The Energy Transition, Air Pollution, and Mortality

Steve Cicala
Tufts University and NBER

Joshua D. Gottlieb*
University of Chicago and NBER

March 2026

Abstract

This paper studies the mortality benefits of the transition from coal-fired power generation in the US. Using hourly operational data of every generator from 1999-2019 combined with a high-frequency atmospheric dispersion model, we estimate the contribution of each plant's generation to downwind air quality. We use the reshuffling of generation across plants to estimate the mortality impacts of acute exposure to particulate matter among the elderly population. We find that 25% of U.S. fine particulates (PM2.5) in 1999 came from power plants, and that 40% of the air quality improvement from 1999–2019 was due to changes in electricity generation. We calculate one day of exposure to 1 $\mu g/m^3$ of PM2.5 raises elderly mortality by 0.12%, and that one month of exposure raises total mortality by 0.5%. Using an empirical dispatch model to estimate counterfactual outcomes, we find that the shift away from coal was worth $50B per year in reduced mortality by 2019.

Keywords: Energy Transition, Air Pollution, Mortality
JEL codes: Q4, Q53, I1, L94

*We are grateful to seminar participants at the University of Chicago, MIT, the University of Zürich, the University of Connecticut, the 2025 Energy Camp at UC Berkeley and Harvard University. Iván Higuera-Mendieta, Daisy Lu, Sifat Bin Raquib, Yann Decressin, and Scott Blatte provided outstanding research assistance. All errors remain our own. e-mail: scicala@gmail.com, jgottlieb@uchicago.edu

1 Introduction

Coal-fired power plants have long been among the largest sources of air pollution in the United States (Muller, Mendelsohn, and Nordhaus, 2011). The particulate matter, sulfur dioxide, and nitrogen oxides they emit travel hundreds of miles downwind, contributing to elevated mortality in counties far from the plants themselves (Holland et al., 2020; Henneman et al., 2023). Over the past two decades, the U.S. electricity sector has undergone a dramatic transformation: coal generation fell by half as natural gas and renewables expanded (Coglianese, Gerarden, and Stock, 2020). This paper asks how much this energy transition improved public health due to reduced fine particulate (PM2.5) pollution.

Answering this question requires solving two problems. First, we must trace emissions from each power plant to pollution exposure in each county—a task complicated by the fact that wind patterns vary from day to day, so the same plant affects different places at different times (Henneman et al., 2019). Second, we must estimate the causal effect of pollution on mortality, accounting for the many confounders that affect both electricity generation and health (Finkelstein, Notowidigdo, and Shi, 2026; Barreca et al., 2016).

We address these problems by combining three elements: hourly data on generation from every U.S. power plant, a high-frequency atmospheric dispersion model that tracks pollution transport under actual weather conditions, and a research design that isolates quasi-experimental variation in pollution exposure. Our instrument exploits the fact that electricity demand must be met in real time, so when one plant generates more, others generate less to keep the grid stable at a given load. Holding total generation fixed, variation in which plants operate—driven by fuel prices, equipment outages, and transmission constraints—reshuffles pollution across downwind counties many miles away. This reshuffling provides variation in pollution exposure that is plausibly unrelated to local determinants of mortality.

We find that particular matter from electricity generation substantially increases mortality, particularly among the elderly. A 1 μg/m^3 increase in PM2.5 for a single day raises

same-day mortality by 0.14% among the elderly population. Extending the analysis to longer horizons for the entire age distribution, we estimate that a 1 μg/m^3 increase in PM2.5 for 1 month raises total mortality by 0.5% over the subsequent five years. The fact that these estimates do not diminish over time suggests that pollution causes deaths that would not otherwise have occurred rather than merely hastening deaths by a few days. We find larger effects in counties without pollution monitors. These places tend to be more rural and have greater barriers to accessing medical care.

Combining these mortality estimates with predicted pollution from each plant, we calculate that electricity generation caused roughly 45,000 deaths per year at the start of our sample in 1999. By 2019, this had fallen to fewer than 20,000 deaths per year—a reduction of more than 50 percent. To isolate the contribution of the energy transition, we construct a counterfactual in which the grid remains frozen at its 2003 configuration. Under this counterfactual, deaths from electricity generation would have risen near 70,000 per year through 2019 due to an aging population. The difference—roughly 50,000 deaths prevented annually by the end of the sample—represents the mortality benefit of the shift from coal to cleaner generation.

This paper makes several contributions. First, we construct a fully-integrated pipeline to both measure pollution from power plants, and credibly estimate subsequent health impacts. Prior work on the health impacts of PM2.5 has either used quasi-experimental designs with coarse measures of pollution source (Deryugina et al., 2019), has been limited to studying impacts in the immediate vicinity of the source (Anderson, 2020), or has employed detailed exposure modeling without quasi-experimental variation (Henneman et al., 2023; Holland et al., 2020). Our approach integrates both elements, using atmospheric dispersion models to transport pollution, and the reshuffling of generation across plants to provide variation to identify the exact LATE of interest.

Second, we estimate mortality effects over horizons ranging from same-day to five years, documenting how effects accumulate and testing for mortality displacement. Short-horizon

studies cannot distinguish between pollution that accelerates imminent deaths and pollution that causes deaths that would not otherwise have occurred. Our distributed lag estimates show that cumulative effects grow and then become stable with the horizon, indicating that the deaths we attribute to pollution are not simply displaced by a few days or weeks.

Third, we quantify the health benefits of a major change in the U.S. energy system. The shift from coal to gas and renewables was driven primarily by market forces—declining natural gas prices following the shale revolution (Hausman and Kellogg, 2015)—rather than by policies targeting health outcomes. Our estimates suggest that this market-driven transition generated health benefits on the order of $50 billion per year by 2019, measured at the EPA's value of a statistical life year (U.S. Environmental Protection Agency, 2020). These benefits accrued over enormous swathes of the US east of the Mississippi River.

The infrastructure and data we develop for this paper yields multiple insights. The first stage of our estimation strategy produces the important finding that changing generation patterns is responsible for about 40% of the decline in ambient PM2.5 over the sample period. It also produces estimates of exposure to power plant pollution throughout the US, including in areas without monitors. Our estimates yield plant- (and equipment-) specific external marginal costs for the entire distribution of plants, both observed, and under various counterfactual scenarios.

The paper proceeds as follows. Section 2 provides background on air pollution from power plants and discusses related literature. Section 3 describes the data. Section 4 presents the estimation strategy. Section 5 reports results, including mortality estimates, aggregate death calculations, and counterfactual simulations. Section 6 concludes.

2 A Framework for Measuring the Mortality Impacts of Economic Activities

We are interested in measuring the external marginal costs of generation from U.S. power plants. In this section we break down the components of such a calculation, discuss prior work on each component, and motivate the need for an integrated estimation strategy.

Let $A_{jt'}$ represent a unit of polluting activity at plant j on day t'. In our setting, this will be the burning of fuel required for 1 MWh of electricity. Its units are MMBTU/MWh. Each MMBTU of fuel burned by plant j releases $e_{jt'}$ kg of particulate matter (primary and secondary). This pollution is then dispersed over time and space in a manner that depends on the weather. This dispersion can be characterized by a time-varying matrix, whose elements $D_{ijt't}$ represent the share of emissions released from plant j on date t' arriving in county i on date t. Dividing by Vol_i, the volume of atmosphere in county i, yields the ambient PM2.5 concentration produced by these emissions. This change in emissions represents the 'dose' of pollution, which, when multiplied by a dose-response estimate, τ, and the exposed population in the county, pop_i yields the estimated mortality impact. Multiplication by a value of statistical life, VSL, yields the monetized mortality cost of this marginal change in generation. Putting these elements together, the external marginal cost of generation from j's emissions on t' can be expressed as:

$$M_{jt'} = \tau \cdot VSL \cdot pop_i \cdot \sum_{i,t} \underbrace{\frac{A_{jt'} e_{jt'} D_{ijt't}}{Vol_i}}_{\Delta PM2.5_{it}} \tag{1}$$

Equation (1) is useful for framing the discussion of work on the health impacts of pollution. When the Environmental Protection Agency (EPA) conducts a regulatory impact assessment (RIA), it combines the various elements of equation (1) to estimate the health benefits of a proposed regulation (National Research Council, 2002; U.S. Environmental Protection Agency, 2024*b*,*a*). Though referred to as integrated assessment models, this process

is conducted piecewise, going back to at least Mendelsohn (1980). Engineering estimates of emissions rates feed into dispersion modeling to estimate changes in ambient pollution. These changes are then multiplied by estimates of the dose-response relationship, drawn from the broader epidemiology literature (U.S. Environmental Protection Agency, 2009). This change in mortality is then combined with economic valuations to produce a final estimate of the regulation's benefits.

In the economics literature, the dispersion step is often simplified by using a time invariant source-receptor matrix, which gives the average contribution of plant j's emissions to ambient PM2.5 in county i over the sample period (Muller and Mendelsohn, 2007). A time-invariant dispersion matrix, combined with prevailing dose-response estimates, is applied to settings such as trains (Clay et al., 2019), power (Holland et al., 2020; Jarvis, Deschenes, and Jha, 2022) and the substitution between electricity and internal combustion for transportation (Holland et al., 2016), among others. This branch of the literature focuses on specific activities of interest, and developing more sophisticated atmospheric chemistry to update static source-receptor matrices.

A second branch of work in economics has focused on estimating dose-response relationships, τ, using quasi-experimental variation in pollution exposure. This includes using changes in wind direction that do not reference pollutant sources (Deryugina et al., 2019; Deryugina and Reif, 2023), wildfires (Miller, Molitor, and Zou, 2024; Dodier, 2024; Cabral, Marika and Dillender, 2024), and changes in economic activity (Chay and Greenstone, 2003b; Currie and Neidell, 2005). The key advantage of this approach is in separating the causal effect of pollution from confounding factors that affect both pollution and health. Because the sources of pollution in these studies are broad-based, subsequent application of their estimates to specific activities relies on a uniform dose-response relationship across all sources of pollution. This is an effective approach when the emissions' source and estimate align (e.g. wildfires), or when the variation comes from the regulation being evaluated (e.g. the Clean Air Act in Chay and Greenstone (2003a) and Sanders, Barreca, and Neidell (2020)).

However, there is reason to believe that PM2.5 from ocean spray, agricultural dust, and automobiles have differing toxicities (Lippmann et al., 2013; Lippmann, 2014). This makes it important to estimate dose-response relationships for the specific sources of pollution being evaluated.

Economists have been able to credibly estimate source-specific dose-response relationships under two extreme cases of pollution dispersion. The first is immediate proximity, where the source of pollution is also the location of exposure. This includes studies of planes (Schlenker and Walker, 2016), automobiles (Anderson, 2020; Knittel, Miller, and Sanders, 2016; Currie and Walker, 2011), and fracking (Currie, Greenstone, and Meckel, 2017; Hill, 2018). The second is global pollutants like CO2, where exposure is uniform, so the dispersion step can be skipped entirely (Hogan and Schlenker, 2024; Carleton et al., 2022).

There is a vast epidemiological literature on the health effects of PM2.5; the U.S. Environmental Protection Agency's 2009 comprehensive assessment is more than 2,200 pages. This work has tended to focus on collecting information individual confounders (e.g. smoking habits), while ignoring the need for exogenous variation in exposure that has been the basis of the 'credibility revolution' in economics. The work most similar to ours is Henneman et al. (2023), who use a simplified dispersion model for SO2 emissions from coal-fired power plants to estimate annual exposure. The estimated relationship is based on the coefficient on this proxy variable, while variables that might drive both generation and mortality (e.g. economic conditions, high frequency weather) are not included.

Our contribution is to develop a fully-integrated pipeline that uses detailed operational data and atmospheric modeling to construct quasi-experimental variation in pollution exposure. Relative to recent work that has used time-invariant dispersion matrices to develop an index of exposure (Barreca, Neidell, and Sanders, 2021), our high-frequency analysis takes the inner product of actual operations with predicted dispersion at the time of those operations. This is important because coal-fired plants require substantial maintenance each year, so they are only capable of operating up to around 75% of the time. Older generators

are likely to be both more polluting, and used less often. These usage patterns depend on electricity demand, which varies with weather: both the wind patterns that affect dispersion, and the temperature that directly affects mortality. Incorporating this analysis into an instrumental variables framework allows us to directly estimate τ for PM2.5. We can then apply this estimate to the components of equation (1) to estimate the mortality impacts of generation from each plant, and to calculate the health benefits of the energy transition.

3 Data

We combine datasets that allow us to capture the full chain from power plant operations to pollution exposure to mortality outcomes from 1999–2019. Our primary sources include hourly load and generation data from the EIA and FERC (among others), atmospheric dispersion modeling using HYSPLIT, ambient PM2.5 measurements from EPA monitors, weather data from ERA5, and mortality records from Medicare and the CDC. We also incorporate economic controls from the BLS. This section describes each data source and our approach to constructing measures of pollution exposure from electricity generation.

3.1 Electricity Generation

Data from the electricity sector extend Cicala (2022) by providing hourly generation for every plant in the contiguous U.S. from 1999 to 2019. Hourly load data are mostly sourced from FERC Form 714, "Annual Electric Balancing Authority and Planning Area Report." These data are filed annually by local balancing authorities. Gaps in reporting are filled with complementary data from regional authorities and markets. Data on hourly generation similarly extend Cicala (2022), synthesizing EPA's CEMS, the Nuclear Regulatory Commission, and various hydropower assets through 2019. Fuel prices, monthly generation, and heat rates are sourced from EIA Form 923, filed annually by generators. Cicala (2022) offers a detailed description of the data construction process, including validation and imputation

methods.

The resulting dataset includes hourly generation (in MWh) for every plant, along with fuel consumption (in MMBTUs). We merge this with data on plant characteristics such as fuel type, generation technology, capacity, vintage, and pollution controls. These data come from both EPA's CEMS database, and EIA forms 860, 767, and 923.

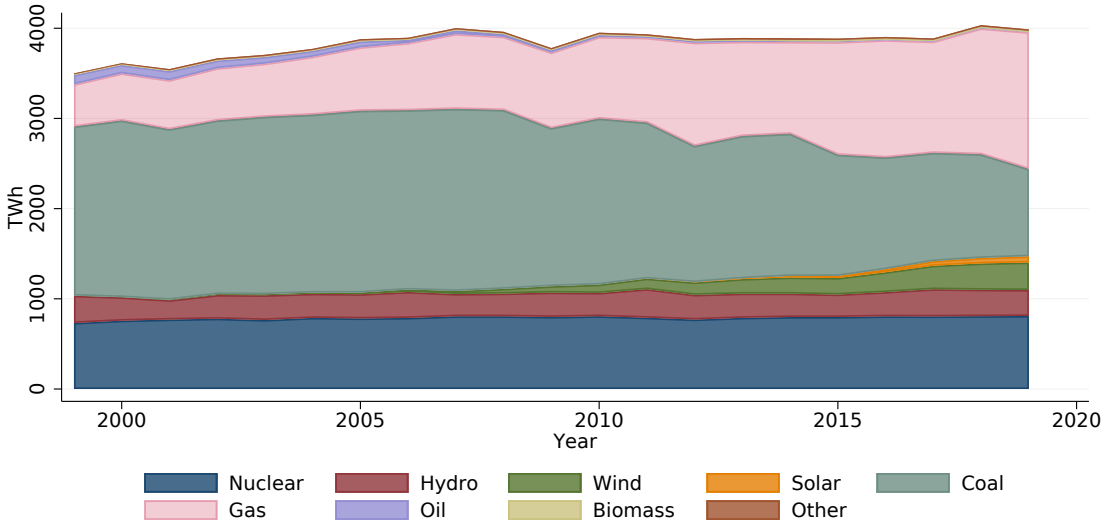
Aggregating these data by fuel type and year yields the trends in generation shown in Figure 1. Between 1999 and 2019, coal generation fell by more than half, from roughly 1,900 TWh to 900 TWh. Natural gas generation more than doubled over the same period, driven by declining gas prices following the shale revolution. Wind and solar generation, negligible in 1999, exceeded 300 TWh by the end of the sample. Notably, total electricity generation has been flat since 2007, with modest growth leading up to that point. This is important for two reasons: First, we will observe a given level of load being met in the data with vastly different generation mixes (and thus pollution levels) over time. This is the basis of the identification strategy, which relies on reshuffling generation across plants holding total electricity consumption fixed. Second, the flat load trend supports the credibility of a counterfactual in which the grid remains fixed at its 2003 configuration. Given the reserve capacity in the U.S. grid, it is plausible that the load patterns observed through 2019 could have been met with the 2003 grid.

Plant locations are geocoded, and we obtain stack heights from the EPA's National Emissions Inventory.

3.2 Atmospheric Dispersion

Translating generation at a power plant into pollution exposure in downwind counties requires modeling how pollutants disperse through the atmosphere. We use HYSPLIT (Hybrid Single-Particle Lagrangian Integrated Trajectory), an atmospheric transport model developed by NOAA's Air Resources Laboratory. HYSPLIT inputs a date, time, location, and elevation (in addition to various pollutant characteristics). It pulls the relevant meteorolog-

Figure 1: Annual Net Generation by Fuel Type, 1999–2019



ical data from the NCEP/NCAR Reanalysis Data product, then simulates the dispersion of individual particles through the atmosphere.

Data on power plant locations and stack heights are sourced from the EPA's National Emissions Inventory (NEI). When plants have multiple stacks, we assume that coal emissions are released from the tallest stack. We estimate stack height for plants without stacks reported in the NEI based on the plant's generation technology, vintage, and capacity.

For each power plant and fuel (associated with stacks of different heights), we simulate four emission releases per day of 500 particles each. We track the dispersion of particles for 48 hours. The model outputs the fraction of emitted particles passing through each county during each hour of the simulation. This is run for every power plant and fuel in the country, regardless of whether the plant is operating or not. Aggregating across hours *in the recipient county* yields a daily dispersion matrix $D_{jt'hit}$ giving the share of emissions from plant j on day t', hour h that arrives in county i on day t.

A key benefit of HYSPLIT is its relatively low computational cost per simulation. Parallelized on a research cluster, we are able to run the full set of simulations for over 4,000 plant-fuels over the 21-year sample period. This allows us to construct a complete dispersion

matrix with about 12TB of storage.

The main drawback of using the HYSPLIT model is the absence of chemical transformation. HYSPLIT tracks the physical dispersion of particles but does not model the chemical reactions that can occur in the atmosphere, such as the formation of secondary PM2.5 from precursor gases. This problem would be especially severe if we were measuring a subset of mass emitted for subsequent conversion into PM2.5 concentrations in the dispersion model itself.[1]

We address this issue in our first stage regression, which converts exposure to shares of production (in MMBTU) from plants of various characteristics, into ambient PM2.5 concentrations. Any secondary PM formation is therefore captured in the estimated coefficients. This approach allows us to leverage the computational efficiency of HYSPLIT while still capturing the relevant variation in pollution exposure for our mortality analysis.[2] Section 4 describes this approach in more detail.

3.3 Air Quality

We measure ambient PM2.5 concentrations using daily readings from the EPA's Air Quality System (AQS). The AQS network includes roughly 2,500 monitors reporting PM2.5 measurements in approximately 1,000 counties. We include counties with at least 1,000 monitor-days of data over the sample period, which includes over 75% of the U.S. population.

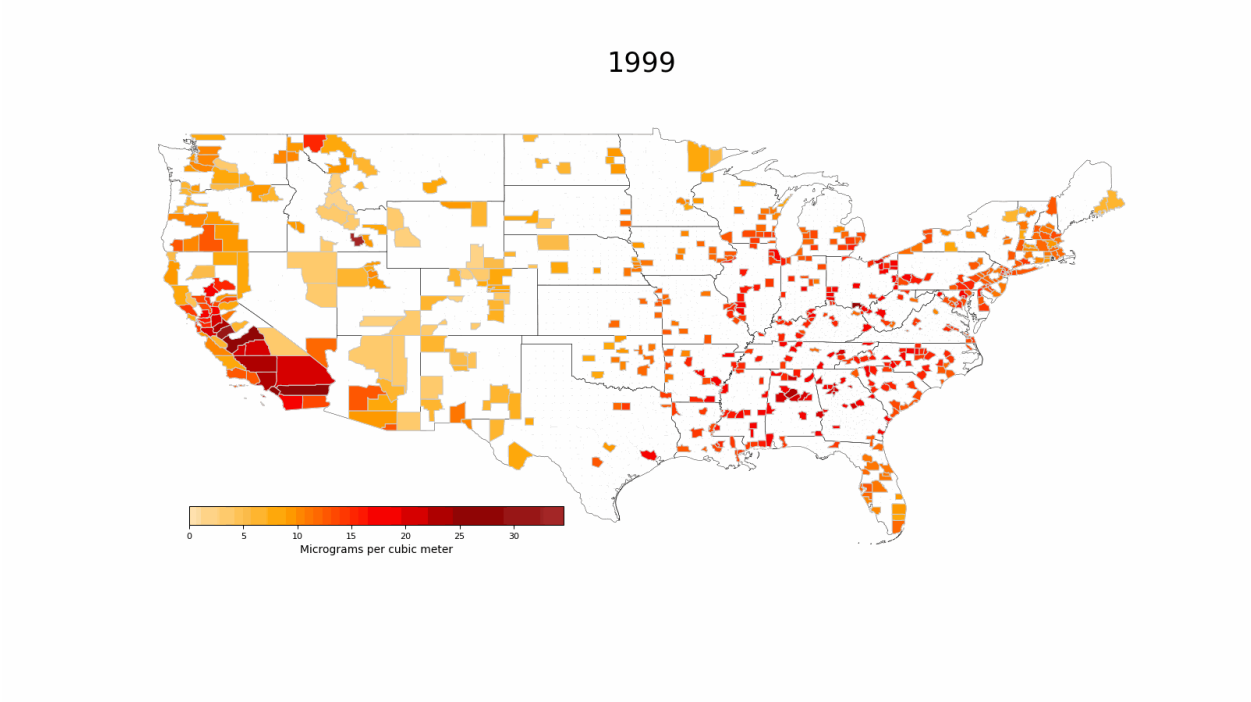
Figure 2 shows how the spatial distribution of PM2.5 has changed between 1999 and 2019. There were hot spots of high PM2.5 in California, the Ohio River Valley, and the Southeast in 1999, with annual averages above 25 μg/m^3 in some locations. This is more than double the current standard of 9.0 μg/m^3 set by the EPA. By 2019, these hot spots had largely disappeared, with most monitored counties reporting annual averages below that standard.

[1] For example, 1 kg of particles evenly dispersed over 1 km^3 would yield a 1 $\mu g/m^3$ increase in concentration. If 0.5 kg of secondary particles formed, this calculation would be off by 50%.

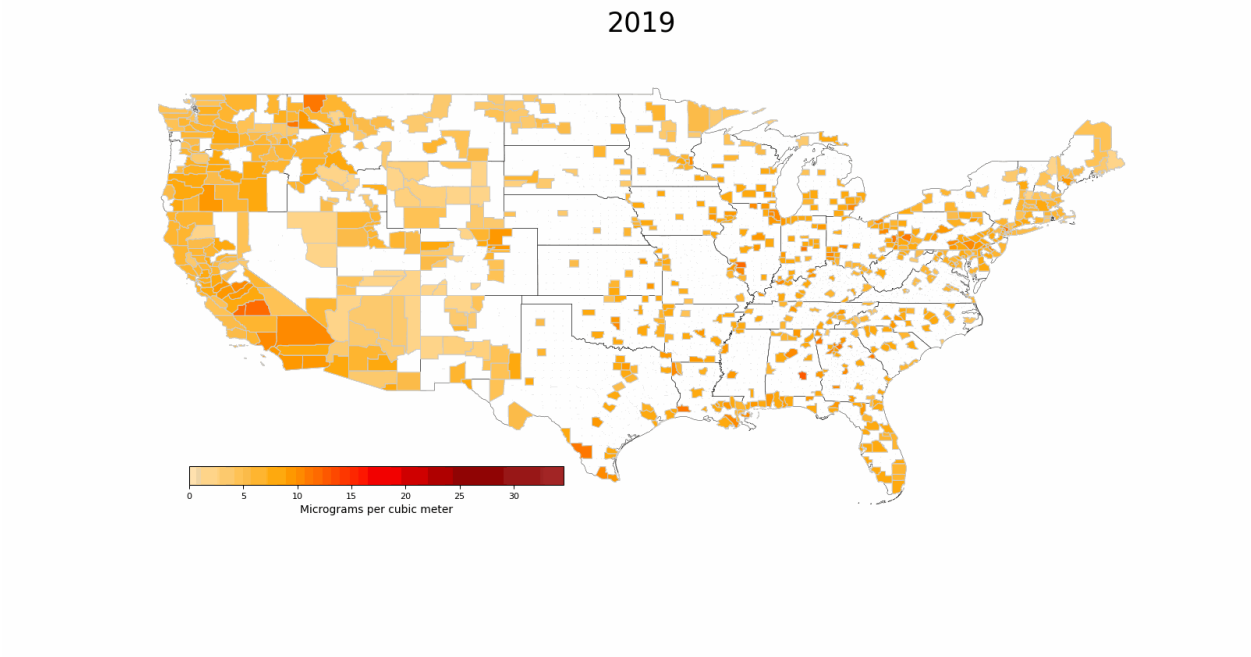
[2] Note that this strategy is necessitated in part by the absence of PM2.5 emissions data collection at U.S. power plants.

Figure 2: Annual Mean Ambient PM2.5 Concentrations, 1999 and 2019

(a)



(b)



Notes: Data from sites with at least 1,000 monitor-days of data over the sample period are included.

Figure 2 also highlights that the monitor network is sparse: most U.S. counties have no PM2.5 monitor, and even monitored counties may have readings on only a subset of days. This sparsity motivates our reduced-form approach, which uses predicted pollution from generation rather than relying solely on monitor readings. The first-stage regression, estimated on the subset of county-days with monitor data, translates exposure to generation by plant characteristics into ambient concentrations. We then apply these coefficients to construct predicted pollution for all counties, whether or not they have monitors. We detail this approach in Section 4.

3.4 Weather

Weather data come from ERA5, the fifth-generation reanalysis produced by the European Centre for Medium-Range Weather Forecasts. We extract daily county-level measures of temperature, humidity, and precipitation. Our primary temperature variable is wet bulb temperature, which combines heat and humidity into a single measure of thermal stress. We calculate heating and cooling degree-days relative to a 65°F baseline.

Weather is a critical control because it affects both pollution concentrations (through atmospheric chemistry and dispersion) and mortality (through heat stress, cold exposure, and respiratory illness). Failing to control adequately for weather would bias estimates of pollution's effect on mortality.

3.5 Mortality

Mortality data come from two sources. For the elderly population, we use the Medicare Beneficiary Summary File, which records the date of death for all Medicare enrollees. These data cover 1999–2019 and allow us to construct daily death counts by county and five-year age group (65–69, 70–74, 75–79, 80–84, 85+).[3] The Medicare data provide high-frequency

[3]For administrative reasons, our access to these data are separated across two contracts at differing institutions. One covers 1999-2013 and the other covers 2014-2019. We run the analyses separately and report inverse-variance weighted estimates. cite XXX

mortality information for the population most vulnerable to air pollution.

For the full population, we use monthly mortality counts from the CDC's National Vital Statistics System, covering 1999–2019. These data report deaths by county, month, and age group, allowing us to examine mortality effects across all ages, albeit at a monthly resolution.

3.6 Economic Controls

Local economic conditions may confound the relationship between generation and mortality. Employment and wages affect both electricity demand (through industrial activity) and health (through income, stress, and access to care). We include monthly county-level employment and average weekly wages from the Bureau of Labor Statistics' Quarterly Census of Employment and Wages. Monthly county-level data on unemployment rates come from the BLS' Local Area Unemployment Statistics.

4 Estimation Strategy

There are two central challenges in estimating the health effects of power plant emissions: the endogeneity of pollution exposure and measurement error in pollution levels. Electricity demand rises with economic activity, which independently affects mortality (Finkelstein, Notowidigdo, and Shi, 2026). Weather conditions simultaneously drive both electricity consumption and health outcomes (Barreca et al., 2016). Aside from measuring the dispersion of pollution to particular locations, the network of monitors is sparse, and many counties have no monitors at all. The measurements from individual monitors are also noisy proxies for county-level exposure, especially in large counties with heterogeneous pollution levels (Sullivan, 2017).

We address these concerns by constructing a predicted measure of pollution exposure that combines high-frequency variation in electricity generation (holding total generation fixed) with atmospheric dispersion modeling. This approach yields quasi-experimental variation in

pollution that is plausibly orthogonal to local determinants of mortality.

4.1 Empirical Framework

We estimate the effect of particulate matter exposure on mortality using a distributed lag specification:

$$y_{it} = \sum_{d=0}^{D} \{P_{i,t-d}\tau_d + X'_{i,t-d}\beta_d + f(L_{i,t-d})\} + \phi_i + \lambda_{ym} + \eta_{sm} + u_{it} \tag{2}$$

where y_{it} is the mortality rate in county i on day t, and $P_{i,t-d}$ is ambient PM2.5 concentration d days prior. The distributed lag structure allows pollution exposure on day $t-d$ to affect mortality on day t, capturing both immediate physiological responses and delayed effects that unfold over subsequent days or weeks.

The vector $X_{i,t-d}$ includes weather controls—wet bulb heating and cooling degrees, and precipitation. The lags parallel those of pollution exposure to account for cumulative heat or cold exposure. We note the importance of always preserving identical lags for pollution and weather controls, as the two are correlated and omitting lags of weather biases the pollution coefficients.

We control flexibly for local electricity load and its corresponding lags, $f(L_{i,t-d})$, using a spline, since load correlates with both generation-related pollution and other mortality determinants like heating and cooling demand. County fixed effects ϕ_i absorb time-invariant differences in population health, industrial composition, and baseline air quality. Year-by-month fixed effects λ_{ym} control for national trends in mortality, seasonality, and macroeconomic conditions. State-by-month fixed effects η_{sm} allow seasonal patterns to differ across states, accommodating regional variation in climate, flu timing, and holiday patterns.

The parameters of interest are the τ_d, which trace out the dynamic mortality response to pollution exposure. The sum $\sum_d \tau_d$ gives the cumulative effect of a sustained increase in PM2.5 on mortality, accounting for any displacement or delayed impacts.

4.2 Constructing the Instrument

We employ an instrumental variables strategy to address the endogeneity of pollution exposure and errors in pollution measurement. The instrument combines three elements: the quantity of fuel burned at each plant, the characteristics of each generating unit, and the atmospheric transport of emissions from plants to downwind counties.

For each plant j and time period t', we observe fuel combustion $A_{jt'}$ in MMBTU. Following the framework in section 2, combining this activity with an emissions rate and dispersion pattern would yield the contribution of plant j to ambient PM2.5 in county i on day t. While SO_2 and NO_x emissions rates are observed for many plants, PM2.5 emissions rates are not. We therefore project emissions rates onto a vector of observed characteristics $K_{jt'}$, leaving a scalar unobserved component $\delta_{jt'}$ that is orthogonal to those characteristics by construction:

$$e_{jt'} = K'_{jt'}\gamma + \delta_{jt'} \tag{3}$$

This vector of unit characteristics includes fuel type (coal, gas, oil), generator technology (boiler, turbine, combined cycle), regulatory epoch of the unit's vintage, categories of generation capacity, and the presence of pollution control equipment (scrubbers, baghouses, selective catalytic reduction).

Substituting this decomposition into the expression for plant j's contribution to pollution in county i on day t due to its generation on day t':

$$p_{itjt'} = A_{jt'} \cdot K'_{jt'}\gamma \cdot D_{jt'it} + \xi_{ijt} \tag{4}$$

Then summing over all plants and periods t' that could affect pollution on day t in county i gives:

$$p_{it} = \sum_{j} \sum_{t'} A_{jt'} \cdot K'_{jt'}\gamma \cdot D_{jt'it} + \xi_{it} \tag{5}$$

We use this inner product of generation, characteristics, and dispersion as the basis for

our instrument:

$$Z_{it} = \sum_{j \notin i} A_{jt'} \cdot K_{jt'} \cdot D_{jt'it} \tag{6}$$

In words, the instrument vector is the amount of heat consumed by upwind plants with specific characteristics that is transported to county i on day t. It is a vector because the characteristics $K_{jt'}$ are multidimensional, and the instrument includes one element for each characteristic: 50 MMBTU from coal-fired boilers without scrubbers built in the 1970s, 100 MMBTU from gas turbines with selective catalytic reduction built in the 2000s, and so on.

4.3 First Stage

Combining the structural equation with the instrument gives the following first-stage regression:

$$P_{it} = Z'_{it}\gamma + X'_{it}\zeta + f(L_{it}) + \phi_i + \eta_{sm} + \nu_{it} \tag{7}$$

The coefficient vector γ translates heat-weighted characteristics into ambient pollution concentrations: Each element gives the PM2.5 concentration (in μg/m^3) associated with full-day exposure to the emissions from 1 MMBTU of heat generated by a unit with that characteristic, holding other characteristics fixed. As an additional precaution against local confounding, we exclude from the instrument any plants located in the same county as the monitor, so that Z_{it} depends only on upwind sources of pollution in other counties.[4]

This first-stage regression serves three purposes. First, it validates the instrument: we test whether the pattern of coefficients is sensible (e.g., coal units contribute more to PM2.5 than gas units). Second, it allows us to construct predicted pollution from all plants, including those in the own county, which we excluded from the instrument. Third, by estimating γ from the data rather than imposing engineering values, we allow the relationship between

[4]This restriction does not materially affect the results.

plant characteristics and ambient pollution to be determined by the quasi-experimental variation in the instrument.

The identifying assumption is that, conditional on controls, upwind plant generation only affects downwind mortality through its impact on ambient pollution.

4.4 Predicted Pollution and Split-Sample IV

With estimates of γ in hand, we construct predicted pollution from electricity generation for each county-day:

$$\widehat{PM}_{it} = \sum_{j} \sum_{t'} A_{jt'} \cdot K'_{jt'} \hat{\gamma} \cdot D_{jt'it} \tag{8}$$

This sum includes all plants, including those in the same county as the outcome, which we excluded from the instrument. The predicted pollution measure $\widehat{PM}_{it}$ isolates the component of ambient PM2.5 attributable to electricity generation.

We estimate the mortality effect of predicted pollution directly:

$$y_{it} = \sum_{d=0}^{D} \{\widehat{PM}_{i,t-d}\tau_d + X'_{i,t-d}\beta_d + f(L_{i,t-d})\} + \phi_i + \lambda_{ym} + \eta_{sm} + u_{it} \tag{9}$$

When equation (9)) is estimated using the same sample of county-days used to estimate $\hat{\gamma}$, the τ_d coefficients are identical to the two-stage least squares estimates from the standard IV regression that uses Z_{it} as the instrument for P_{it}. However, we can also estimate equation (9)) on the full sample of county-days, including those without monitors. In this case, the τ_d coefficients are equal to the IV estimates, but they still have a causal interpretation under the same identifying assumptions. We are referring to these estimates as "split sample IV" because we are estimating the first stage in one part of the sample (county days with monitors), and estimating the second stage using predicted values for the whole sample. This approach allows us to leverage all available data on mortality while still using quasi-experimental variation in pollution exposure.

4.5 Interpretation

Local average treatment effect. Our estimates identify the mortality effect of PM2.5 from electricity generation, not PM2.5 from all sources. If the dose-response relationship differs by source—perhaps because generation-related PM2.5 has different chemical composition than wildfire smoke or vehicle exhaust—our estimates may not generalize. That said, the policy-relevant question is precisely what we estimate: how much would mortality fall if generation-related pollution declined by reallocating generation across plants?

Multiple pollutants. Power plants emit not only primary particulate matter but also nitrogen oxides, which form ozone when combined with volatile organic compounds and heat. Our estimates capture the combined effect of all pollutants whose emissions covary with generation and unit characteristics. In the appendix, we show that our estimates are robust to controlling for predicted ozone from generation, suggesting that the effects we attribute to PM2.5 are not driven by correlated changes in ozone.

Inference. While inference when estimating equation (2) is straightforward, in the split sample specification, the regressor $\widehat{PM}_{it}$ is generated from first-stage estimates. Standard errors should account for estimation uncertainty in $\hat{\gamma}$. Our baseline split sample specifications do not make this adjustment; we are developing bootstrap procedures that propagate first-stage uncertainty to the mortality estimates.

5 Results

This section presents estimates of the mortality effects of PM2.5 from electricity generation and applies them to quantify the health benefits of the energy transition. We begin by validating the first stage: confirming that generation from plants with different characteristics contributes differentially to ambient pollution. We then estimate mortality effects at daily and monthly frequencies, documenting how effects accumulate over time. Finally, we com-

bine these estimates with predicted pollution to calculate aggregate deaths attributable to generation and to simulate how mortality would have evolved had the grid remained frozen at its 2003 configuration.

5.1 First Stage: From Generation to Ambient Pollution

The first-stage regression translates the instrument—heat input from upwind plants, differentiated by unit characteristics and weighted by atmospheric dispersion—into ambient PM2.5 concentrations. The coefficient vector $\hat{\gamma}$ converts marginal generation from plants with specific characteristics into downwind pollution. Validation of the first stage requires not only statistical precision, but also that these coefficients align with engineering expectations: coal plants should contribute more than gas plants, and plants without pollution controls should contribute more than those with scrubbers, etc.

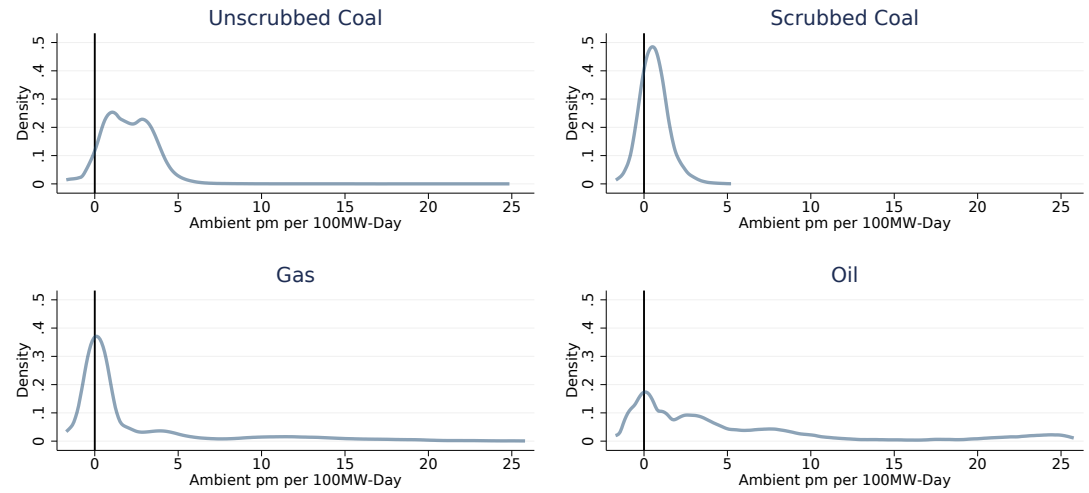
Figure 3 displays kernel densities of predicted plant-level contributions to ambient PM2.5, calculated as $A_{jt} \cdot K'_{jt}\hat{\gamma}$ for each plant-month.[5] To better facilitate interpretation (where the mean ambient PM2.5 reading is about 10 $\mu g/m^3$, the marginal impact is multiplied by 24 hours and 100 MWh. The x-axis therefore represents the change in ambient PM2.5 experienced by exposure to the complete pollution emitted by a 100MW plant with characteristics K_{jt} generating for 1 day. Ambient air quality is presented for an average-sized county. The distributions confirm that generation technology matters substantially for pollution. Coal-fired plants dominate the right tail: the median coal plant-day contributes several times more PM2.5 per MMBTU than the median gas plant-day. Within coal, plants without scrubbers show the largest contributions.

It is also worth recognizing the magnitude of these contributions. In 1999, there were over 125 coal-fired power plants with at least 1GW of capacity. 80% of these plants had no scrubbers, and overall there was nearly 250GW of unscrubbed coal capacity. Exposure to the emissions of a single, typical unscrubbed coal-fired power plant of 1GW capacity would raise

[5] Plant characteristics, including pollution controls, heat rates, and even fuel type change over time.

ambient PM2.5 to more than triple the national average. Scrubbing the coal clearly improves this situation, but does not eliminate the problem: typical emissions from a 1GW coal plant with a scrubber still raises ambient PM2.5 by about 10 $\mu g/m^3$. In contrast, emissions from gas plants tend to be tightly clustered around zero, with small positive values for the few gas plants using internal combustion. The estimates for oil-fired plants is also notable: oil has been largely squeezed out of the generation mix, but was an important source of peaking capacity, especially earlier in the sample. These units tended to be located in urban areas and without tall stacks so their contributions to local PM2.5 was likely concentrated in the manner calculated for this figure. The various bumps reflect types of prime movers, with internal combustion diesel units showing the largest contributions. Overall, the first stage results align well with expectations and provide confidence that the instrument captures meaningful variation in pollution exposure.

Figure 3: Distribution of First-Stage Contributions by Fuel Type



Notes:

Kernel densities of predicted plant-level contributions to ambient PM2.5, calculated as $A_{jt} \cdot K'_{jt}\hat{\gamma}$. Each observation is a plant-month.

In Figure 4, we operationalize the first stage with the complete generation and dispersion estimates to predict the spatial dispersion of generation's contribution to ambient air quality in 1999 and 2019. Let $Q_{t'}$ represent the j-length vector of all power plant generation at time

t' in MWh. Continuing with the matrix representation of equation (8), we can calculate the i-length vector of predicted PM2.5 concentrations according to:

$$\widehat{PM}_t = Q_{t'} A_{t'} K_{\prime t} \hat{\gamma} D_{t't} \tag{10}$$

Figure 4 then reproduces Figure 2, by taking the annual means of equation (10)). Comparing with the previous figure, the first notable feature is that (10)) can be calculated for *all* counties, regardless of whether it contains an ambient PM2.5 monitor.[6]

Figure 4 highlights the astounding change in predicted pollution from power plants between 1999 and 2019. The entire US has seen decreases, lead by the coal-intensive regions of the Ohio River valley, Texas, and Northern Alabama and Georgia. The places experiencing the highest pollution levels in 1999 were not merely the sites of power plants, but large areas that experienced their emissions downwind. It is also worth noting the parallels between the spatial patterns in Figures 2 and 4. Because these are annual averages, one might be concerned that some longer-term economic decline might be an omitted variable driving the patterns in both figures. However, the coefficients estimated in the first stage are based on high-frequency exposure to pollution from hourly generation patterns. While economic decline may yield less electricity generation ($Q_{t'}$), it would not be expected to change the physical relationship between generation and ambient pollution ($\hat{\gamma}$) or the atmospheric dispersion of pollution ($D_{t't}$). The close alignment of the spatial patterns in the two figures provides further validation that our predicted measure is capturing relevant variation in ambient pollution.

We get a better sense of the overall importance of generation-related pollution by plotting annual trends in observed ambient PM2.5 and predicted PM2.5 from generation in

[6]This does require the assumption that the physical relationship between predicted receipt of pollution from upwind plants and ambient air quality is the same between monitored and unmonitored counties. Given that monitors are located based on population, not geography, this seems reasonable enough. This will be violated if, for example, populations locate to places where the geography is such that a given pollution deposition can be detected in a manner that is substantially different from how it would be detected in sparsely populated areas.

Figure 5. Population-weighted annual average PM2.5 concentrations declined from nearly 15 $\mu g/m^3$ in 1999 to about 8 $\mu g/m^3$ in 2019, about a 50% decline. Figure 5 shows the decline in generation-related pollution has accounted for nearly 3 $\mu g/m^3$ in aggregate air quality improvements. This represents 40% of the total improvement in ambient PM2.5 over the sample period. The figure also shows that this improvement has come essentially entirely from reductions in coal generation. While oil generation may also be dirty, it is also a small part of the generation mix that does not travel far.

5.2 Mortality Effects: Daily Estimates

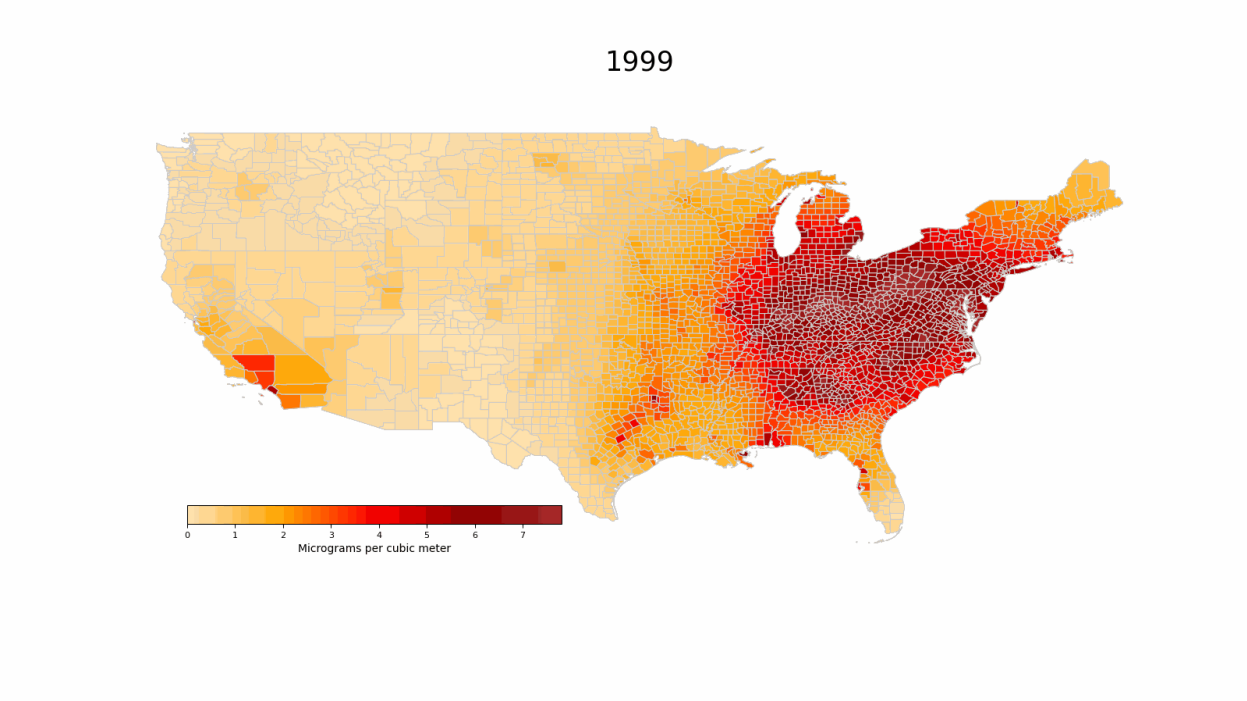
We estimate the mortality effects of PM2.5 using Medicare data for the population aged 65 and older, covering 1999–2019. We begin by estimating same-day impacts, so the lags of equation (2)) are removed. It is important to highlight this specification does *not* represent the causal effect of a single day of exposure to a given amount of PM2.5. This is because there is a substantial amount of serial correlation in both pollution and temperature. Days of prior exposure become omitted variables in equation (2)) and will load to some degree on both the contemporaneous pollution and weather controls. It nonetheless serves as a useful starting point for the analysis.

Table 1 presents estimates for the 85+ age group across specifications with progressively richer controls. Panel A presents the results of an OLS regression for different versions of equation (2)), ignoring the electricity instrument. Column (1) includes only county and year-month fixed effects. Column (2) adds state-month and day-of-week fixed effects. Column (3) adds weather controls. Column (4) adds economic controls, including flexible controls for load.

The OLS and IV estimates tell a consistent story, but differ in magnitude. In the most saturated specification (column 4), OLS yields a coefficient of 0.33 deaths per million, while IV yields 0.51. There are a couple prospective explanations for this difference. First, the LATE measured in the IV reflects pollution from power plants. There is good reason to

Figure 4: Annual Average Predicted PM2.5 from Generation, 1999 and 2019

(a)



(b)

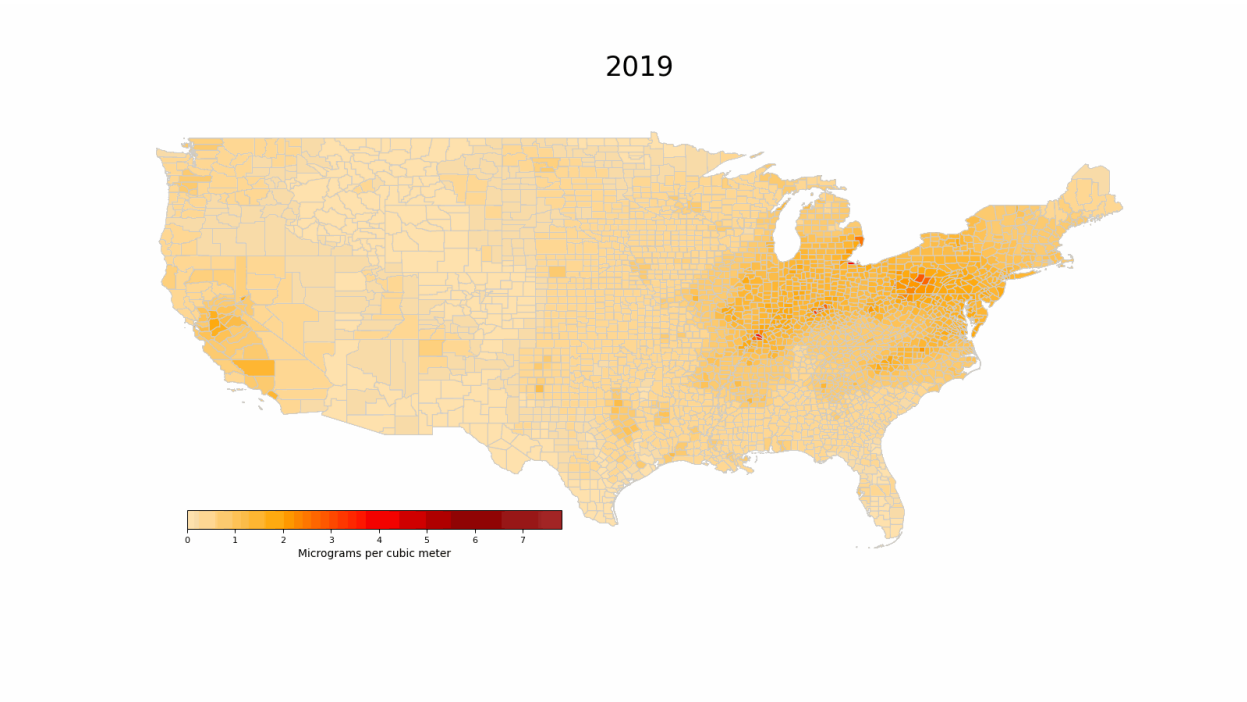


Table 1: Same-Day Mortality Effects of PM2.5: Ages 85+

Panel A: OLS

	(1)	(2)	(3)	(4)
PM 2.5	0.287***	0.356***	0.346***	0.333***
	(0.046)	(0.046)	(0.050)	(0.049)
Dependent Variable Mean	377	377	377	377
Estimate Relative to Mean, Percent	0.076	0.094	0.092	0.088
State-Month, DOW FE		✓	✓	✓
Weather Controls			✓	✓
Econ Controls				✓
Observations	2,173,632	2,173,632	2,173,632	2,173,632

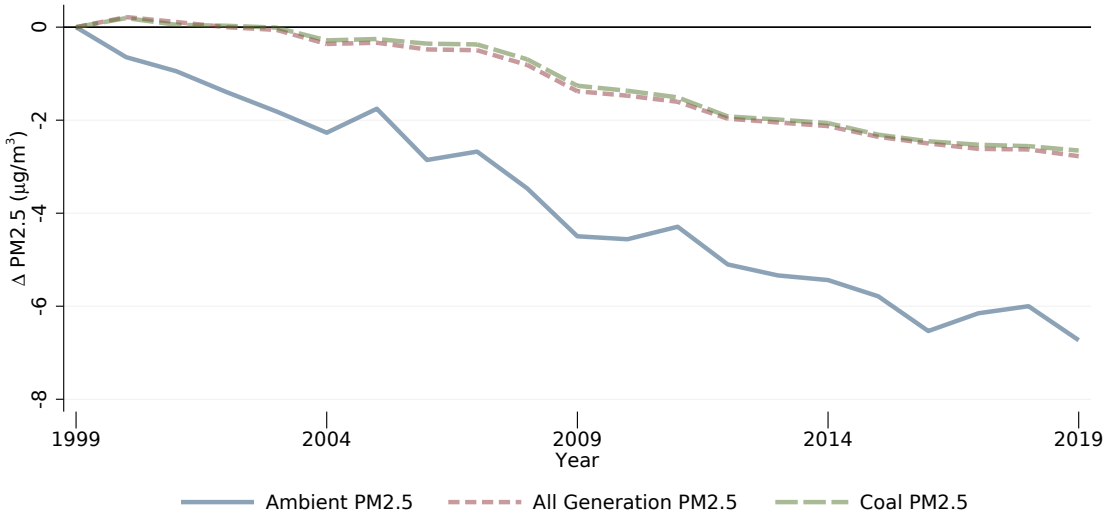
Note: All specifications include county and year-month fixed effects. Standard errors clustered by county in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Panel B: IV

	(1)	(2)	(3)	(4)
PM 2.5	0.493***	0.680***	0.548***	0.511***
	(0.099)	(0.088)	(0.098)	(0.101)
Dependent Variable Weighted Mean	377	377	377	377
Effect relative to mean, percent	0.131	0.180	0.145	0.135
F-statistic, PM 2.5	3,033.2	2,863.2	2,460.7	2,274.5
State-Month, DOW FE		✓	✓	✓
Weather Controls			✓	✓
Econ Controls				✓
Observations	2,173,632	2,173,632	2,173,632	2,173,632

Note: All specifications include county and year-month fixed effects. Standard errors clustered by county in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Figure 5: Changes in Observed and Predicted PM2.5 from Electricity Generation, 1999–2019



Notes:

Population-weighted annual average PM2.5 concentrations. Observed ambient PM2.5 from EPA monitors and predicted PM2.5 from electricity generation, constructed from first-stage estimates. Magnitudes are plotted relative to 1999 values.

believe such pollution is more toxic than generic PM2.5, which may be less harmful. Second, the EPA monitor network is sparse, and daily readings from a single monitor are noisy proxies for population-average exposure within a county. Instrumenting for ambient air quality with exposure to pollution from power plants helps address this measurement error. The stability of both OLS and IV estimates as controls are added suggests that omitted variable bias is not a dominant force driving the estimates. The first-stage F-statistic exceeds 2,000 in all IV specifications, indicating that weak instruments are not a concern. The preferred IV specification indicates that a 1 μg/m^3 increase in same-day PM2.5 raises mortality among those aged 85 and older by 0.14 percent.

Table 2 reports IV estimates across age groups. The mortality effect rises with age in absolute terms: the coefficient increases from 0.05 deaths per million for the 65–69 group to 0.51 for the 85+ group—an order of magnitude larger. This gradient reflects that the elderly are more vulnerable to the cardiovascular and respiratory insults caused by fine particulate matter. When scaled by baseline mortality rates, however, the proportional effects are more similar across groups, ranging from 0.12 to 0.15 percent per μg/m^3. This suggests that a

pollution shock poses similar proportional risk across the elderly population, but translates into more deaths in absolute terms for groups with higher baseline mortality.

Table 2: Same-Day Mortality Effects of PM2.5 by Age Group: IV Estimates

	65-69	70-74	75-79	80-84	85+
	(1)	(2)	(3)	(4)	(5)
PM 2.5	0.052**	0.088***	0.161***	0.197***	0.511***
	(0.023)	(0.028)	(0.037)	(0.045)	(0.101)
Dependent Variable Mean	44	66	105	170	377
Effect relative to mean, percent	0.119	0.133	0.153	0.116	0.135
F-statistic, PM 2.5	2,086.2	2,152.7	2,211.9	2,244.9	2,274.5
Observations	2,173,632	2,173,632	2,173,632	2,173,632	2,173,632

Note: All specifications include county, year-month, state-month and day of week fixed effects, as well as controls for weather, load, and employment. Standard errors clustered by county in parentheses. * p<0.1, ** p<0.05, *** p<0.01

5.3 Split Sample Estimates Using Predicted Pollution

The split sample specification regresses mortality directly on predicted PM2.5 from generation, as defined in equation (8). This approach allows us to use all counties, not just those with PM2.5 monitors, substantially expanding the sample.

Table 3 reports split sample estimates of same-day mortality effects by age group. The coefficients on predicted PM2.5 are substantially larger than the IV coefficients on measured PM2.5. For the 85+ group, the split sample coefficient is 1.06 deaths per million per $\mu g/m^3$ of predicted pollution. Recall that this method would replicate the point estimates of Table 2 if the sample were restricted to that generating the first stage. What accounts for the difference? One prospective reason is that counties without a monitor are more sparsely populated. The elderly in these counties will typically have to travel much further to seek care, and the care they receive may be less effective than in the major population centers where monitors are located. These estimates suggest that the mortality effects of pollution are substantially larger in rural areas, but the absence of monitoring has prevented this from being estimated in prior work.

Table 3: Same-Day Mortality Effects of Predicted PM2.5 by Age Group

	65-69	70-74	75-79	80-84	85+
	(1)	(2)	(3)	(4)	(5)
Predicted PM2.5	0.088***	0.157***	0.271***	0.362***	1.06***
	(0.016)	(0.021)	(0.031)	(0.041)	(0.092)
Dependent Variable Weighted Mean	46	69	109	177	388
Effect relative to mean, percent	0.191	0.227	0.249	0.204	0.274
Observations	17,015,969	17,015,969	17,015,969	17,015,969	17,015,969

Note: All specifications include county, year-month, state-month and day of week fixed effects, as well as controls for weather, load, and employment. Standard errors clustered by county in parentheses. * p<0.1, ** p<0.05, *** p<0.01

The IV estimates have a direct policy interpretation. Each coefficient captures the mortality consequences of exposure to predicted pollution from electricity generation, integrating the emissions intensity of different plant types, atmospheric transport, and the dose-response relationship. These are the objects needed to calculate the mortality cost of generation from specific plants or to simulate counterfactual policies.

5.4 Dynamic Effects and Mortality Displacement

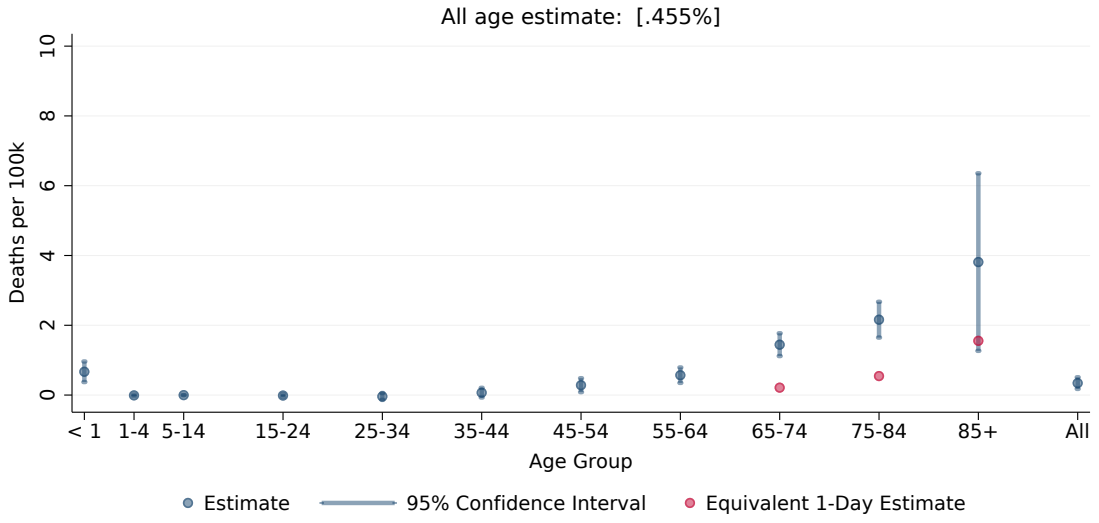
Same-day mortality estimates capture only part of the health impact of pollution exposure. Aside from autocorrelation creating omitted variables bias when estimating same-day effects, health consequences may accumulate or manifest with delay. Failing to account for these dynamics can bias estimates of the effect of a sustained pollution shock.

We address this concern by estimating distributed lag models at monthly frequencies using CDC mortality data covering 1999–2019. The monthly specification allows us to trace out the mortality response over horizons from one month to five years, testing whether short-horizon effects are offset by subsequent reductions in mortality (harvesting) or whether effects accumulate over longer horizons. An additional benefit of using the CDC data is that it includes all deaths, not just those of Medicare beneficiaries.

Figure 6 plots contemporaneous (same-month) mortality effects by age group. Point estimates increase with age, ranging from near zero for children to nearly 4 deaths per 100k

for the 85+ population. The age gradient mirrors that of the daily estimates, confirming that the elderly bear the largest burden of generation-related pollution. We are also able to detect impacts on infants, as well as adults as young as 45. Between these age groups we precisely estimate zero mortality effects. In order to compare the accumulated monthly estimates against the same day estimates in the Medicare data, the figure also plots the 30 times the daily estimates of Table 3. The monthly estimates are larger than the daily estimates, suggesting that the mortality response to a pollution shock is not fully realized on the same day. This is consistent with the idea that pollution causes cardiovascular and respiratory insults that may not be immediately fatal, but can lead to death in subsequent days or weeks. Over all ages, we find that a month of exposure to 1 μg/m^3 of predicted PM2.5 from generation raises mortality by 0.455%.

Figure 6: Contemporaneous Mortality Effect of One Month of Exposure to 1 μg/m^3 PM2.5

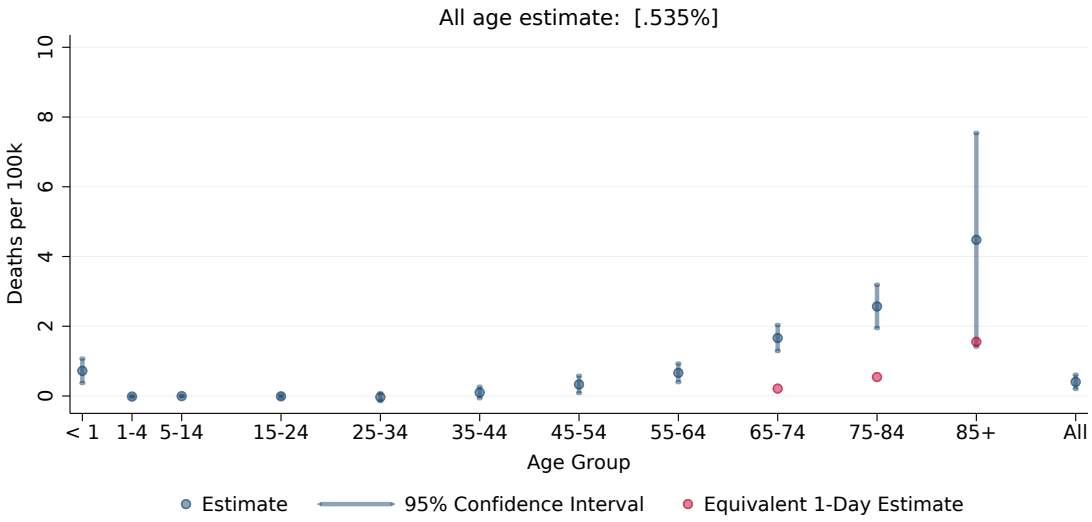


Notes:

Each point represents the coefficient from a separate regression of monthly mortality on contemporaneous predicted PM2.5 from generation. Bands show 95 percent confidence intervals. All specifications include county, year-month, and state-month fixed effects, as well as weather and economic controls.

Figure 7 displays cumulative mortality effects over a three-month horizon, summing coefficients on the contemporaneous month and two monthly lags. For most age groups, the three-month cumulative effect modestly exceeds the contemporaneous effect, indicating that the mortality response continues after the initial exposure.

Figure 7: Cumulative Mortality Effect Over Three Months



Notes:

Cumulative mortality effect of one month of exposure to 1 μg/m^3 of predicted PM2.5, summing coefficients on the contemporaneous month and two monthly lags. Bands show 95 percent confidence intervals.

Figure 8 extends the horizon to five years. The five-year cumulative effect for the 85+ group is nearly 6 deaths per 100k, larger than the contemporaneous or three-month effects. Any displacement is more than offset by the greater likelihood of death following longer-run exposure. This finding is important for welfare calculations, as it implies that pollution-related deaths represent genuine losses of life rather than short-term shifts in timing.

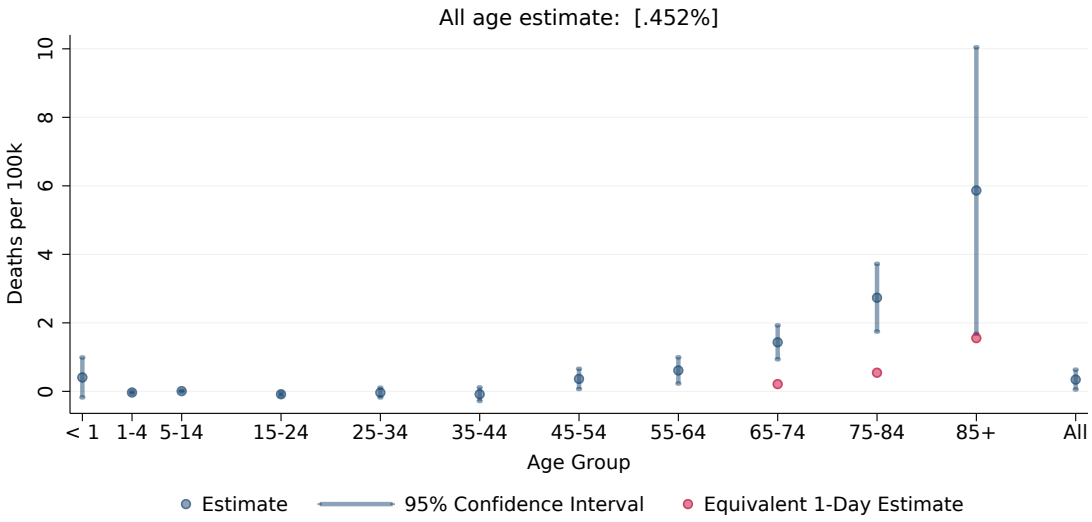
5.5 Aggregate Mortality from Electricity Generation

We combine the mortality estimates with predicted pollution from each plant-day to calculate aggregate deaths attributable to electricity generation. For each county-day, we multiply predicted PM2.5 by age-specific mortality coefficients and age-specific population, then sum across age groups and counties:

$$\text{Deaths}_t = \sum_i \sum_g \widehat{PM}_{it} \cdot \hat{\tau}_g \cdot \text{pop}_{itg} \tag{11}$$

where g indexes age groups and $\hat{\tau}_g$ is the estimated mortality effect for group g.

Figure 8: Cumulative Mortality Effect Over Five Years



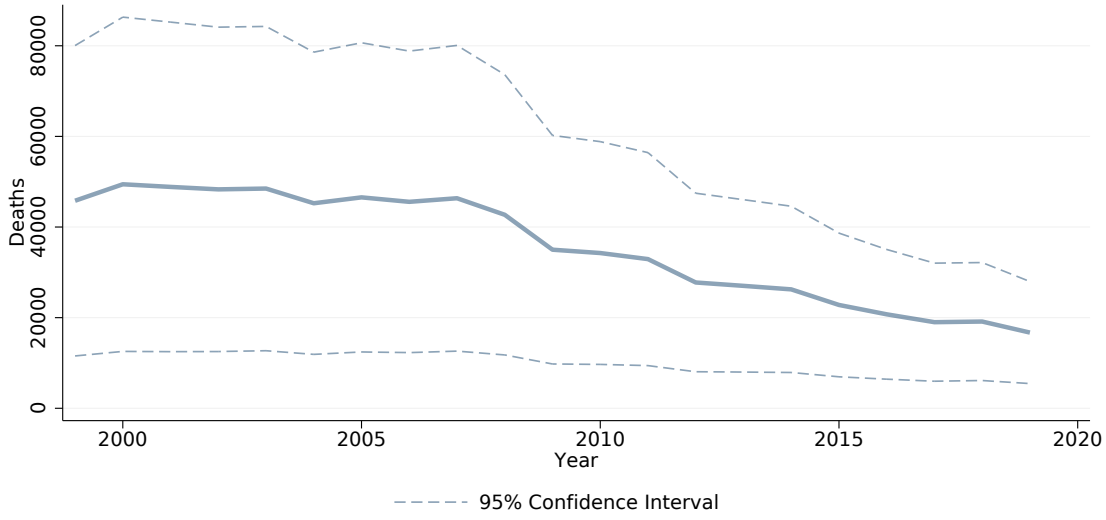
Notes:

Cumulative mortality effect of one month of exposure to 1 μg/m^3 of predicted PM2.5, summing coefficients from zero to five years of lags. Bands show 95 percent confidence intervals.

Figure 9 plots annual deaths attributable to generation-related PM2.5 from 1999 to 2019. At the start of the sample, electricity generation caused nearly 50,000 deaths per year. By 2019, this had fallen to fewer than 20,000 deaths per year. This decline reflects both the fall in coal generation documented in Figure 1 and the retirement of the oldest, dirtiest plants—those without modern pollution controls, and the installation of pollution controls on the plants that remained.

The decline in deaths is not simply a reflection of falling electricity consumption. As documented in Section 3, total electricity generation was essentially flat from 2007 onward. The mortality reduction stems from changes in which plants operated and how much they ran: the shift from coal to natural gas and renewables, and the retirement of uncontrolled coal units. The remaining coal fleet in 2019 was both smaller and cleaner than the fleet in 1999.

Figure 9: Annual Deaths Attributable to Electricity Generation



Notes:

Annual deaths calculated by multiplying county-day predicted PM2.5 from generation by age-specific mortality coefficients and population, then summing across counties and days.

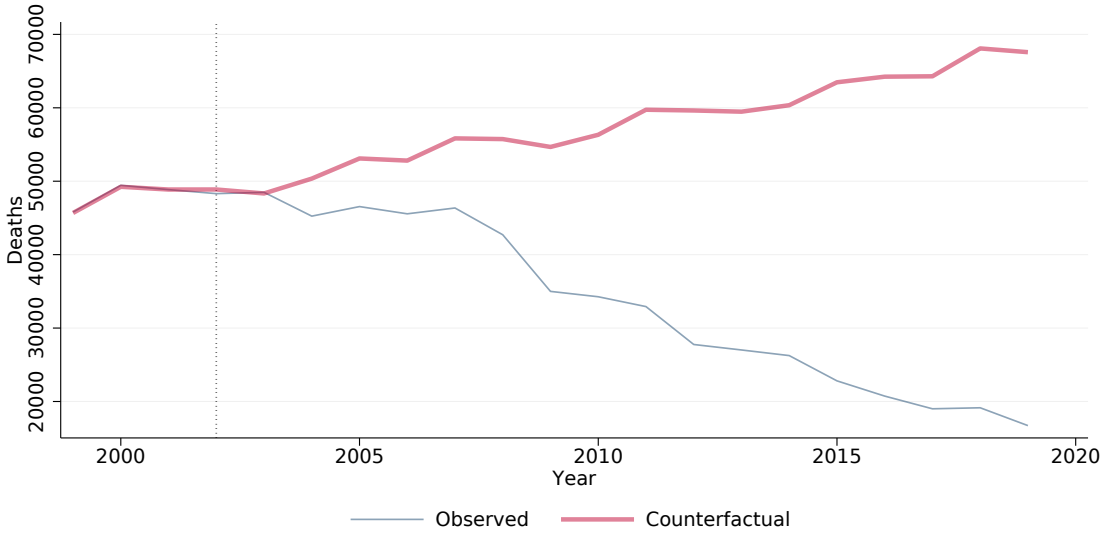
5.6 Counterfactual: The Health Benefits of the Energy Transition

To isolate the contribution of the energy transition to mortality reductions, we construct a counterfactual in which the electricity grid remains frozen at its 2003 configuration. We use random forests to estimate dispatch rules for each PCA separately using unit-level generation data from 1999–2002. This yields a set of predictions for which plants would have operated in each hour given observed electricity demand, fuel prices, and weather. We then apply these rules to the grid as it existed in 2003, had it experienced the observed patterns of load observed 2003–2019. In effect, we ask what pollution would have been if the grid had been frozen in its 2003 state. We then apply these counterfactual generation quantities to the HYSPLIT dispersion predictions produced earlier to calculate pollution exposure.

Figure 10 compares observed and counterfactual mortality. The two series track each other closely through 2003. Note that the production patterns in 2003 are generated out of sample from the rules estimated in 1999-2002. Thereafter, the gap widens as the actual grid becomes progressively cleaner than the frozen 2003 grid. By 2019, the counterfactual grid would have caused approximately 70,000 deaths per year. This figure is higher than the 1999

level despite using a cleaner 2003 grid because the U.S. population aged substantially over this period: the population aged 65 and older grew by more than 50 percent between 2003 and 2019, and the population aged 85 and older—the group most vulnerable to pollution—grew even faster. Under the counterfactual, this demographic shift would have increased the death toll from electricity generation.

Figure 10: Observed vs. Counterfactual Mortality from Electricity Generation



Notes:

The thin line shows estimated deaths from observed generation; the thick line shows deaths from counterfactual generation holding the grid fixed at its 2003 configuration.

The difference between the two series represents deaths prevented by the energy transition. By 2019, this gap reached approximately 50,000 deaths per year. Valuing the associated life-years lost at the EPA's value of a statistical life-year yields annual health benefits on the order of $50 billion per year by 2019. These benefits accrued disproportionately to counties downwind of major coal plants, particularly in the Ohio River Valley and the industrial Midwest, where both coal generation and population density are high.

The shift from coal to gas and renewables was driven primarily by market forces—declining natural gas prices following the shale revolution and falling costs of wind and solar generation—rather than by policies explicitly targeting health outcomes (Coglianese, Gerarden, and Stock, 2020). These market-driven changes produced health benefits as a

byproduct. Second, our counterfactual holds electricity consumption fixed at observed levels. The relevant comparison is not between generating electricity and not generating it, but between generating it with the 2003 fleet and generating it with the fleet that actually operated in each year. This comparison isolates the health value of cleaner generation.

6 Conclusion

This paper develops an integrated framework for estimating the mortality effects of power plant pollution and applies it to measure the health benefits of the U.S. energy transition. By combining hourly generation data, atmospheric dispersion modeling, and quasi-experimental variation from the reshuffling of electricity production across plants, we trace emissions from individual power plants to pollution exposure in downwind counties and estimate the causal effect on mortality.

Our estimates indicate that PM2.5 from electricity generation substantially increases mortality. Among Medicare beneficiaries, a 1 μg/m^3 increase in same-day PM2.5 raises mortality by 0.14 percent. The IV estimates exceed OLS estimates by roughly 50 percent, consistent with either measurement error attenuation or greater toxicity of power plant emissions relative to ambient PM2.5 from other sources. Extending the analysis to unmonitored counties reveals larger effects, suggesting that rural populations—who lack nearby monitors and face greater barriers to medical care—bear a disproportionate burden from power plant pollution.

The distributed lag estimates address a limitation of same-day specifications: the inability to distinguish mortality displacement from genuine life-years lost. Cumulative effects over five years substantially exceed contemporaneous effects, indicating that pollution causes deaths that would not otherwise have occurred. We therefore find that the mortality impacts accumulate over time, rather than simply hastening imminent deaths by a few days or weeks.

Applying these estimates to predicted pollution from each plant yields aggregate mor-

tality figures. Electricity generation caused nearly 50,000 deaths per year at the start of our sample in 1999. By 2019, this had fallen below 20,000—a reduction driven by the shift from coal to natural gas and renewables, the retirement of the oldest and dirtiest plants, and the installation of pollution controls on the coal fleet that remained. The decline in generation-related PM2.5 accounts for 40 percent of the total improvement in ambient air quality over the sample period.

Our counterfactual analysis isolates the contribution of the energy transition. Had the grid remained frozen at its 2003 configuration, deaths from electricity generation would have risen to approximately 70,000 per year by 2019, as an aging population faced continued exposure to pollution from coal-fired plants. The difference—roughly 50,000 deaths prevented annually—represents the mortality benefit of cleaner generation. Valued at the EPA's estimate of a statistical life-year, these benefits reached approximately $50 billion per year by the end of the sample.

The health benefits we document were largely a byproduct of market forces rather than policies targeting public health. Declining natural gas prices following the shale revolution made gas-fired generation more competitive than coal; falling costs of wind and solar expanded the role of renewables. These market-driven changes produced substantial improvements in air quality and reductions in mortality for large swathes of the country.

These findings have implications for energy and environmental policy. The substantial health benefits of the coal-to-gas transition suggest that policies accelerating the retirement of remaining coal plants—particularly those without modern pollution controls—would generate meaningful public health gains. We highlight that these benefits are due entirely to the reduction in exposure to local pollutants; the climate benefits of reduced coal generation are not included in our estimates.

References

Anderson, Michael L. 2020. "As the Wind Blows: The Effects of Long-Term Exposure to Air Pollution on Mortality." *Journal of the European Economic Association*, 18(4): 1886–1927.

Barreca, Alan I., Karen Clay, Olivier Deschênes, Michael Greenstone, and Joseph S. Shapiro. 2016. "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the 20th Century." *Journal of Political Economy*, 124(1): 105–159. arXiv: 1011.1669v3 ISBN: 9781604138795.

Barreca, Alan I., Matthew Neidell, and Nicholas J. Sanders. 2021. "Long-run pollution exposure and mortality: Evidence from the Acid Rain Program." *Journal of Public Economics*, 200: 104440.

Cabral, Marika, and Marcus Dillender. 2024. "Air Pollution, Wildfire Smoke, and Worker Health."

Carleton, Tamma, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Robert E Kopp, Kelly E McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon Seo, Arvid Viaene, Jiacan Yuan, and Alice Tianbo Zhang. 2022. "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits." *The Quarterly Journal of Economics*, 137(4): 2037–2105.

Chay, Kenneth Y, and Michael Greenstone. 2003*a*. "Air quality, infant mortality, and the Clean Air Act of 1970." Publisher: National Bureau of Economic Research Cambridge, Mass., USA Publication Title: Environmental Protection Issue: 10053.

Chay, Kenneth Y, and Michael Greenstone. 2003*b*. "The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession." *Quarterly Journal of Economics*, 118(3): 1121–1167.

Cicala, Steve. 2022. "Imperfect Markets versus Imperfect Regulation in U.S. Electricity Generation." *American Economic Review*, 112(2).

Clay, Karen, Akshaya Jha, Nicholas Muller, and Randall Walsh. 2019. "External Costs of Transporting Petroleum Products: Evidence from Shipments of Crude Oil from North Dakota by Pipelines and Rail." *The Energy Journal*, 40(1): 55–72.

Coglianese, John, Todd D. Gerarden, and James H. Stock. 2020. "The Effects of Fuel Prices, Environmental Regulations, and Other Factors on U.S. Coal Production, 2008-2016." *The Energy Journal*, 41(1): 55–82.

Currie, J., and M. Neidell. 2005. "Air Pollution and Infant Health: What Can We Learn from California's Recent Experience?" *The Quarterly Journal of Economics*, 120(3): 1003–1030.

Currie, Janet, and W. Reed Walker. 2011. "Traffic Congestion and Infant Health: Evidence from E-ZPass." *American Economic Journal-Applied Economics*, 3(1): 65–90.

Currie, Janet, Michael Greenstone, and Katherine Meckel. 2017. "Hydraulic fracturing and infant health: New evidence from Pennsylvania." *Science Advances*, , (3): 1–9.

Deryugina, Tatyana, and Julian Reif. 2023. "The Long-run Effect of Air Pollution on Survival." National Bureau of Economic Research w31858, Cambridge, MA.

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

Dodier, Matthew. 2024. "Smoke Signals: Examining the impact of smoke plumes on hospital care use in the state of Oregon."

Finkelstein, Amy, Matthew Notowidigdo, and Steven Shi. 2026. "Trading Goods for Lives: NAFTA's Mortality Impacts and Implications." National Bureau of Economic Research w34855, Cambridge, MA.

Hausman, Catherine H., and Ryan Kellogg. 2015. "Welfare and Distributional Implications of Shale Gas." *Brookings Papers on Economic Activity*.

Henneman, Lucas, Christine Choirat, Irene Dedoussi, Francesca Dominici, Jessica Roberts, and Corwin Zigler. 2023. "Mortality risk from United States coal electricity generation." *Science*, 382(6673): 941–946.

Henneman, Lucas R.F., Christine Choirat, Cesunica Ivey, Kevin Cummiskey, and Corwin M. Zigler. 2019. "Characterizing population exposure to coal emissions sources in the United States using the HyADS model." *Atmospheric Environment*, 203(February): 271–280.

Hill, Elaine L. 2018. "Shale gas development and infant health : Evidence from Pennsylvania." *Journal of Health Economics*, 61: 134–150.

Hogan, Dylan, and Wolfram Schlenker. 2024. "Empirical approaches to climate change impact quantification." In *Handbook of the Economics of Climate Change*. Vol. 1, 53–111. Elsevier.

Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J Yates. 2016. "Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors." *American Economic Review*, 106(12): 3700–3729. ISBN: 8024432080.

Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates. 2020. "Decompositions and Policy Consequences of an Extraordinary Decline in Air Pollution from Electricity Generation." *American Economic Journal: Economic Policy*, 12(4): 244–274.

Jarvis, Stephen, Olivier Deschenes, and Akshaya Jha. 2022. "The Private and External Costs of Germany's Nuclear Phase-Out." *Journal of the European Economic Association*, 20(3): 1311–1346.

Knittel, Christopher R., Douglas L. Miller, and Nicholas J. Sanders. 2016. "Caution, Drivers! Children Present: Traffic, Pollution, and Infant Health." *Review of Economics and Statistics*, 98(2): 350–366.

Lippmann, Morton. 2014. "Toxicological and epidemiological studies of cardiovascular effects of ambient air fine particulate matter ($PM_{2.5}$) and its chemical components: Coherence and public health implications." *Critical Reviews in Toxicology*, 44(4): 299–347.

Lippmann, Morton, Lung Chi Chen, Terry Gordon, Zakuhiko Ito, and George D. Thurston. 2013. "National Particle Component Toxicity (NPACT) Initiative: integrated epidemiologic and toxicologic studies of the health effects of particulate matter components." Health Effects Institute 177.

Mendelsohn, Robert. 1980. "An economic analysis of air pollution from coal-fired power plants." *Journal of Environmental Economics and Management*, 7(1): 30–43.

Miller, Nolan H., David Molitor, and Eric Zou. 2024. "The Nonlinear Effects of Air Pollution on Health: Evidence from Wildfire Smoke."

Muller, Nicholas Z, and Robert Mendelsohn. 2007. "Measuring the damages of air pollution in the United States." *Journal of Environmental Economics and Management*, 54(1): 1–14.

Muller, Nicholas Z, Robert Mendelsohn, and William D. Nordhaus. 2011. "Environmental Accounting for Pollution in the United States Economy." *The American Economic Review*, 101(5): 1649–1675.

National Research Council. 2002. *Estimating the Public Health Benefits of Proposed Air Pollution Regulations.* Washington, D.C.:National Academies Press. Pages: 10511.

Sanders, Nicholas J., Alan I. Barreca, and Matthew J. Neidell. 2020. "Estimating Causal Effects of Particulate Matter Regulation on Mortality." *Epidemiology*, 31(2): 160–167.

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

Sullivan, Daniel. 2017. "The True Cost of Air Pollution: Evidence from the Housing Market."

U.S. Environmental Protection Agency. 2009. "Final Report: Integrated Science Assessment for Particulate Matter." Washington, D.C. Issue: December 2009.

U.S. Environmental Protection Agency. 2020. "Estimating Particulate Matter Emissions for eGRID."

U.S. Environmental Protection Agency. 2024*a*. "Estimating PM2.5- and Ozone-Attributable Health Benefits: 2024 Update." Office of Air and Radiation, Research Technical Support Document, Triangle Park, NC.

U.S. Environmental Protection Agency. 2024*b*. "Final Regulatory Impact Analysis for the Reconsideration of the National Ambient Air Quality Standards for Particulate Matter." Office of Air Quality Planning and Standards EPA-452/R-24-006, Research Triangle Park, NC.

The authors are grateful for access to the Medicare data through grant NIA P01 AG005842 from the National Institute on Aging. We also appreciate support from the Tufts University High Performance Compute Cluster the University of Chicago's Research Computing Center, and the NBER. The authors have no financial or non-financial conflicts of interest to disclose.