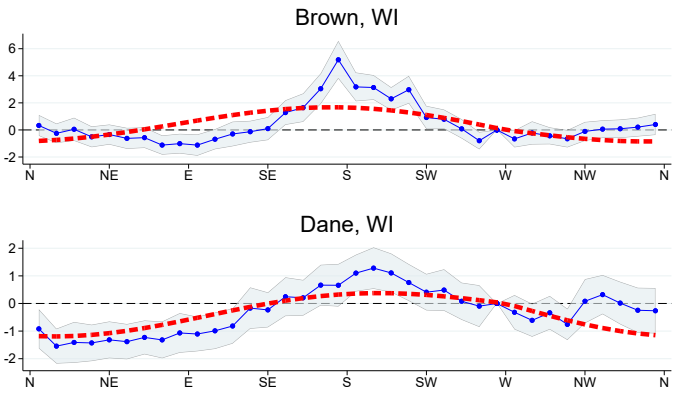

Notes: The graphs plot group-specific non-parametric coefficients (blue points) and 95% confidence intervals (blue shaded area) from a first-stage regression of SO_2 on wind direction measured in 10-degree angle bins. The red dashed lines report the parametric sine fit from Equation (2). The x-axis on the graphs reports direction from which the wind is blowing ("N" = North, "NE" = Northeast, etc.). The y-axis reports the corresponding relative changes in SO_2 (parts per billion). Regressions include county-by-month and month-by-year fixed effects and flexible weather controls. Standard errors are robust to heteroskedasticity. The plots for "Los Angeles, CA" and "Baltimore, MD" are reproduced as "Greater Philadelphia area" and "Southern California area" in Figure 1.

Figure A.12: IV estimates of effect of acute (1-day) SO_2 exposure on mortality, by detailed cause of death

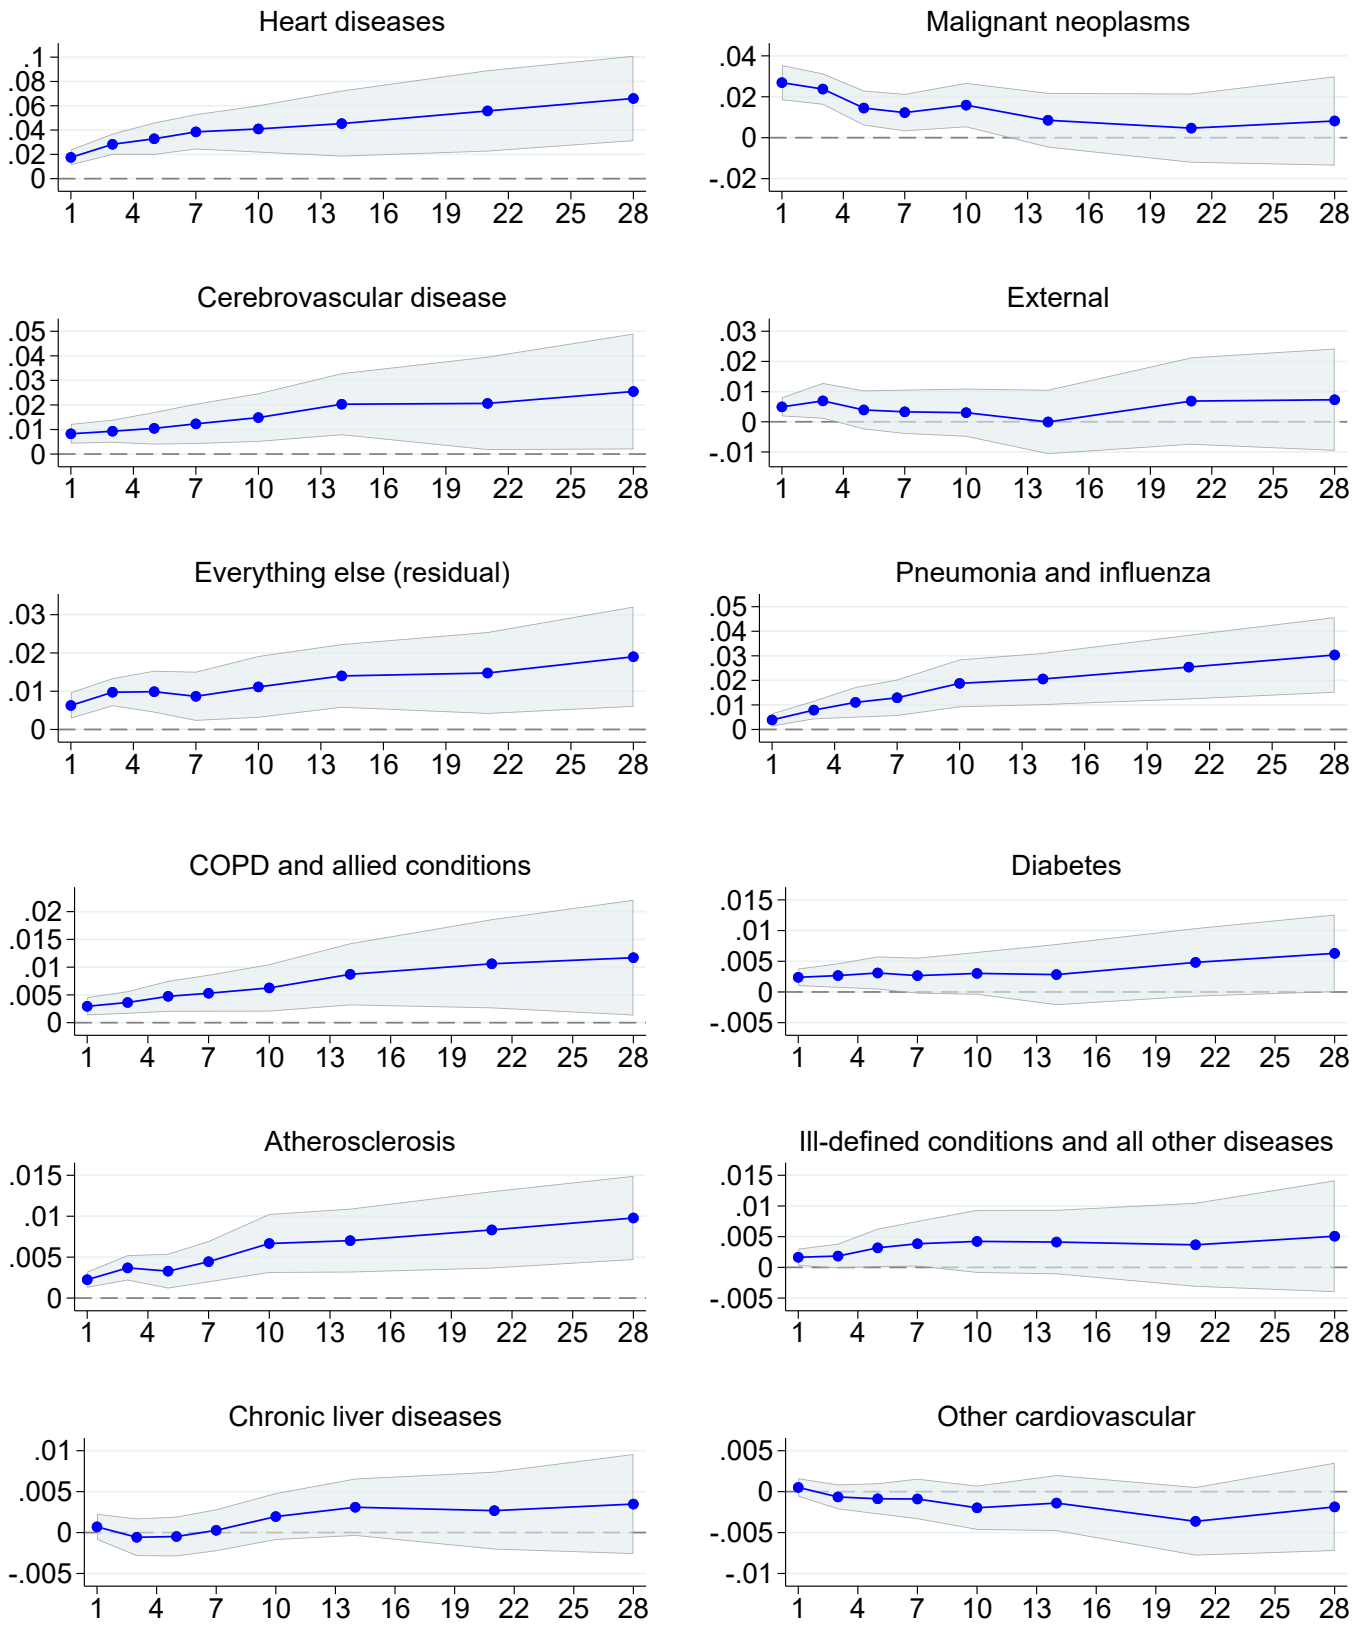

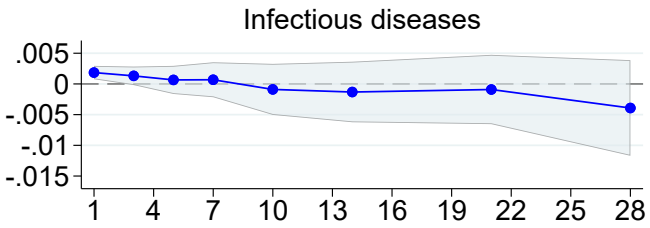
Infectious diseases
.005
0
-.005
-.01
-.015
1 4 7 10 13 16 19 22 25 28

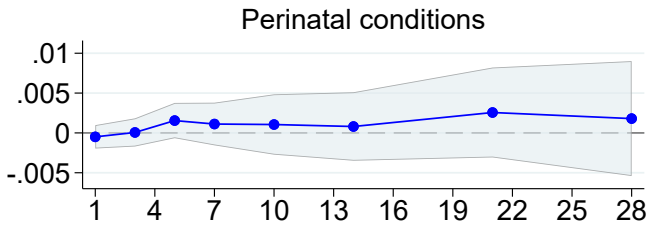
Perinatal conditions
.01
.005
0
-.005
1 4 7 10 13 16 19 22 25 28

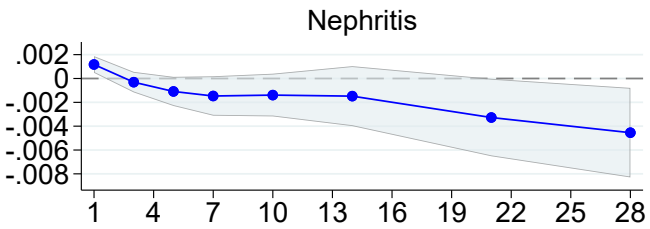
Nephritis
.002
0
-.002
-.004
-.006
-.008
1 4 7 10 13 16 19 22 25 28

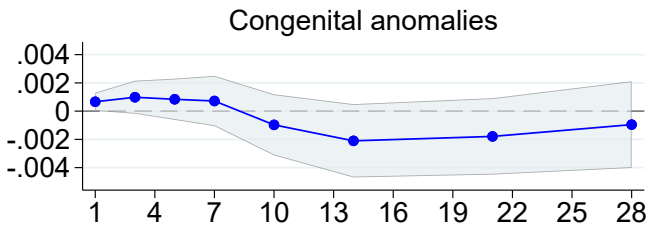
Congenital anomalies
.004
.002
0
-.002
-.004
1 4 7 10 13 16 19 22 25 28

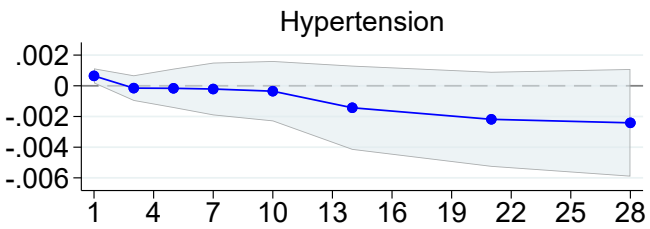
Hypertension
.002
0
-.002
-.004
-.006
1 4 7 10 13 16 19 22 25 28

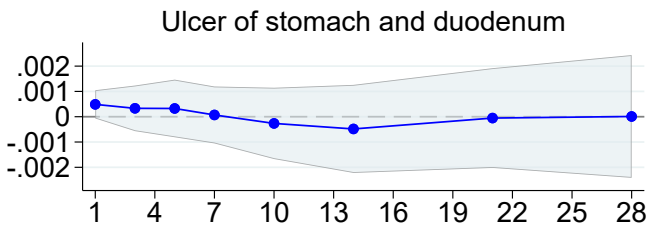
Ulcer of stomach and duodenum
.002
.001
0
-.001
-.002
1 4 7 10 13 16 19 22 25 28

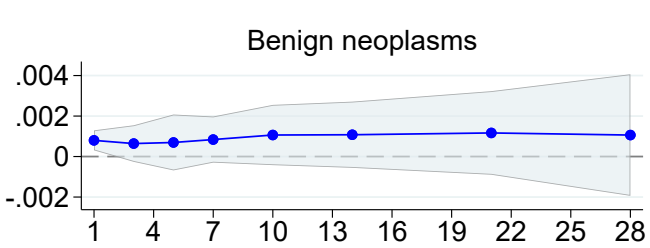
Benign neoplasms
.004
.002
0
-.002
1 4 7 10 13 16 19 22 25 28

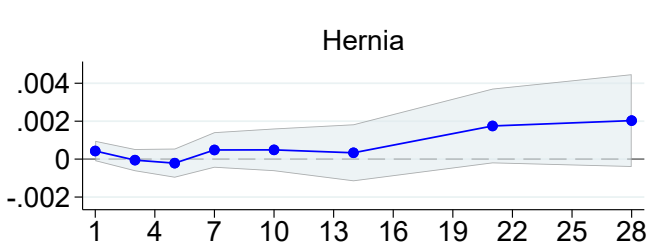
Hernia
.004
.002
0
-.002
1 4 7 10 13 16 19 22 25 28

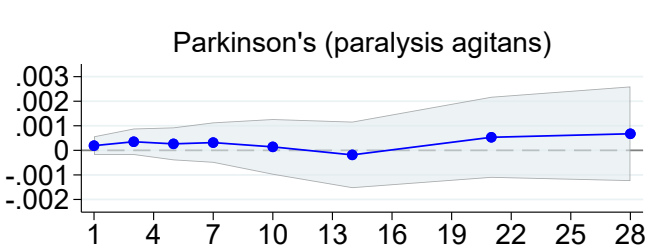
Parkinson's (paralysis agitans)
.003
.002
.001
0
-.001
-.002
1 4 7 10 13 16 19 22 25 28

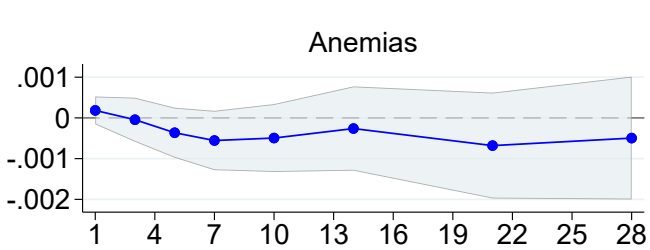
Anemias
.001
0
-.001
-.002
1 4 7 10 13 16 19 22 25 28

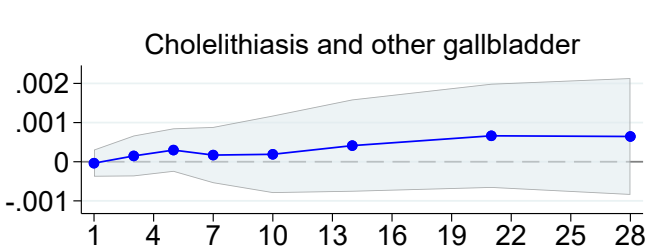
Cholelithiasis and other gallbladder
.002
.001
0
-.001
1 4 7 10 13 16 19 22 25 28

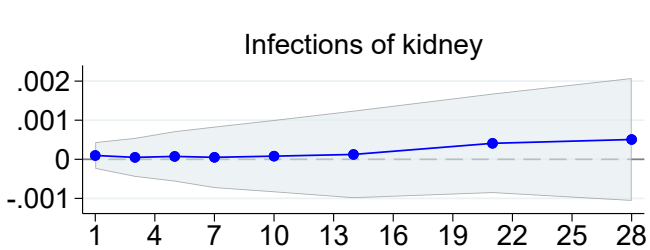
Infections of kidney
.002
.001
0
-.001
1 4 7 10 13 16 19 22 25 28

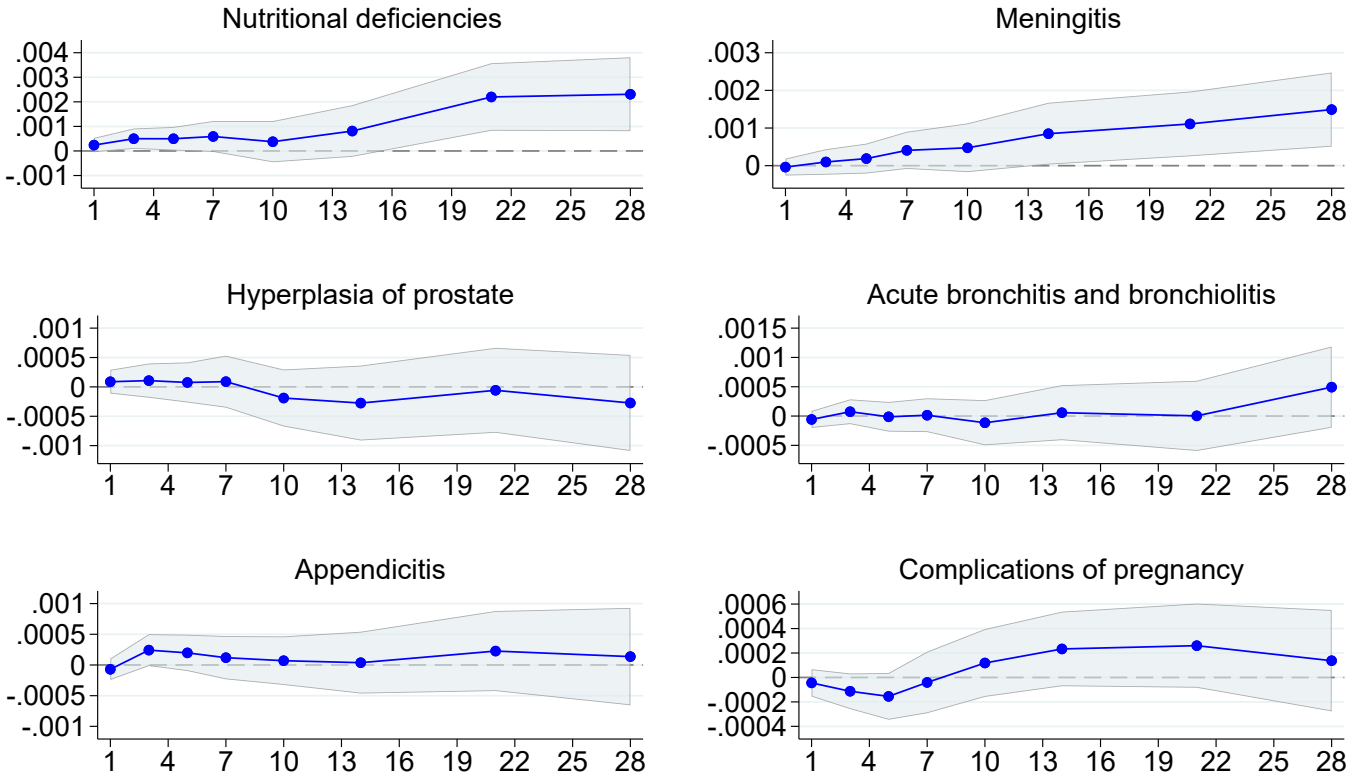

Notes: Each point reports an IV estimate from Equation 1 of the effect of acute (1-day) sulfur dioxide (SO_2) exposure on mortality (deaths per million), for thirty different causes of death. Mortality is measured with a time window ranging from 1 to 28 days as indicated by the x-axis. Shaded areas correspond to 95% confidence intervals. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors are clustered by county.

Table A.1: Correlations between county-level characteristics and the size of the first stage

	Population	Percent 65+	Percent Black	Per-capita income (dollars)	Per-capita transfers (dollars)	Employment rate	Mean SO_2 (ppb)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
First stage size	4,884	0.048	–1.4**	90**	32**	–0.55**	1.4**
	(8,728)	(0.053)	(0.15)	(28)	(4.4)	(0.17)	(0.087)
Mean outcome	240,068	11	9.2	8,857	1,117	47	7.3
Sample size	9,714	9,714	9,613	9,573	9,573	9,573	9,714
R-squared	0.0094	0.060	0.069	0.77	0.78	0.045	0.30

Notes: Dependent variable shown above each column. Indepedent variable is the difference in SO_2 changes between the most and least polluting wind direction for the monitor group to which the county belongs, measured in parts per billion. All regressions include year fixed effects. County characteristics are obtained from the Regional Economic Information System (REIS) dataset published by the Bureau of Economic Analysis (BEA). Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table A.2: OLS and IV estimates of effect of SO_2 on mortality, for different outcome windows

	(1)	(2)
Outcome window	OLS	IV
1 day	0.0077*	0.084**
	(0.0031)	(0.013)
3 days	0.021*	0.10**
	(0.0091)	(0.013)
5 days	0.029	0.099**
	(0.015)	(0.016)
7 days	0.031	0.11**
	(0.020)	(0.017)
10 days	0.035	0.12**
	(0.028)	(0.019)
14 days	0.042	0.13**
	(0.039)	(0.024)
21 days	0.049	0.16**
	(0.056)	(0.035)
28 days	0.046	0.19**
	(0.073)	(0.039)

Notes: Dependent variable is number of deaths per million people during the outcome window. Each estimate comes from a separate regression. IV estimates in Column (2) are shown in Figure 2. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table A.3: IV estimates of effect of SO_2 on mortality, for different age groups

	(1)	(2)
Age group	Absolute effect, deaths per million	Relative effect, percent
0–1	0.066	0.20
	(0.061)	(0.18)
1–19	0.0028	0.19
	(0.0022)	(0.14)
20–44	0.017**	0.38**
	(0.0054)	(0.12)
45–59	0.044**	0.22**
	(0.011)	(0.054)
60–64	0.099**	0.21**
	(0.034)	(0.073)
65–69	0.31**	0.44**
	(0.046)	(0.066)
70–74	0.23**	0.22**
	(0.070)	(0.066)
75–79	0.48**	0.31**
	(0.097)	(0.062)
80–84	1.1**	0.44**
	(0.17)	(0.070)
85+	2.3**	0.51**
	(0.44)	(0.10)

Notes: Dependent variable is number of deaths per million people in the given age group on the day of exposure. Each estimate comes from a separate regression. Relative effect is calculated as the percent of the age group's mean one-day mortality rate. Estimates are also reported in Figure 3. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table A.4: IV estimates of effect of SO_2 on mortality, for different age groups and outcome windows

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Outcome window	0–1	1–19	20–44	45–59	60–64	65–69	70–74	75–79	80–84	85+
1 day	0.066	0.0028	0.017**	0.044**	0.099**	0.31**	0.23**	0.48**	1.1**	2.3**
	(0.061)	(0.0022)	(0.0054)	(0.011)	(0.034)	(0.046)	(0.070)	(0.097)	(0.17)	(0.44)
3 days	0.21**	0.00049	0.020**	0.038*	0.13**	0.22**	0.36**	0.64**	1.4**	3.2**
	(0.080)	(0.0033)	(0.0063)	(0.016)	(0.043)	(0.065)	(0.089)	(0.12)	(0.25)	(0.56)
7 days	0.32*	–0.0047	0.023**	0.0010	0.12	0.31**	0.55**	1.0**	1.4**	3.3**
	(0.15)	(0.0061)	(0.0077)	(0.021)	(0.073)	(0.098)	(0.11)	(0.18)	(0.35)	(0.64)
14 days	0.27	–0.0076	0.0082	–0.0050	0.098	0.54**	0.78**	1.3**	2.3**	4.4**
	(0.27)	(0.0073)	(0.010)	(0.036)	(0.098)	(0.15)	(0.18)	(0.26)	(0.54)	(1.1)
28 days	0.53	0.012	0.0056	–0.0040	0.36*	0.85**	1.4**	1.9**	2.9**	6.5**
	(0.43)	(0.011)	(0.019)	(0.048)	(0.14)	(0.25)	(0.24)	(0.50)	(0.73)	(1.7)

Notes: Dependent variable is number of deaths per million people in the given age group over the given number of days following exposure. Each estimate comes from a separate regression. These estimates are used to calibrate the dynamic production model of health. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table A.5: IV estimates of effect of SO_2 on mortality, for different causes of death and outcome windows

	(1)	(2)	(3)	(4)
Outcome window	Cardio	Cancer	Other	External
1 day	0.029**	0.027**	0.023**	0.0050**
	(0.0045)	(0.0044)	(0.0045)	(0.0016)
3 days	0.041**	0.024**	0.030**	0.0070*
	(0.0054)	(0.0039)	(0.0046)	(0.0030)
5 days	0.046**	0.015**	0.035**	0.0041
	(0.0079)	(0.0043)	(0.0071)	(0.0033)
7 days	0.054**	0.012**	0.037**	0.0034
	(0.0089)	(0.0046)	(0.0078)	(0.0037)
10 days	0.060**	0.016**	0.045**	0.0031
	(0.011)	(0.0055)	(0.0097)	(0.0041)
14 days	0.070**	0.0085	0.052**	0.000064
	(0.015)	(0.0068)	(0.011)	(0.0054)
21 days	0.079**	0.0047	0.066**	0.0069
	(0.019)	(0.0086)	(0.016)	(0.0074)
28 days	0.097**	0.0082	0.077**	0.0074
	(0.020)	(0.011)	(0.019)	(0.0086)

Notes: Dependent variable is number of deaths per million people from the given cause of death over the given number of days following exposure. Each estimate comes from a separate regression. Estimates are also reported in Figure 4. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table A.6: IV estimates of effect of SO_2 on 1-day mortality, controlling for all pollutants except TSP

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
SO_2, ppb	0.085**	0.077**	0.086**	0.082**	0.077**	0.084**	0.078**
	(0.0085)	(0.012)	(0.010)	(0.0091)	(0.012)	(0.011)	(0.013)
NO_2, ppb		0.010			0.011		0.013
		(0.010)			(0.015)		(0.015)
Ozone, ppb			–0.0059			–0.0052	–0.0093
			(0.020)			(0.020)	(0.020)
CO, ppm				0.097	–0.012	0.094	–0.039
				(0.14)	(0.20)	(0.14)	(0.20)
First-stage F-statistic	225	72	34	91	41	34	30
Mean outcome	26	26	26	26	26	26	26
Sample size	275,690	275,690	275,690	275,690	275,690	275,690	275,690

Notes: Dependent variable is number of deaths per million people on the day of exposure. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table A.7: IV estimates of effect of SO_2 on mortality, using different fixed effects

	(1)	(2)	(3)	(4)	(5)	(6)
SO_2, ppb	0.084**	0.076**	0.079**	0.081**	0.080**	0.078**
	(0.013)	(0.0093)	(0.0089)	(0.0095)	(0.0094)	(0.0093)
Fixed effects	county-month, month-year	county, year, month	county, state-year-month	county-year, state-month	county, year, state-month	county, month-year, state-month
First-stage F-statistic	42	32	25	24	34	34
Mean outcome	25	25	25	25	25	25
Sample size	2,032,338	2,032,367	2,032,343	2,032,348	2,032,367	2,032,367

Notes: Dependent variable is number of deaths per million people on the day of exposure. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table A.8: IV estimates of effect of SO_2 on mortality, using different numbers of instrument leads and lags

	(1)	(2)	(3)	(4)	(5)	(6)
SO_2, ppb	0.084**	0.087**	0.090**	0.084**	0.083**	0.083**
	(0.013)	(0.013)	(0.014)	(0.013)	(0.013)	(0.013)
# of instrument leads	2	0	1	4	4	6
# of instrument lags	2	0	1	4	6	6
First-stage F-statistic	48	30	53	55	55	49
Mean outcome	25	25	25	25	25	25
Sample size	2,023,456	2,028,921	2,025,000	2,020,429	2,017,485	2,017,485

Notes: Dependent variable is number of deaths per million people on the day of exposure. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table A.9: IV estimates of effect of SO_2 on mortality, using different instrument specifications

	(1)	(2)	(3)	(4)	(5)	(6)
SO_2, ppb	0.084**	0.098**	0.074**	0.080**	0.081**	0.078**
	(0.013)	(0.014)	(0.012)	(0.016)	(0.014)	(0.015)
# of monitor groups	50	25	100	50	50	50
Wind angle spec	sines	sines	sines	10-degree bins	40-degree bins	60-degree bins
First-stage F-statistic	42	16	306	30	143	95
Mean outcome	25	25	25	25	25	25
Sample size	2,032,338	2,032,338	2,032,338	2,032,338	2,032,338	2,032,338

Notes: Dependent variable is number of deaths per million people on the day of exposure. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table A.10: 2SLS and LIML estimates of effect of SO_2 on mortality, for different outcome windows

	(1)	(2)	(3)	(4)	(5)	(6)
SO_2, ppb	0.084**	0.084**	0.11**	0.10**	0.19**	0.20**
	(0.013)	(0.013)	(0.017)	(0.013)	(0.039)	(0.028)
IV method	2SLS	LIML	2SLS	LIML	2SLS	LIML
Outcome window (days)	1	1	7	7	28	28
First-stage F-statistic	48	26	73	28	65	28
Mean outcome	25	25	173	25	691	25
Sample size	2,023,456	2,022,134	2,023,435	2,022,113	2,023,369	2,022,046

Notes: Dependent variable is number of deaths per million people over the time window specified in each column. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table A.11: Placebo and falsification tests of the effect of SO_2 on mortality

	(1)	(2)	(3)	(4)
SO_2, ppb	–0.079	0.18	–0.041	
	(0.062)	(0.23)	(0.49)	
SO_2 on day $t+1$, ppb				–0.0036
				(0.0048)
First-stage F-statistic	2.0	1.9	1.9	28
Outcome window, days	1	7	28	1
Mean outcome	25	173	691	25
Sample size	2,023,456	2,023,435	2,023,369	2,031,165

Notes: Dependent variable is number of deaths per million people over a window of 1, 7, or 28 days. Columns (1)–(3) report placebo tests, where wind direction is randomly generated. Column (4) reports a falsification test of the effect of SO2 on the previous day's mortality rate. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table A.12: IV estimates of effect of SO_2 on all-cause and cancer-related 1-day mortality, ages 65 and over

	(1)	(2)
Age group	All causes	Cancer-related causes
65–69	0.31**	0.17**
	(0.046)	(0.028)
70–74	0.23**	0.14**
	(0.070)	(0.034)
75–79	0.48**	0.13**
	(0.097)	(0.040)
80–84	1.1**	0.18**
	(0.17)	(0.065)
85+	2.3**	0.17*
	(0.44)	(0.084)

Notes: These estimates are used by the survival model to calibrate the effect of air pollution exposure on mortality. Dependent variable is number of deaths per million people in the given age group on the day of pollution exposure. Each estimate comes from a separate regression. All regressions include county-by-month and month-by-year fixed effects, as well as flexible controls for maximum temperature, precipitation, and wind speed; leads of these weather controls; and two leads and two lags of the instruments. Estimates are weighted by the county population. Standard errors, clustered by county, are reported in parentheses. A */** indicates significance at the 5%/1% level.

Table A.13: Baseline parameter values for the dynamic production model of health

	(1)	(2)
Parameter	Annual data	Daily data
I	0.74773	0.0020521
α	1.53762	1.537619
$\ln \delta$	–5.83878	–11.74124
μ_H	10.39737	11.43803
σ_e	2.25247	0.1178985
N	1,000,000	100,000
SSE	57.80479	20880.82

Notes: This table reports baseline parameter values for the dynamic production model of health (3). Column (1) reports parameter estimates when the model is fitted to annual survival data from a 1972 period life table. Column (2) reports corresponding estimates for daily data. The parameters $\underline{H}$ and σ_H are normalized to 0 and 1, respectively. N is the number of individuals in the simulation, and SSE is the sum of squared errors. The resulting fit for Column (2) is shown in Figure A.7. The mortality data underlying the life table are counts, not samples. Because there is no sampling error, we do not report standard errors for these parameter estimates.

Table A.14: Effect of permanent change in SO_2 on life expectancy (years)

	(1)	(2)	(3)
	IV extrapolation	Aging model 1 (α)	Aging model 2 (δ)
1-ppb decrease	0.16 [–0.03, 0.35]	1.18 [0.33, 2.21]	1.32 [0.36, 2.51]
2-ppb decrease	0.31 [–0.06, 0.70]	2.41 [0.55, 4.90]	2.67 [0.60, 5.57]
3-ppb decrease	0.47 [–0.10, 1.06]	3.74 [0.94, 7.56]	4.10 [1.02, 8.54]
1-ppb increase	–0.16 [–0.34, 0.03]	–1.05 [–1.99, –0.30]	–1.20 [–2.33, –0.33]
2-ppb increase	–0.31 [–0.68, 0.07]	–2.15 [–3.78, –0.54]	–2.46 [–4.43, –0.60]
3-ppb increase	–0.46 [–1.01, 0.10]	–3.09 [–5.34, –0.90]	–3.55 [–6.28, –1.00]

Notes: Each value in this table reports the estimated change in life expectancy (years) caused by a permanent, lifetime change in SO_2 of up to 3 part per billion (ppb) for individuals born in 1972. Baseline life expectancy in 1972 is 71.32 years. Confidence intervals for the 5th and 95th percentiles are reported in brackets. Estimates in Column (1), "IV extrapolation", are calculated by applying interpolated, age-specific 28-day IV estimates to the whole life-cycle. Estimates in Columns (2) and (3) assume that changes in daily mortality are governed by the dynamic production model of health (3). Those estimates assume that the effect of pollution on cancer-related deaths are governed by changes in $\underline{H}$, which produces mortality displacement, and that all remaining deaths are governed by changes in health capital. Columns (2)–(3) report estimates under different assumptions about which aging parameter drives the changes in health capital. Figure 6 shows how the 1-unit decrease estimates are distributed across the life cycle.