APPENDICES AVAILABLE ON THE HEI WEB SITE

Research Report 187

Causal Inference Methods for Estimating Long-Term Health Effects of Air Quality Regulations

Corwin M. Zigler et al.

Appendix A. Technical Details for Case Study 1: PM_{10} Nonattainment Designations

Appendix B. Sensitivity Analysis to the Pruning of Observations in Case Study 1

Appendix C. Technical Details for Case Study 2: Power-Plant Emissions Controls

Appendix D. Results from the Power-Plant Case Study with 75-km Data Linkage

Appendix E. Preliminary Extension of the Power-Plant Case Study to Health Outcomes

These appendices were reviewed by the HEI Health Review Committee but did not undergo the full HEI scientific editing and production process. They were reviewed solely for spelling, grammar, and cross-references to the main text.

Correspondence may be addressed to Dr. Corwin Matthew Zigler, Harvard T.H. Chan School of Public Health, Building 2, 4th Floor, 655 Huntington Ave., Boston, MA 02115; e-mail: czigler@hsph.harvard.edu.

Although this document was produced with partial funding by the United States Environmental Protection Agency under Assistance Award CR–83467701 to the Health Effects Institute, it has not been subjected to the Agency's peer and administrative review and therefore may not necessarily reflect the views of the Agency, and no official endorsement by it should be inferred. The contents of this document also have not been reviewed by private party institutions, including those that support the Health Effects Institute; therefore, it may not reflect the views or policies of these parties, and no endorsement by them should be inferred.

Appendices for Causal Inference Methods for Estimating Long Term Health Effects of Air Quality Regulations

Corwin Matthew Zigler* Chanmin Kim * Christine Choirat † John Barret Hansen†

Yun Wang * Lauren Hund ‡ Jonathan Samet§ Gary King†

Francesca Dominici*

*Harvard TH Chan School of Public Health
†Harvard University
‡University of New Mexico
§Keck School of Medicine, University of Southern California

Contents

Appendix A. Technical Details for Case Study 1: PM_{10} Nonattainment Designations

A.1 Missing Covariate Data in the Analysis of Case Study 1

Rather than exclude observations with missing covariate data, we employ spatial hierarchical model to impute the missing 1990 ambient PM_{10} measurements and use these imputations in our analysis. Specifically, we fit the same type of spatial hierarchical model as detailed below, with the outcome specified as the log-transformed PM_{10} concentration during the year 1990. Covariates included in this model were the same as those in the model of the main text (See Table 1 of the main text), with the addition of a covariate denoting whether a location lies in an attainment or nonattainment area. In total, ambient average PM_{10} during 1990 was imputed for 131 nonattainment and 153 attainment locations in the entire (non-pruned) sample using posterior-predictive means from this model and treated as fixed covariates in the analysis.

A.2 Locations of Discarded Monitoring Locations

As discussed in the main text, the propensity score strategy identified 52 locations that did not "overlap" with locations in the opposite treatment group. These locations were discarded for the analysis. Figure A.1 depicts the locations of all monitoring locations, the locations retained for the analysis, and the locations of those excluded due to lack of overlap.

Figure A.1: Locations of all 547 PM_{10} monitoring locations available for analysis and for the 495 locations retained after propensity score pruning

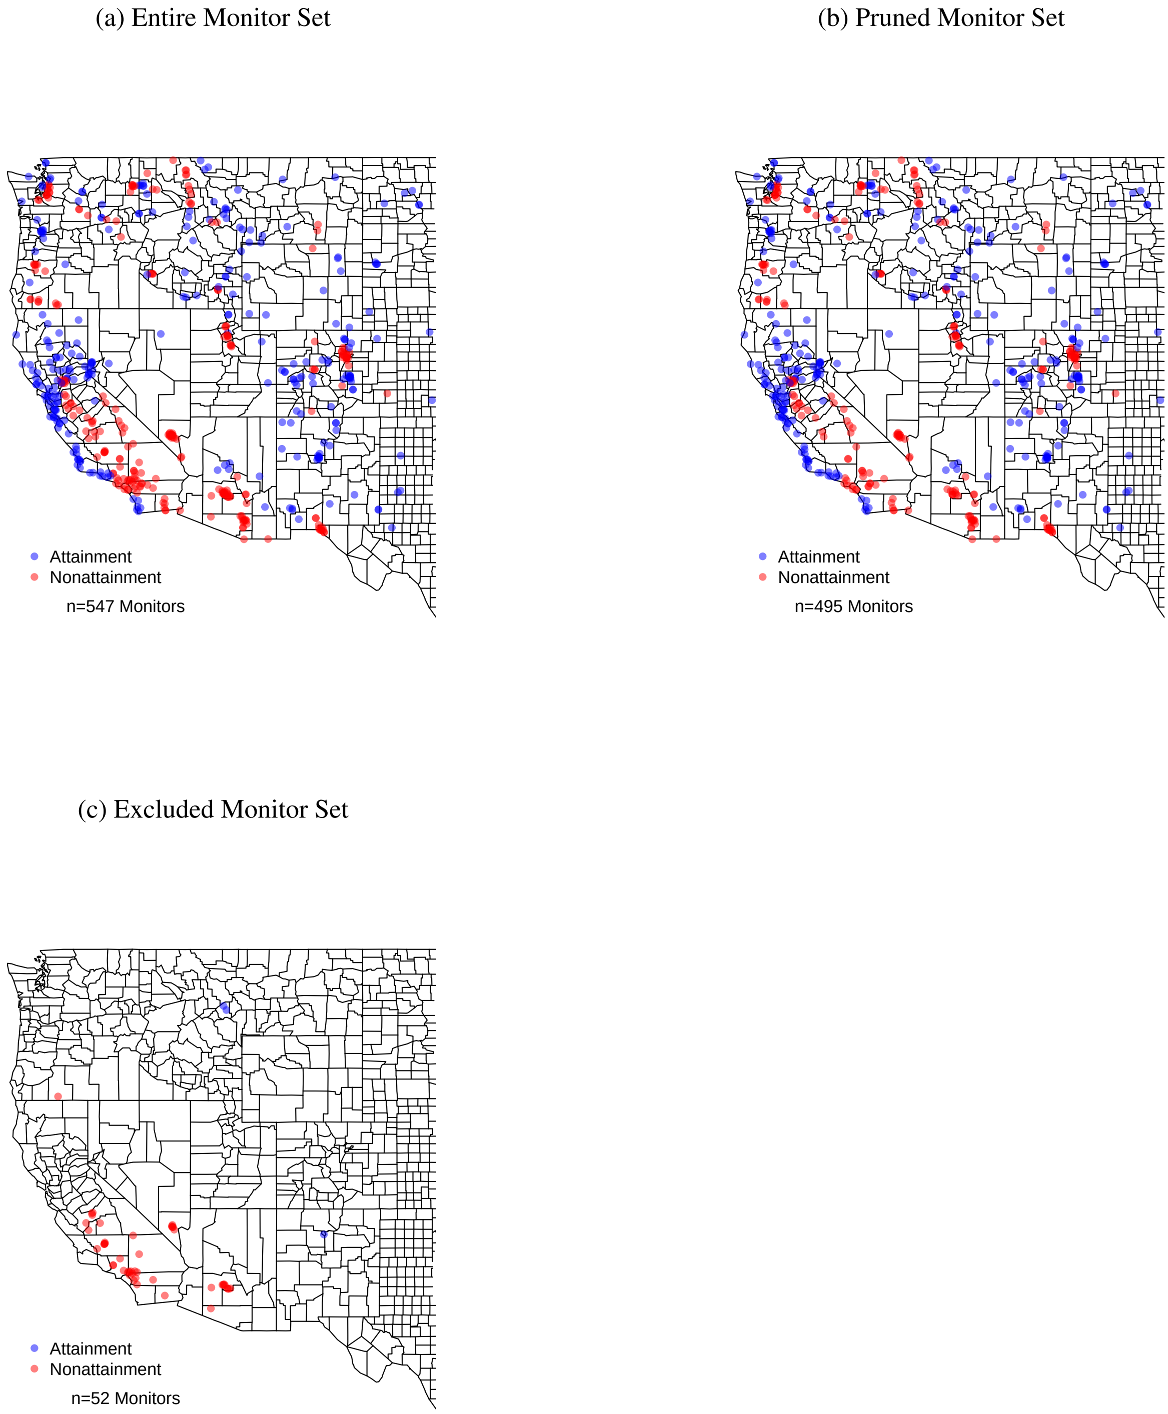

A.3 Models for Air Pollution and Medicare Health Outcomes

A.3.1 Notation and Ignorability Assumption

For any hypothetical allocation of nonattainment designations to the $n = 495$ pruned locations, let $\mathbf{A} \equiv [A(s_i)]_{i=1}^{n}$ be the vector of indicators denoting whether each of $n = 495$ monitoring locations would fall in a nonattainment county, with $A(s_i) = 1$ denoting nonattainment for the i^{th} location, and $A(s_i) = 1$ denoting attainment. We refer to the entire vector $\mathbf{A}$ as a regulation program. We denote a specific regulation program with $\mathbf{A} = \mathbf{a}$, and the observed program representing the actual allocation of nonattainment designations with $\mathbf{A} = \mathbf{a^{obs}}$.

Let $Y_{\mathbf{a}}(s)$ denote health outcome in 2001 (either all-cause mortality, CVD hospitalizations, or respiratory hospitalizations) at location s that would potentially occur under regulation program $\mathbf{A} = \mathbf{a}$. Let $X_{\mathbf{a}}(s)$ denote the average ambient PM_{10} concentration that would potentially be observed during 1999-2001 under regulation program $\mathbf{A} = \mathbf{a}$. Note that the only observed potential outcomes are $(X_{\mathbf{a^{obs}}}(s), Y_{\mathbf{a^{obs}}}(s))$; all others are considered missing data. Let $Z(s)$ denote the vector of covariate values for location s (i.e., those listed in Table 1 of the main text), and also assume that $Z(s)$ contains, in addition to those covariates, indicators of propensity score subclass membership.

We confine attention to the regulation programs $\mathbf{A} = \mathbf{0}$ and $\mathbf{A} = \mathbf{a}^{\mathbf{obs}}$ and define the monitor-level causal effect of $\mathbf{A}$ on ambient PM_{10} as the comparison between the potential pollution concentration of that pollutant under the observed nonattainment designations, $X_{a^{obs}}(s)$, and the potential concentration under the setting with no nonattainment regulations, $X_0(s)$, among the nonattainment areas in the pruned data set. We similarly define the causal effect of nonattainment designations on a given health outcome for location s as the comparison between $Y_{a^{obs}}(s)$ and $Y_0(s)$.

We assume that assignment to the initial nonattaiment designations is strongly ignorable conditional on covariates. In other words, there is no unmeasured confounding in the sense that $Z(s)$ contains all factors that tend to differ between nonattainment and attainment locations and that also impact potential pollution and Medicare health outcomes.

A.3.2 Spatial Hierarchical Model for Air Pollution

For $f(X_0(s), X_{a^{obs}}(s)|Z(s))$, we propose the following spatial hierarchical model:

$$X(s) = Z^T(s)\beta + W(s) + \varepsilon(s), \tag{1}$$

where $X(s) = (X_0(s), X_{a^{obs}}(s))^T$ is the vector of potential pollution concentrations under each possible nonattainment status, $W(s)$ is a vector of spatially-varying random intercepts, and $\varepsilon(s)$ represents nonspatial "nugget" error (e.g., measurement error). We assume $\varepsilon(s) \sim MVN(0, \Psi)$, and Ψ diagonal. $Z^T(s)$ is a $2 \times p$ matrix of time-invariant covariates, where p is the number of covariates (including indicators of propensity score subclass) used to adjust for confounding. The analysis presented here assumes that, conditional on $Z(s)$, potential pollution concentrations under $\mathbf{A} = \mathbf{a^{obs}}$ and $\mathbf{A} = \mathbf{0}$ are conditionally independent. Details of the spatial correlation structure can be found in Zigler et al. (2012) and Banerjee et al. (2008).

A.3.3 Log-linear Model for Mortality

For $f(Y_0(s), Y_{a^{obs}}(s)|X_0(s), X_{a^{obs}}(s), Z(s))$, we make use of two additional assumptions. We assume conditional independence of potential health outcomes, conditional on covariates and air pollution: $Y_0(s) \perp\!\!\!\perp Y_{a^{obs}}(s) | X_0(s), X_{a^{obs}}(s), Z(s)$.

We also assume that under a given designation, after conditioning on pollution under that designation (and covariates), health outcomes are independent of pollution under the opposite designation: $f(Y_\mathbf{a}(s)\big|X_0(s), X_{a^{obs}}(s), Z(s)) = f(Y_\mathbf{a}(s)\big|X_\mathbf{a}(s), Z(s))$, for $\mathbf{a} = 0, a^{obs}$. This assumption reflects a belief that knowledge of both $(X_0(s), X_{a^{obs}}(s))$ does not contribute any information pertaining to $Y_\mathbf{a}(s)$ above and beyond that contained in $X_\mathbf{a}(s)$ alone. As a result of these assumptions, we write $f(Y_0(s), Y_{a^{obs}}(s)|X_0(s), X_{a^{obs}}(s), Z(s)) = \prod f(Y_\mathbf{a}(s)|X_\mathbf{a}(s), Z(s))$ for $\mathbf{a} = 0, a^{obs}$, and model the terms of this product with the following log-linear models:

$$\log(E[Y_\mathbf{a}(s)]) = \alpha_0^\mathbf{a} + Z^T(s)\alpha_1^\mathbf{a} + X_\mathbf{a}(s)\alpha_2^\mathbf{a} + \log(N(s)), \tag{2}$$

where $\mathbf{a} = 0, a^{obs}$, $N(s)$ is the total number of Medicare beneficiaries (for the mortality outcome) or person-years (for the hospitalization outcomes) living near location s , $\alpha_1^{\mathbf{a}}$ captures relative risks associated with differences in $Z(s)$ under nonattainment program $\mathbf{A} = \mathbf{a}$, and $\alpha_2^{\mathbf{a}}$ captures relative risks associated with differences in post-regulation ambient pollution concentrations under regulation program $\mathbf{A} = \mathbf{a}$.

A.4 Bayesian Estimation

Recall that $X(s) = (X_0(s), X_{a^{obs}}(s))^T$ and let $Y(s) = (Y_0(s), Y_{a^{obs}}(s))^T$. The full joint density of the data can be written as:

$$f(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) = \int \prod_{i=1}^{n} f(Z(s_i), X(s_i), Y(s_i) \big| \theta) p(\theta) d\theta, \tag{3}$$

where θ is a generic parameter with prior distribution $p(\theta)$. Distinguishing between the missing (*mis*) and observed (*obs*) quantities in $X(s)$ and $Y(s)$, the posterior distribution of θ is proportional to:

$$p(\theta) f(\mathbf{Z}) \int \int \prod_{i=1}^{n} f(X^{mis}(s_i), X^{obs}(s_i), Y^{mis}(s_i), Y^{obs}(s_i) \big| Z(s_i), \theta) dY^{mis}(s_i) dX^{mis}(s_i). \tag{4}$$

Inference from (4) is difficult because of the integration over missing potential outcomes, leading us to focus instead on the following joint posterior distribution:

$$p(\theta, \mathbf{X^{mis}}, \mathbf{Y^{mis}} \big| \mathbf{X^{obs}}, \mathbf{Y^{obs}}, \mathbf{Z}) \propto p(\theta) \prod_{i=1}^{n} f(X^{mis}(s_i), X^{obs}(s_i), Y^{mis}(s_i), Y^{obs}(s_i) | Z(s_i), \theta), \tag{5}$$

which is convenient for its proportionality to the standard posterior distribution of θ had all of the potential outcomes been observed (Jin and Rubin, 2008). Thus, our computational strategy will consist of a Markov chain Monte Carlo (MCMC) data augmentation algorithm that iteratively samples missing potential outcomes conditional on observed data and parameters, then samples parameters and calculates causal effect estimates conditional on "complete" data with identified principal strata.

A.4.1 Prior Distributions and Outline of MCMC Strategy

The mechanics of model (1) rely on two key features: the relationship among potential pollution concentrations within a location, and the decay of their correlations across space. For the relationship among pollutants within a location, note that the cross covariance within a location, $K(s,s)$, is in fact a covariance matrix of the 2 random effects corresponding to potential pollution concentrations at a common location. We write $K(s,s) = LL^T$, where L is the lower-triangular Cholesky square root of this covariance matrix, and assume that $K(s,s)$ is the same for all s, that is, that the process is stationary.

For the spatial decay, we define a simpler MVGP, $\tilde{W}(s)$, such that $Var(\tilde{W}_k(s)) = 1$ and the cross covariance is diagonal: $\tilde{K}(s_i,s_j;\nu) = diag\{\rho_k(s_i,s_j;\nu_k)\}$, where $\rho_k(s_i,s_j;\nu_k)$ represents a function for the spatial decay of the correlation between the k^{th} element of $\tilde{W}(s)$ across space. We assume isotropic exponential covariance functions that depend only on the Euclidean distance between locations s_i and s_j ($||s_i - s_j||$), with $\rho_k(s_i,s_j) = e^{-\nu_k||s_i-s_j||}$. The covariance matrix of $\tilde{W} = [\tilde{W}(s_i)]_{i=1}^n$ can be written as $\Sigma_{\tilde{W}} = [\tilde{K}(s_i,s_j)]_{i,j=1}^n$.

Rather than model $K(s_i,s_j;\nu)$ directly, we separately specify $\tilde{K}(s_i,s_j;\nu)$ and LL^T, and define $W(s) = L\tilde{W}(s)$, which implies that the spatial random effects in (1) are a linear transformation of the simpler MVGP, with transformation defined by the relationships among the pollutants. With this specification, $K(s_i,s_j;\nu) = L\tilde{K}(s_i,s_j;\nu)L^T$.

Let $\mathbf{X}^T = [(X_0(s_i)^T, X_1(s_i)^T)]_{i=1}^n$ be the $2n \times 1$ pollution vector and $\mathcal{Z}$ be the $2n \times p$ matrix of regressors. The above model can, after marginalization over $\tilde{W}$, be equivalently stated as

$$\mathbf{X} \sim MVN(\mathcal{Z}\beta, \mathcal{L}\Sigma_{\tilde{W}}\mathcal{L}^T + I_n \otimes \Psi) \tag{6}$$

where $\mathcal{L} = I_n \otimes L$, and $\otimes$ is the Kronecker product. Details for this model formulation as well as generalizations can be found in Zigler et al. (2012), Wackernagel (2003), Finley et al. (2007), and Banerjee et al. (2008).

For the MCMC, $K(s,s)$ is updated via updates of $L_\mathbf{a}$, which are the lower-triangular Cholesky roots of the $q \times q$ diagonal blocks of $K(s,s)$ that are informed by the observed data. The off-diagonal blocks of $K(s,s)$ are updated according to the pre-specified value of a sensitivity parameter, ω, and the values of $L_\mathbf{a}$, subject to a positive-definiteness constraint. For this analysis, we fix $\omega = 0$. Under the prior specification detailed below, the posterior distribution for β is multivariate normal, with samples drawn using a fully conditional Gibbs step. All other parameters and missing data are updated with a Metropolis step using normal proposal distributions, with appropriate transformations for all variables having restricted support. Each diagonal element of Ψ, each ν, and each missing quantity are updated individually, with block updating carried out for β, $\alpha^0 = (\alpha_0^0, \alpha_1^0, \alpha_2^0)$, $\alpha^1 = (\alpha_0^1, \alpha_1^1, \alpha_2^1)$, and (L_0, L_1). For each model, MCMC chains are run for 32,000 iterations. After discarding the first 5,000 iterations as burn in, inference is based on every 10^{th} posterior sample.

As pointed out in Finley et al. (2007), values of ν are only weakly identifiable and require reasonably informative priors for satisfactory MCMC behavior, but the model decomposition described above entails adequate structure to identify $\Sigma_{\tilde{W}}, L, \Sigma_W$, and Ψ. We treat the parameters $\beta, \Psi, L_0, L_1, \nu, \alpha^0$, and α^1, as *a priori* independent. We specify flat priors for β, α^0, and α^1. For the diagonal elements of Ψ, we specify independent inverse-gamma distributions with shape parameters set to 2 and scale parameters set to 0.5. For ν_k, we specify uniform prior distributions on the interval $(0.45, 3.38)$. Parameters for the prior distributions of Ψ and ν are meant to reflect diffuse prior information within the range of plausible parameter values. For each diagonal element of $K(s,s)$, we specify an inverse-gamma prior distribution with shape parameter set to 2 and scale parameter set to 0.5.

A.5 Assumption about Interference between Locations

Mortality outcomes and pollution levels are only observed under the program $\mathbf{A} = \mathbf{a^{obs}}$. Therefore, we require assumptions to relate observed potential outcomes to those that would have been observed under the hypothetical scenario with no nonattainment designations, which we denote with $\mathbf{A} = \mathbf{0}$. Typically, this would be achieved with the assumption of no interference between observational units (or the Stable Unit Treatment Value Assumption, Rubin [1980]), which states that potential outcomes for a given location are unrelated to designations of all other locations. This assumption implies that there are exactly two sets of potential outcomes for each location: pollution and mortality if that location is regulated, and pollution and mortality if that location is unregulated. Thus, with no interference, potential outcomes under any hypothetical program $\mathbf{A} = \mathbf{a}$ could be considered on a location-by-location basis, with $(X_{\mathbf{a}}(s), Y_{\mathbf{a}}(s)) = (X_{\mathbf{a^{obs}}}(s), Y_{\mathbf{a^{obs}}}(s))$ as long as $\mathbf{a}$ and $\mathbf{a^{obs}}$ entail the same regulation status for location s.

In studies of air pollution, however, the assumption of no interference does not likely hold because regulations at a given location likely impact air quality at nearby locations. Thus, knowing the observed potential outcomes at location s under $\mathbf{A} = \mathbf{a^{obs}}$ does not imply knowledge of the potential outcomes under any other $\mathbf{A} = \mathbf{a}$ because potential outcomes for s can differ when regulations are allocated differently to other locations. In fact, with no assumptions regarding interference, potential outcomes for each location are distinctly defined for each different possible regulation program because changing the regulation designation of any location could impact potential outcomes at all other locations.

We liken investigation of the nonattainment designations to previously-considered problems of "partial interference" (Sobel, 2006) where observations within a clearly-defined group (e.g., residents of a particular neighborhood) interfere with one another, but observations in different groups (e.g., residents of distant neighborhoods) do not. Unlike previously considered partial-interference settings, there are no clearly defined interference sets for analyzing the PM_{10} nonattainment designations (e.g., assuming no interference between locations in different counties might be too restrictive, especially for observations near county borders). We argue that a unique feature of the present context is that nonattainment designations were "assigned" with some implicit regard to interference because one criterion for a nonattainment designation was contribution to a NAAQS violation in a nearby area. That is, if weather patterns or mere proximity led pollution in one location to affect pollution in another location, the EPA ensured that these two locations shared the same regulation designation.

Let $\mathcal{R}^{\mathbf{a^{obs}}}$ and $\mathcal{U}^{\mathbf{a^{obs}}}$ respectively denote the set of 219 nonattainment (i.e., "regulated") and 276 attainment (i.e., "unregulated") locations under the program $\mathbf{A} = \mathbf{a^{obs}}$. We adopt what we term the *assignment group interference assumption (AGIA)* to reflect the notion that locations within $\mathcal{R}^{\mathbf{a^{obs}}}$ do not interfere with those in $\mathcal{U}^{\mathbf{a^{obs}}}$. Thus, changing the regulation designation of any location in $\mathcal{U}^{\mathbf{a^{obs}}}$ would not change the potential outcomes of locations in $\mathcal{R}^{\mathbf{a^{obs}}}$ (and *vice versa*). A consequence of this assumption is that $(X_{\mathbf{1}}(s), Y_{\mathbf{1}}(s)) = (X_{\mathbf{a^{obs}}}(s), Y_{\mathbf{a^{obs}}}(s))$ for $s \in \mathcal{R}^{\mathbf{a^{obs}}}$ and $(X_{\mathbf{0}}(s), Y_{\mathbf{0}}(s)) = (X_{\mathbf{a^{obs}}}(s), Y_{\mathbf{a^{obs}}}(s))$ for $s \in \mathcal{U}^{\mathbf{a^{obs}}}$. Figure A.2 graphically depicts the implication of AGIA. In practice, this assumption implies that, in all attainment areas ($\mathcal{U}^{\mathbf{a^{obs}}}$), observed potential outcomes are the same as those that would have occurred if the EPA had not designated any other area (i.e., if there were no nonattainment designations). The AGIA is discussed further in Zigler et al. (2012).

Figure A.2: Structure of potential outcomes for different regulation programs under the Assignment Group Interfer-ence Assumption. Points represent monitor locations in counties contained in portions of California and Arizona.

(a) $\mathbf{A} = \mathbf{a}^{\mathbf{obs}}$. All pollution and mortality outcomes are observed under this program.

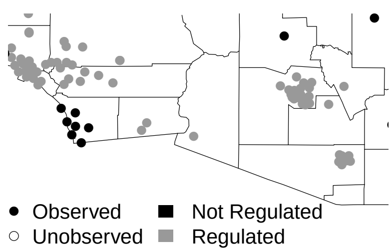

(b) $\mathbf{A} = \mathbf{0}$. Pollution and mortality outcomes are observed for locations in $\mathcal{U}^{\mathbf{a}^{\mathbf{obs}}}$ and unobserved for those in $\mathcal{R}^{\mathbf{a}^{\mathbf{obs}}}$.

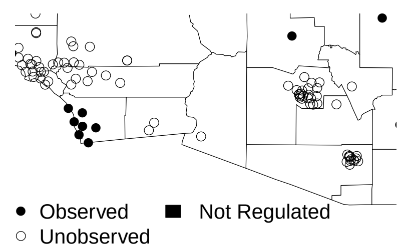

(c) $\mathbf{A} = \mathbf{1}$. Pollution and mortality outcomes are observed for locations in $\mathcal{R}^{\mathbf{a}^{\mathbf{obs}}}$ and unobserved for those in $\mathcal{U}^{\mathbf{a}^{\mathbf{obs}}}$.

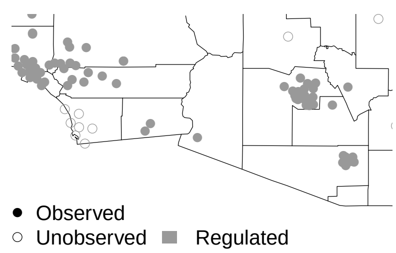

Appendix B. Sensitivity Analysis to the Pruning of Observations in Case Study 1

As a sensitivity analysis to the analysis of PM_{10} nonattainment designations presented in the Case Study 1: Accountability Assessment of PM_{10} Nonattainment Designations in the Western US of the main text, we conduct an analogous analysis but without pruning observations based on estimated propensity scores. That is, we estimate average causal effects of PM_{10} nonattainment designations on average annual ambient PM_{10} during 1999-2001 and on Medicare health outcomes using data on all 547 monitoring locations (268 of which are located in nonattainment areas).

Table B.1: Causal Effect Estimates for overall, associative, and dissociative effects in the sensitivity analysis of PM_{10} nonattainment designations carried out among all 547 monitoring locations, without pruning based on the estimated propensity score.

	Overall Average Causal Effect			Average Dissociative Effect			Average Associative Effect		
	Mean	2.5%	97.5%	Mean	2.5%	97.5%	Mean	2.5%	97.5%
Ambient PM10	-3.27	-15.76	5.38						
All-Cause Mortality	-5.26	-7.94	-2.24	-4.22	-8.5	0.02	-5.85	-9.91	-1.81
CVD Hospitalization	4.61	-1.7	10.06	5.48	-3.05	14.25	-0.92	-8.8	6.11
Respiratory Hospitalization	-0.63	-4.43	2.58	0.58	-3.3	4.29	-3.19	-8.68	1.45

Table B.1 presents point estimates and 95% posterior intervals of the overall average causal effects of the nonattainment designations, as well as estimates of average dissociative and associative effects. The estimates of average causal effects on Medicare health outcomes (overall ATT, average dissociative effects, and average associative effects) are depicted in Figure B.1, which is analogous to Figure 9 of the main text. In comparison with the results in the main text, the analysis without propensity score pruning estimates more pronounced effects on mortality and CVD-related hospitalizations, and similar results for respiratory-related hospitalizations. However, the results of the analysis without propensity score pruning necessarily extrapolate inferences beyond the observed data, as the analysis entails inferences for nonattainment locations that have no comparable attainment location in the data.

Figure B.1: Posterior mean point estimates and 95% posterior probability intervals for overall, associative, and dissociative effects in the sensitivity analysis of PM_{10} nonattainment designations carried out among all 547 monitoring locations, without pruning based on the estimated propensity score.

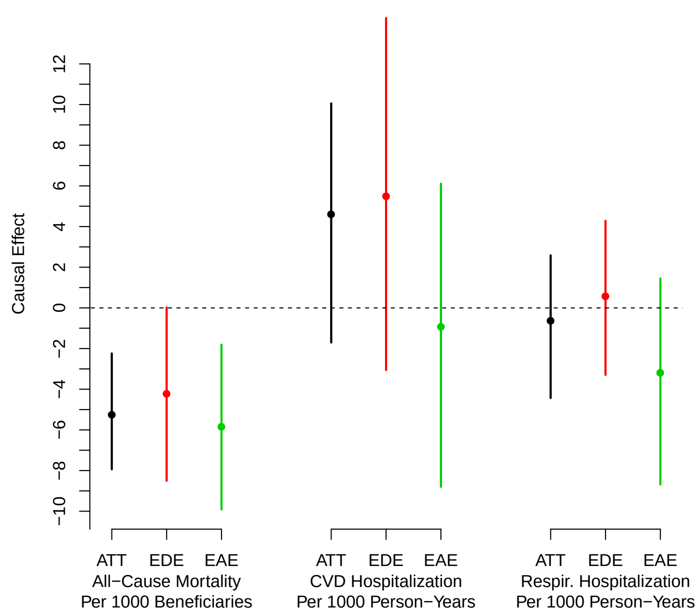

Appendix C. Technical Details for Case Study 2: Power Plant Emissions Controls

C.1 Defining the Treatment: Scrubber Installations on Power Plants

Figure C.1 shows the distribution of the % of each power plant's total heat input among EGUs with a scrubber. Note that the vast majority of power plants have scrubbers on all or none of their EGUs.

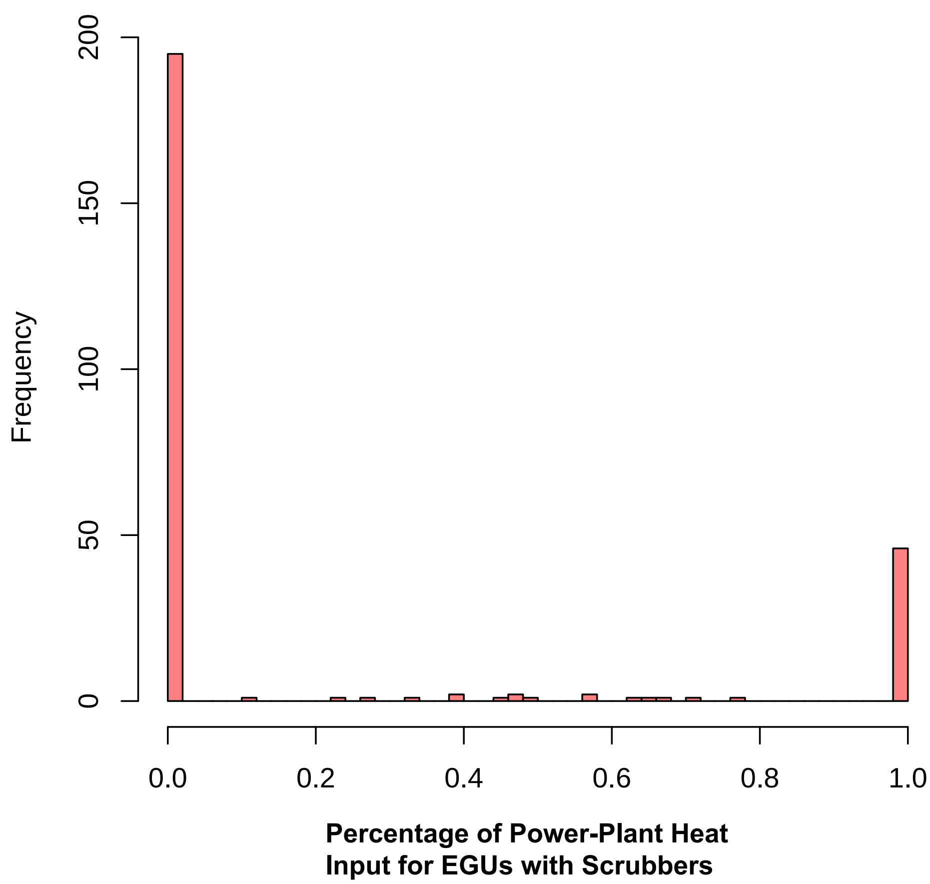

Figure C.1: Histogram of the % of heat input generated from an EGU with a scrubber for each power-generating facility.

C.2 Bayesian Nonparametric Models for Observed Distributions of Mediators and Outcome

We try to minimize parametric assumptions in our specification of the models for the observed data. In particular, we specify Dirichlet process mixtures of multivariate normals (Escobar and West, 1995; Mller et al., 1996; Jara et al., 2011) for the distribution of each mediator (emission). For each intervention $z = 0,1$ and baseline covariates $\boldsymbol{X} = \boldsymbol{x}$, the conditional distribution of the k-th observed mediator is specified as

$$
\begin{aligned}
M_{k,i}|Z=z,\boldsymbol{X}=\boldsymbol{x} &\sim N(\beta^z_{k0,i}+\boldsymbol{x}^\top\boldsymbol{\beta}^z_{k1},\ \tau^z_{k,i}), \qquad k=1,2,3; i=1,\cdots,n^z\\
\beta^z_{k0,i},\tau^z_{k,i} &\sim F^z_k,\\
F^z_k &\sim DP(\lambda^z_k,\ \mathcal{F}^z_k),
\end{aligned}
$$

where subscript k indicates the k-th mediator and i indicates the i-th observation and superscript z indicates the intervention received such that $\beta^z_{k0,i}$ and $\tau^z_{k,i}$ denote the intercept and precision parameters for the i-th observation of the k-th mediator, respectively. Here, DP denotes the Dirichlet process with two parameters, a mass parameter (λ^z_k) and a base measure ($\mathcal{F}^z_k$) for each mediator k and intervention z. To not overly complicate the model we only 'mixed' over the intercept parameters and precisions in the conditional distributions, $\beta^z_{k0,i}$ and $\tau^z_{k,i}$. The base distribution $\mathcal{F}^z_k$ is taken to be the conjugate normal-Gamma distribution, $N(\mu^z_k,S^z_k)G(a^z_k,b^z_k)$, where S^z_k is the precision parameters and the Gamma is parametrized as the mean to be a^z_k/b^z_k and we set a Gamma prior $G(1,1)$ on λ^z_k, the inverse of the mass parameter λ^z_k (Rasmussen, 1999). For the hyper-priors, we follow the specification from Taddy (2008) such that $\mu^z_k \sim (\mu^{z\star}_k,S^{z\star}_k), S^z_k \sim G(a^{z\star}_k,b^{z\star}_k)$ and $a^z_k \sim \text{Unif}(0.1,10), b^z_k = a^z_k \times \hat{\Sigma}^z_k/2$ where $\hat{\Sigma}^z_k$ is the MLE of the variance of $M_k(z)$. $S^{z\star}_k$ is set to $2/\hat{\Sigma}^z_k$ and $\mu^{z\star}_k$ is set to the mean of the data. And $a^{z\star}_k \sim \text{Unif}(0.1,10), b^{z\star}_k = a^z_k \times \hat{\Sigma}^z_k/2$. From these specifications, $E(\tau^z_{k,i}) = E(S^z_k) = E(S^{z\star}_k) = \hat{\Sigma}^z_k/2$ (i.e., the expected variance components are an attenuated value of the MLE of the variance of the data) in order to avoid fitting only one global distribution over the whole data points.

Similarly, we can specify Dirichlet process mixtures of multivariate normals for the distribution of each outcome. For each intervention $z = 0,1$, the conditional distribution of the observed outcome (ambient $PM_{2.5}$) is specified as

$$\begin{aligned} Y_i|Z=z,\boldsymbol{X}=\boldsymbol{x} &\sim N(\gamma_{0,i}^z + \boldsymbol{x}^\top \boldsymbol{\gamma}_1^z,\ \omega_i^z), \qquad i=1,\cdots,n^z \\ \gamma_{0,i}^z,\omega_i^z &\sim G^z, \\ G^z &\sim DP(\eta^z,\ \mathcal{G}^z), \end{aligned}$$

where all details follow in the same manner as the specification of the mediator models.

These observation models can be represented as the stick-breaking construction (Sethuraman, 1994) which can be approximated by a finite mixture of normals such that, for example, the conditional distribution of M_1 under intervention $z = 1$ can be represented as

$$f_{M_1}(m|z=1,\boldsymbol{x}) = \sum_{n=1}^{N} \theta_n N(m;\, \beta_{10,n}^{z=1} + \boldsymbol{x}^\top \boldsymbol{\beta}_{11}^{z=1}, \tau_{1,n}^{z=1}),$$

where $\theta_n = \theta'_n \prod_{h<n}(1-\theta'_h), \theta'_h \sim \text{Beta}(1,\lambda_1)$, and $(\beta_{10,n}^{z=1}, \tau_{1,n}^{z=1}) \overset{iid}{\sim} \mathcal{F}_1^{z=1}$ and N is a maximum number of clusters.

With the marginal distributions specified as above, we can specify the joint distribution of the outcome and the mediators under the same intervention via a Gaussian copula model (Nelsen, 1999) and each conditional distribution of the mediator and the outcome above. Specifically, it has the following form for each intervention $Z = 0,1$:

$$\begin{aligned} &F_{M_1,M_2,M_3,Y}(m_1,m_2,m_3,y\,|\,Z=z,\boldsymbol{X}=\boldsymbol{x}) = \\ &\quad \Phi_4[\Phi_1^{-1}\{F_{M_1}(m_1|Z=z,\boldsymbol{x})\}, \Phi_1^{-1}\{F_{M_2}(m_2|Z=z,\boldsymbol{x})\}, \Phi_1^{-1}\{F_{M_3}(m_3|Z=z,\boldsymbol{x})\}, \Phi_1^{-1}\{F_Y(y|Z=z,\boldsymbol{x})\}], \end{aligned}$$

where Φ_1 is the univariate standard normal CDF and Φ_4 is the multivariate normal CDF with mean $\mathbf{0}$ and a correlation matrix R. Since the outcome and the mediators under the same intervention are observed by the data, the correlation matrix R can be identified by the observed data.

C.3 Assumptions for Identification of Principal Causal Effects and Mediation Effects

The models in Section C.2 are for the mediator and outcome values that are actually observed. Identification of causal effects requires estimation of both observed and unobserved potential outcomes for each power plant. This requires assumptions related to the unobserved but *observable* outcomes for the principal causal effects, and additional assumptions related to unobserved and *unobservable* outcomes for the mediation effects.

The conditional distribution $[Y(z;\mathbf{M}(z_1,z_2,z_3))\,|\,\mathbf{M}(0,0,0)=\mathbf{m}_{0,0,0},\mathbf{M}(1,1,1)=\mathbf{m}_{1,1,1},\boldsymbol{X}=\boldsymbol{x}]$ is denoted by $f_{z,\mathbf{M}(z_1,z_2,z_3)}(y\,|\,\mathbf{m}_{0,0,0},\mathbf{m}_{1,1,1},\boldsymbol{x})$ where $\mathbf{m}_{z_1,z_2,z_3}$ is a vector of realized values of SO_2, NO_x and CO_2 under the interventions z_1,z_2,z_3. The conditional distribution $[\mathbf{M}(z_1,z_2,z_3)|\boldsymbol{X}=\boldsymbol{x}]$ is denoted by $f_{\mathbf{M}(z_1,z_2,z_3)}(\mathbf{m}_{z_1,z_2,z_3}|\boldsymbol{x})$. Other conditional distributions are defined using similar notation.

C.3.1 Assumptions for the Principal Causal Effects

Assumption 1. (Ignorability of treatment)

$$\{Y(z;\mathbf{M}(z,z,z)),\mathbf{M}(z,z,z)\}\perp\!\!\!\perp Z|\boldsymbol{X}=\boldsymbol{x},$$

for $z=0,1$, which assumes that SO_2 scrubber installation status is "randomized" conditional on $\boldsymbol{X}=\boldsymbol{x}$. With this assumption, we can identify the conditional distributions of potential $PM_{2.5}$ outcomes and the conditional distributions of emissions outcomes under each intervention based on observed data using the models from Section C.2 denoted with $f_{z,\mathbf{M}(z,z,z)}(y|\boldsymbol{x})$ and $f_{M_k(z)}(m|\boldsymbol{x})$.

Assumption 2. *The joint distribution of potential mediators and outcomes conditional on covariates follows a Gaussian copula model (Nelsen, 1999):*

$$\begin{aligned}&F_{\mathbf{M}(0,0,0),\mathbf{M}(1,1,1),(0,\mathbf{M}(0,0,0)),(1,\mathbf{M}(1,1,1))}(\mathbf{m}_{0,0,0},\mathbf{m}_{1,1,1},y_0,y_1)=\\&\quad\Phi_8[\Phi_1^{-1}\{F_{M_1(0)}(m_1)\},\Phi_1^{-1}\{F_{M_2(0)}(m_2)\},\Phi_1^{-1}\{F_{M_3(0)}(m_3)\},\Phi_1^{-1}\{F_{M_1(1)}(m_1)\},\\&\quad\Phi_1^{-1}\{F_{M_2(1)}(m_2)\},\Phi_1^{-1}\{F_{M_3(1)}(m_3)\},\Phi_1^{-1}\{F_{(0,\mathbf{M}(0,0,0))}(y_0)\},\Phi_1^{-1}\{F_{(1,\mathbf{M}(1,1,1))}(y_1)\}]\end{aligned}$$

where y_z indicates a value of potential outcome under intervention $Z = z$ and Φ_1 is the univariate standard normal CDF and Φ_8 is the multivariate normal CDF with mean $\mathbf{0}$ and a correlation matrix R.

Here, for notational simplicity, we omit covariates, $\boldsymbol{X}$, when noting conditional distributions of mediators and outcomes. In Section C.2, the joint distribution of the observed data (outcome and mediators under the same scrubber status) was specified as a separate Gaussian copula model for each intervention $Z = z$, and we are now assuming the joint distribution of potential outcomes and mediators under both interventions $Z = 0,1$. Through this model, we have a benefit of flexibility in the marginal distributions (which we modeled earlier using Bayesian nonparametric methods). Since any pair of $M_j(0)$ and $[M_k(1), Y(1;\mathbf{M}(1,1,1))]$ or any pair of $Y(0;\mathbf{M}(0,0,0))$ and $[M_k(1), Y(1;\mathbf{M}(1,1,1))]$ given covariates $\boldsymbol{X} = \boldsymbol{x}$ cannot be observed at the same time, we cannot identify part of the correlation structure of the joint distribution from the observed data. To identify these nonidentifiable pieces, we first adopt strategy from Zigler et al. (2012) to specification of the correlations between mediators under different interventions. Specifically, we set

(1) $cor(M_j(0), M_k(1)) = \frac{cor(M_j(0),M_k(0))+cor(M_j(1),M_k(1))}{2} \times \rho_1, \quad \text{for } j,k = 1,2,3,$

where ρ_1 is a sensitivity parameter. This strategy implies that (a) the correlation between the same mediator $(j = k)$ under opposite interventions is ρ_1, and (b) the correlation between different mediators $(j \neq k)$ under opposite interventions is attenuated version of the correlation observed separately under each intervention. Furthermore, it is reasonable to assume that magnitude of correlations (b) do not exceed the observed correlations between two mediators under the same intervention. Thus, we set a constraint $|\rho_1| \leq 2 \times \frac{\min\{|cor(M_j(0),M_k(0))|,|cor(M_j(1),M_k(1))|\}}{|cor(M_j(0),M_k(0))+cor(M_j(1),M_k(1))|}$ for all $j \neq k$ combinations. We use one sensitivity parameter to specify these correlations, but a different parameter could be specified for each mediator separately. For the analysis presented in the main text, we specify a uniform prior distribution over $(0, 0.75)$ for ρ_1.

We also need to specify correlations between $Y(0;\mathbf{M}(0,0,0))$ and $M_k(1)$ and between $Y(1;\mathbf{M}(1,1,1))$ and $M_k(0)$ for all $k = 1,2,3$. This can be done by the same strategy from the above specification for $k = 1,2,3$

(2) $cor(M_k(z'), Y(z;\mathbf{M}(z,z,z))) = \frac{cor(M_k(0),Y(0;\mathbf{M}(0,0,0)))+cor(M_k(1),Y(1;\mathbf{M}(1,1,1)))}{2} \times \rho_2,$

where $z' = 1 - z$ for $z \in \{0,1\}$ and ρ_2 is another sensitivity parameter. This specification implies that the correlation between mediators under intervention $z = 1$ ($z = 0$) and the outcome under intervention $z = 0$ ($z = 1$) is attenuated version of the correlations between the mediators and the outcome under the same interventions. We also restrict a range of possible values as $|\rho_2| \leq 2 \times \frac{\min\{|cor(M_k(0),Y(0;\mathbf{M}(0,0,0)))|,|cor(M_k(1),Y(1;\mathbf{M}(1,1,1)))|\}}{|cor(M_k(0),Y(0,\mathbf{M}(0,0,0)))+cor(M_k(1),Y(1,\mathbf{M}(1,1,1)))}$. For the analysis presented in the main text, we specify a uniform prior distribution over $(0,0.15)$ for ρ_2.

Note that we do not need to specify a third sensitivity parameter for the nonidentifiable association between $(Y(0;\mathbf{M}(0,0,0),Y(1;\mathbf{M}(1,1,1))$. Although this would be required to construct the full joint distribution of all observable potential outcomes, the estimands we consider do not depend on this correlation. This is due in part to the following assumption:

Assumption 3. (Conditional independence I) $Y(0;\mathbf{M}(0,0,0))$ *and* $Y(1;\mathbf{M}(1,1,1))$ *are conditionally independent given all potential mediators under* $z = 0$ *and* $z = 1$ *and covariates* $\boldsymbol{X}$.

This assumption states that the potential values of $PM_{2.5}$ under both interventions are independent of each other conditional on all potential values of SO_2, NO_x and CO_2 under both interventions and baseline covariates. This is not necessary to estimate the posterior means of the principal causal effects but necessary to estimate other features of the posterior distribution of the principal causal effects such as bounds for posterior variances. In this paper, the mean differences are our primary causal estimands and, therefore, this assumption is not needed.

C.3.2 Assumptions for the mediation effects

Assumptions 1-3 pertain exclusively to observed outcomes and unobserved but *observable* outcomes of ambient $PM_{2.5}$ and emissions. As noted in the main text, identification of natural indirect effects requires assumptions that relate observed outcomes to *unobservable* potential outcomes of $PM_{2.5}$ and emissions that are simultaneously subject to different scrubber statuses. We specify such assumptions here.

Assumption 3*. (Conditional Independence II) $\{Y(z;\mathbf{M}(z_1,z_2,z_3)) : (z,z_1,z_2,z_3) \in \{0,1\}^{\otimes 4}\}$ *are conditionally independent given all potential mediators and covariates* $\boldsymbol{X}$.

It is important to note that Assumption 3 is a specific case of Assumption 3*. Again, this assumption is not necessary to estimate posterior means of NDE, JNIEs and other indirect effects but necessary to estimate other features of the posterior distribution of mediation effects such as an upper bound on the posterior variance of the indirect effect assuming a non-negative correlation between potential outcomes.

Assumption 4. *For a given intervention $Z = 1$, the conditional distribution of the potential outcome given all potential mediators and covariates is the same whether corresponding mediators were induced by $Z = 1$ or $Z = 0$.*

This assumption implies that the unobservable potential outcomes $Y(1;\mathbf{M}(0,0,0))$ and the observed potential outcomes $Y(1;\mathbf{M}(1,1,1))$ have the same conditional distribution,

$$\begin{aligned} & f_{1,\mathbf{M}(0,0,0)}(y\,|\,\mathbf{M}(0,0,0)=\mathbf{m},\mathbf{M}(1,1,1),\boldsymbol{x}) \\ = \quad & f_{1,\mathbf{M}(1,1,1)}(y\,|\,\mathbf{M}(0,0,0),\mathbf{M}(1,1,1)=\mathbf{m},\boldsymbol{x}), \end{aligned} \tag{7}$$

where the conditional distribution of the RHS has **m** as a vector of realized values of when an SO_2 scrubber is installed. That is, the conditional distribution of the unobservable potential $PM_{2.5}$ concentration when a scrubber is installed but emissions are set to the value **m** that they would have been absent the scrubber is equal to the observed distribution of $PM_{2.5}$ observed around power plants that had scrubbers and were observed to have emissions **m**.

This assumption is also defined for cases of any two mediators or single pollutant(s) emitted under different interventions. For instance, the potential outcomes of $PM_{2.5}$ $Y(1;\mathbf{M}(0,1,0))$ and $Y(1;\mathbf{M}(1,1,1))$ have the same conditional distribution regardless of whether corresponding emissions values arose under a scrubber ($Z = 1$) or absent a scrubber ($Z = 0$),

$$\begin{aligned} & f_{1,\mathbf{M}(0,1,0)}(y\,|\,\mathbf{M}(0,1,0)=\mathbf{m},\mathbf{M}(1,0,1),\boldsymbol{x}) \\ = \quad & f_{1,\mathbf{M}(1,1,1)}(y\,|\,\mathbf{M}(0,0,0),\mathbf{M}(1,1,1)=\mathbf{m},\boldsymbol{x}). \end{aligned}$$

Effects	Assumptions	Note
Principal Causal Effect	A1, A2, A3	
Mediation Effect	A1, A2, A3$^\bigstar$, A4	A3$^\bigstar$ implies A3

Table C.1: Assumptions needed for identifying each effect. 'A' indicates assumption.

Similarly, the potential outcomes $Y(1;\mathbf{M}(1,1,0))$ and $Y(1;\mathbf{M}(1,1,1))$ have the same conditional distribution,

$$\begin{aligned} & f_{1,\mathbf{M}(1,1,0)}(y\,|\,\mathbf{M}(1,1,0)=\mathbf{m},\mathbf{M}(0,0,1),\boldsymbol{x}) \\ = \quad & f_{1,\mathbf{M}(1,1,1)}(y\,|\,\mathbf{M}(0,0,0),\mathbf{M}(1,1,1)=\mathbf{m},\boldsymbol{x}). \end{aligned}$$

The key point is that unobservable $PM_{2.5}$ concentrations under a certain hypothetical emissions amount (**m**) are assumed to have the same distribution as that of the $PM_{2.5}$ concentrations among power plants that had, in reality, emissions of **m** observed, regardless of observed scrubber status.

It is worth noting that this Assumption 4 induces the following property

$$Y(1;\mathbf{M}(1,1,1)) \perp\!\!\!\perp \mathbf{M}(0,0,0)|\mathbf{M}(1,1,1)=\mathbf{m},\boldsymbol{X}=\boldsymbol{x} \tag{8}$$

for all **m** and $\boldsymbol{x}$ values since Equality (7) holds regardless of $\mathbf{M}(0,0,0)$ in the conditioning part of the RHS for all realized values of $\mathbf{M}(1,1,1)$ and $\boldsymbol{X}$. This property simplifies posterior computation.

Also, note that Assumption 4 is consistent with Assumption 2 (the joint distribution of outcomes and mediators) since Assumption 2 only impacts on the conditional distribution of $Y(1;\mathbf{M}(1,1,1))$ which is the RHS of Assumption 4.

We summarize assumptions proposed in the previous section and this section with Table C.1. Assumption 1-3 are sufficient to identify the principal causal effects. To identify the mediation effects, we replace Assumption 3 with stronger assumption, Assumption 3$^\bigstar$ and propose an additional assumption, Assumption 4. It is again worth noting that Assumption 1-3 only contain the observable outcomes while Assumption 3$^\bigstar$-4 additionally contain the unobservable outcomes. Thus, we modularize assumptions so that estimation of principal causal effects rely on Assumptions 1, 2, and 3, while estimation of mediation effects relies on

Assumptions 1, 2, $3^\bigstar$, and 4.

C.3.3 Proofs of Identification

In the following, we will prove that Assumptions 1-3 are sufficient to identify the distribution of the principal causal effects and Assumption 1,2,$3^\bigstar$ and 4 are sufficient to identify the distributions of NDE, JNIEs and NIEs. Here, we proceed with 3 mediators, SO_2, NO_x, and CO_2, but it is straightforward to extend this to $K \geq 3$ cases.

Theorem 1. *The posterior distributions of the principal causal effects (associative effect and dissociative effect) are identified under Assumption 1-4.*

Proof :

To obtain the posterior distribution of the principal causal effect, it is sufficient to identify the conditional joint distribution of the potential outcomes and mediators $[Y(0;\mathbf{M}(0,0,0)) = y_0, Y(1;\mathbf{M}(1,1,1)) = y_1, \mathbf{M}(0,0,0) = \mathbf{m}_{0,0,0}, \mathbf{M}(1,1,1) = \mathbf{m}_{1,1,1}]$,

$$
\begin{aligned}
& \mathrm{f}(y_0, y_1, \mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}) && (9)\\
= & \int f_{(0,\mathbf{M}(0,0,0)),(1,\mathbf{M}(1,1,1))}(y_0, y_1 | \mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}, \boldsymbol{x}) f_{\mathbf{M}(0,0,0),\mathbf{M}(1,1,1)}(\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1} | \boldsymbol{x}) dF_{\boldsymbol{X}}(\boldsymbol{x}) \\
= & \int \left\{ f_{0,\mathbf{M}(0,0,0)}(y_0 | \mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}, \boldsymbol{x}) f_{1,\mathbf{M}(1,1,1)}(y_1 | \mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}, \boldsymbol{x}) \right\} f_{\mathbf{M}(0,0,0),\mathbf{M}(1,1,1)}(\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1} | \boldsymbol{x}) dF_{\boldsymbol{X}}(\boldsymbol{x})
\end{aligned}
$$

where the equality in (9) from Assumption 3 all terms are identified by Assumption 1 and 2 with the observed data model. Then, the posterior distribution of the principal causal effect is a function of (9).

Theorem 2. *The posterior distributions of* NDE, JNIE_{123}*,* JNIE_{jk}*, and* NIE_k *are identified under Assumptions 1, 2,* $3^\bigstar$ *and 4.*

Proof :

To obtain the posterior distributions of NDE, JNIE's and NIE's, it is sufficient to identify the joint distribution of the potential outcomes $[Y(1;\mathbf{M}(1,1,1)) = y_1, Y(0;\mathbf{M}(0,0,0)) = y_2, Y(1;\mathbf{M}(0,0,0)) = y_3, Y(1;\mathbf{M}(1,0,0)) =$

$y_4, Y(1;\mathbf{M}(0,1,0)) = y_5, Y(1;\mathbf{M}(0,0,1)) = y_6, Y(1;\mathbf{M}(1,1,0)) = y_7, Y(1;\mathbf{M}(1,0,1)) = y_8, Y(1;\mathbf{M}(0,1,1)) = y_9]$,

$$\begin{aligned}
&f(y_1,y_2,y_3,y_4,y_5,y_6,y_7,y_8,y_9) = \\
&\quad \int \Big\{ f(y_1,y_2,y_3,y_4,y_5,y_6,y_7,y_8,y_9 \,|\, \mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}, \boldsymbol{x}) \\
&\qquad \times f_{\mathbf{M}(0,0,0),\mathbf{M}(1,1,1)}(\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1} \,|\, \boldsymbol{x}) \Big\}\, d\mathbf{m}_{0,0,0}\, d\mathbf{m}_{1,1,1}\, dF_{\boldsymbol{X}}(x),
\end{aligned}$$

where $\mathbf{m}_{z_1,z_2,z_3}$ denotes a vector of realized values of the mediators $\{M(z_1), M(z_2), M(z_3)\}$. With omitting $\boldsymbol{x}$ for notation simplicity, the second term in the RHS can be identified by Assumption 2. The first term in the RHS can be factored as

$$\begin{aligned}
&f(y_1,y_2,y_3,y_4,y_5,y_6,y_7,y_8,y_9|\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}) \\
&= f_{1,\mathbf{M}(1,1,1)}(y_1|\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1})\, f_{0,\mathbf{M}(0,0,0)}(y_2|\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}) \\
&\quad \times f_{1,\mathbf{M}(0,0,0)}(y_3|\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1})\, f_{1,\mathbf{M}(1,0,0)}(y_4|\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}) \\
&\quad \times f_{1,\mathbf{M}(0,1,0)}(y_5|\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1})\, f_{1,\mathbf{M}(0,0,1)}(y_6|\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}) \\
&\quad \times f_{1,\mathbf{M}(1,1,0)}(y_7|\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1})\, f_{1,\mathbf{M}(1,0,1)}(y_8|\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}) \\
&\quad \times f_{1,\mathbf{M}(0,1,1)}(y_9|\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}) \qquad (10) \\
&= f_{1,\mathbf{M}(1,1,1)}(y_1|\mathbf{M}(1,1,1) = \mathbf{m}_{1,1,1})\, f_{0,\mathbf{M}(0,0,0)}(y_2|\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}) \\
&\quad \times f_{1,\mathbf{M}(1,1,1)}(y_3|\mathbf{M}(1,1,1) = \mathbf{m}_{0,0,0})\, f_{1,\mathbf{M}(1,1,1)}(y_4|\mathbf{M}(1,1,1) = \mathbf{m}_{1,0,0}) \\
&\quad \times f_{1,\mathbf{M}(1,1,1)}(y_5|\mathbf{M}(1,1,1) = \mathbf{m}_{0,1,0})\, f_{1,\mathbf{M}(1,1,1)}(y_6|\mathbf{M}(1,1,1) = \mathbf{m}_{0,0,1}) \\
&\quad \times f_{1,\mathbf{M}(1,1,1)}(y_7|\mathbf{M}(1,1,1) = \mathbf{m}_{1,1,0})\, f_{1,\mathbf{M}(1,1,1)}(y_8|\mathbf{M}(1,1,1) = \mathbf{m}_{1,0,1}) \\
&\quad \times f_{1,\mathbf{M}(1,1,1)}(y_9|\mathbf{M}(1,1,1) = \mathbf{m}_{0,1,1}) \qquad (11)
\end{aligned}$$

where the first equality (10) follows from Assumption $3^\star$ and the second equality (11) follows from Assumption 4. All terms in (11) are identified by Assumption 2. Note that, in (11), all conditional distributions for potential outcomes except one for $Y(0, \mathbf{M}(0,0,0))$ are independent of $\mathbf{M}(0,0,0)$ which simplify the posterior computation and result in efficient estimates. Then, the posterior distributions of NDE, JNIEs and NIEs are functions of (11).

C.4 Posterior Computation

In this section, we provide the MCMC computational algorithm for estimating principal causal effects and the mediation effects. To compute the posterior distribution of the principal causal effects, for example the associative effect defined on all three pollutants (SO_2, NO_x, CO_2), we obtain N posterior samples of the parameters in the observation models via Stan (Guo et al., 2015). Then, for each set of sampled parameters with sensitivity parameters ρ_1 and ρ_2, we do the following steps.

1. Generate S samples of $[y_0, y_1, \mathbf{m}_0, \mathbf{m}_1]$ from the distribution of $[Y(0,\mathbf{M}(0,0,0)), Y(1,\mathbf{M}(1,1,1)), \mathbf{M}(0,0,0), \mathbf{M}(1,1,1)|\boldsymbol{X} = \boldsymbol{x}]$ where y_0, y_1, $\mathbf{m}_{0,0,0}$ and $\mathbf{m}_{1,1,1}$ are realized samples of $Y(0,\mathbf{M}(0,0,0))$, $Y(1,\mathbf{M}(1,1,1))$, $\mathbf{M}(0,0,0)$ and $\mathbf{M}(1,1,1)$, respectively.

2. Compute $E[Y(1;\mathbf{M}(1,1,1)) \,\big|\, |\mathbf{M}(1,1,1) - \mathbf{M}(0,0,0)|_{\mathcal{K}} > C^A_{\mathcal{K}}, \boldsymbol{X} = \boldsymbol{x}]$ as follows,

$$E[Y(1;\mathbf{M}(1,1,1)) \,\big|\, |\mathbf{M}(1,1,1) - \mathbf{M}(0,0,0)|_{\mathcal{K}} > C^A_{\mathcal{K}}, \boldsymbol{X} = \boldsymbol{x}] = \frac{1}{S}\sum_{i=1}^{S} y_{1i} I(|\mathbf{m}_{1,1,1,i} - \mathbf{m}_{0,0,0,i}|_{\mathcal{K}} > C^A_{\mathcal{K}}),$$

   where $I()$ is an indicator function and y_{1i}, $\mathbf{m}_{0,0,0,i}$, and $\mathbf{m}_{1,1,1,i}$ indicate the i-th sample of the outcome under intervention $z = 1$ and sets of the mediators under intervention $z = 0$ and intervention $z = 1$, respectively. Analogously, compute $E[Y(0;\mathbf{M}(0,0,0)) \,\big|\, |\mathbf{M}(1,1,1) - \mathbf{M}(0,0,0)|_{\mathcal{K}} > C^A_{\mathcal{K}}, \boldsymbol{X} = \boldsymbol{x}]$.

3. Compute the associative effect,

$$AE = \int E[Y(1;\mathbf{M}(1,1,1)) - Y(0;\mathbf{M}(0,0,0)) \,\big|\, |(\mathbf{M}(1,1,1) - \mathbf{M}(0,0,0))|_{\mathcal{K}} > C^A_{\mathcal{K}}, \boldsymbol{x}] dF_{\boldsymbol{X}}(\boldsymbol{x})$$

Then, we iterate steps 1-3 N times and estimate posterior means and standard deviations of the associative effects. The dissociative effect is computed in the exact same way with some threshold $C^D_{\mathcal{K}}$.

To compute the posterior distribution of the JNIE's, NDE and all mediator-specific indirect effects, using the same N sets of parameters in the observation models from the above step, for each set of sampled parameters and either fixed values for sensitivity parameters ρ_1 and ρ_2, we do the following steps.

1. Generate S sets of samples $(\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1})$ from the joint distribution $[\mathbf{M}(0,0,0), \mathbf{M}(1,1,1)|\boldsymbol{X}=\boldsymbol{x}]$ where $\mathbf{m}_{0,0,0}$ and $\mathbf{m}_{1,1,1}$ are realized samples of $\mathbf{M}(0,0,0)$ and $\mathbf{M}(1,1,1)$, respectively.

2. Compute $f_{1,\mathbf{M}(0,0,0)}(y|\boldsymbol{x})$ via Monte Carlo integration using S sets of $(\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1})$ as follows,

$$\begin{aligned}
& f_{1,\mathbf{M}(0,0,0)}(y|\boldsymbol{x}) \\
= & \int \left\{ f_{1,\mathbf{M}(0,0,0)}(y|\mathbf{M}(0,0,0)=\mathbf{m}_{0,0,0}, \mathbf{M}(1,1,1)=\mathbf{m}_{1,1,1}, \boldsymbol{x}) f_{\mathbf{M}(0,0,0),\mathbf{M}(1,1,1)|\boldsymbol{X}=\boldsymbol{x}}(\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}) \right\} d\mathbf{m}_{0,0,0}\, d\mathbf{m}_{1,1,1} \\
= & \int \left\{ f_{1,\mathbf{M}(1,1,1)}(y|\mathbf{M}(1,1,1)=\mathbf{m}_{0,0,0}, \boldsymbol{x}) f_{\mathbf{M}(0,0,0),\mathbf{M}(1,1,1)|\boldsymbol{X}=\boldsymbol{x}}(\mathbf{m}_{0,0,0}, \mathbf{m}_{1,1,1}) \right\} d\mathbf{m}_{0,0,0}\, d\mathbf{m}_{1,1,1} \quad \text{(A4)} \\
\approx & \frac{1}{S}\sum_{i=1}^{S} f_{1,\mathbf{M}(1,1,1)}(y|\mathbf{M}(1,1,1)=\mathbf{m}_{0,0,0,i}, \boldsymbol{x}),
\end{aligned}$$

where $\mathbf{m}_{0,0,0,i}$ indicates the i-th sample of mediators, $\mathbf{M}(0,0,0)$, and 'A4' denotes Assumption 4. To compute $f_{1,\mathbf{M}(1,1,1)}(y|\mathbf{M}(1,1,1)=\mathbf{m}_{0,0,0}, \boldsymbol{x})$, we note that this conditional distribution is equivalent to

$$\frac{f(Y(1;\mathbf{M}(1,1,1))=y, \mathbf{M}(1,1,1)=\mathbf{m}_{0,0,0}, \boldsymbol{x})}{f(\mathbf{M}(1,1,1)=\mathbf{m}_{0,0,0}, \boldsymbol{x})},$$

where all the terms are identifiable by Assumption 2 and the model specifications in Section C.2. All other distributions of unobservable outcomes such as $f_{1,\mathbf{M}(1,1,0)}(y|\boldsymbol{x})$, $f_{1,\mathbf{M}(1,0,1)}(y|\boldsymbol{x})$, $f_{1,\mathbf{M}(0,1,1)}(y|\boldsymbol{x})$, $f_{1,\mathbf{M}(1,0,0)}(y|\boldsymbol{x})$, $f_{1,\mathbf{M}(0,1,0)}(y|\boldsymbol{x})$ and $f_{1,\mathbf{M}(0,0,1)}(y|\boldsymbol{x})$ can be computed in the same manner.

3. Compute NDE, JNIEs, NIEs,

$$\text{NDE} = E[Y(1,\mathbf{M}(1,1,1)) - Y(0,\mathbf{M}(0,0,0))] = \int y \left\{ f_{1,\mathbf{M}(0,0,0)}(y|\boldsymbol{x}) - f_{0,\mathbf{M}(0,0,0)}(y|\boldsymbol{x}) \right\} dy\, dF_{\boldsymbol{X}}(\boldsymbol{x}),$$

$$\text{JNIE}_{123} = E[Y(1,\mathbf{M}(1,1,1)) - Y(1,\mathbf{M}(0,0,0))] = \int y \left\{ f_{1,\mathbf{M}(1,1,1)}(y|\boldsymbol{x}) - f_{1,\mathbf{M}(0,0,0)}(y|\boldsymbol{x}) \right\} dy\, dF_{\boldsymbol{X}}(\boldsymbol{x}),$$

$$\text{JNIE}_{12} = E[Y(1,\mathbf{M}(1,1,1)) - Y(1,\mathbf{M}(0,0,1))] = \int y \left\{ f_{1,\mathbf{M}(1,1,1)}(y|\boldsymbol{x}) - f_{1,\mathbf{M}(0,0,1)}(y|\boldsymbol{x}) \right\} dy\, dF_{\boldsymbol{X}}(\boldsymbol{x}),$$

and

$$\text{NIE}_1 = E[Y(1,\mathbf{M}(1,1,1)) - Y(1,\mathbf{M}(0,1,1))] = \int y \left\{ f_{1,\mathbf{M}(1,1,1)}(y|\boldsymbol{x}) - f_{1,\mathbf{M}(0,1,1)}(y|\boldsymbol{x}) \right\} dy\, dF_{\boldsymbol{X}}(\boldsymbol{x}).$$

Other mediator specific effects, $\text{NIE}_2, \text{NIE}_3$, and joint natural indirect effects of other pairs of mediators, $\text{JNIE}_{13}, \text{JNIE}_{23}$, are computed similarly.

We iterate steps 1-3 N times and estimate posterior means and standard deviations of the effects.

C.5 Assessment of Overlap of Natural Indirect Effects

A key benefit of our proposed methods relative to other approaches for multiple mediators is that we allow possible overlapping between NIEs such that one can examine whether there is actually overlap between NIEs or additivity of NIEs holds. We evaluate the relationship between the joint effects JNIE_{jk} and the mediator-specific effects NIE_1, NIE_2, NIE_3 through

$$\begin{aligned}
(\text{NIE}_1 + \text{NIE}_2) - \text{JNIE}_{12} &= 0.000(-0.002, 0.003), \\
(\text{NIE}_1 + \text{NIE}_3) - \text{JNIE}_{13} &= 0.000(-0.001, -0.003), \\
(\text{NIE}_2 + \text{NIE}_3) - \text{JNIE}_{23} &= 0.000(-0.003, 0.001)
\end{aligned}$$

which give no evidence of overlap between NIEs. That is, it appears as though the causal effect of an SO_2 scrubber on ambient $PM_{2.5}$ that is mediated through a given emission does not depend on the value of other emissions, i.e., there is no evidence of synergistic effects of joint reductions in multiple emissions.

Appendix D. Results from the Power-Plant Case Study with 75-km Data Linkage

As a sensitivity analysis to the analysis of the 150km radius used to link power plants to ambient monitors in the analysis of the Case Study 2: Accountability Assessment of Power Plant Emissions Controls Section of the main text, we conduct the same analysis but with power plants linked to ambient monitors within a 75km radius. Table D.1 shows the characteristics of the 53 power plants with scrubbers and 181 power plants without scrubbers in this analysis. Table D.2 presents posterior estimates of causal effects of scrubbers on emissions. Table D.3 presents posterior mean (and 95% interval) estimates of principal causal effects, and Figure D.1 depicts point estimates of principal effects analogous to Figure 13 of the main text. Table D.4 contains posterior estimates of natural direct and indirect effects, and Figure D.2 depicts boxplots of the posterior distributions of these mediation effects. Overall, the results of this sensitivity analysis are very similar to those of the main analysis, but exhibit wider uncertainty intervals (due to the inclusion of fewer power plants).

Table D.1: Summary statistics for covariates and outcomes available for the analysis of SO_2 controls when power plants are linked to monitors within a 75km radius.

	Plants with scrubbers (n=53)		Plants w/o scrubbers (n=181)	
	Mean	SD	Mean	SD
Monitor Data				
Average Ambient $PM_{2.5}$ 2005	12.10	4.00	13.30	2.50
Average Temperature 2004	12.80	4.50	13.10	3.70
Average Barometric Pressure 2004	725.80	47.90	744.70	20.60
Power Plant Level Data				
Total SO_2 Emission 2005	1390.00	1922.90	2173.80	2542.40
Total NO_x Emission 2005	936.30	796.90	588.10	558.20
Total CO_2 Emission 2005	560497.50	469587.80	369568.10	369898.60
Unit Level Data				
Selective Catalytic or Selective Non Catalytic Reduction Jan. 2004	0.30	0.40	0.20	0.40
Number of NO_x Controls Jan. 2004	1.20	0.60	1.00	0.60
Average Heat Input 2004	4578925.40	3602538.60	3605378.20	
Average Percent Operating Capacity	20.50	11.60	17.90	10.30
Phase II Indicator 2004	0.80	0.40	0.80	0.40
Total Operating Time 2004	7636.90	715.30	7347.10	1042.50
Sulfur Content in Coal 2004	1.50	1.10	0.80	0.60

Table D.2: Posterior mean and 95% probability intervals for causal effects of SO_2 controls on emissions when power plants are linked to monitors within a 75km radius.

	SO_2	NO_x	CO_2
Mean	-0.925	0.118	0.066
95% C.I.	(-1.289, -0.572)	(-0.151, 0.418)	(-0.152, 0.256)

Table D.3: Posterior mean and 95% probability intervals for expected associative and dissociative effects of SO_2 controls when power plants are linked to monitors within a 75km radius.

		SO_2	NO_x	CO_2	SO_2 & NO_x	SO_2 & CO_2	NO_x & CO_2	SO_2 & NO_x & CO_2
EAE_1	Mean	-0.518	-0.180	-0.373	-0.224	-0.438	-0.200	-0.231
	95% P.I.	(-1.683, 0.662)	(-1.447, 1.018)	(-1.432, 0.818)	(-1.524, 0.893)	(-1.633, 0.753)	(-1.385, 0.921)	(-1.507, 0.932)
EDE	Mean	-0.386	-0.522	-0.420	-0.407	-0.352	-0.522	-0.401
	95% P.I.	(-1.572, 0.990)	(-1.676, 0.650)	(-1.564, 0.930)	(-1.663, 1.043)	(-1.515, 1.089)	(-1.739, 0.667)	(-1.660, 1.077)
EAE_2	Mean	-0.075	-0.502	-0.499	-0.107	-0.191	-0.589	-0.236
	95% P.I.	(-1.522, 1.730)	(-1.891, 1.051)	(-1.690, 0.892)	(-1.649, 1.903)	(-1.595, 1.612)	(-1.967, 1.018)	(-1.686,1.659)

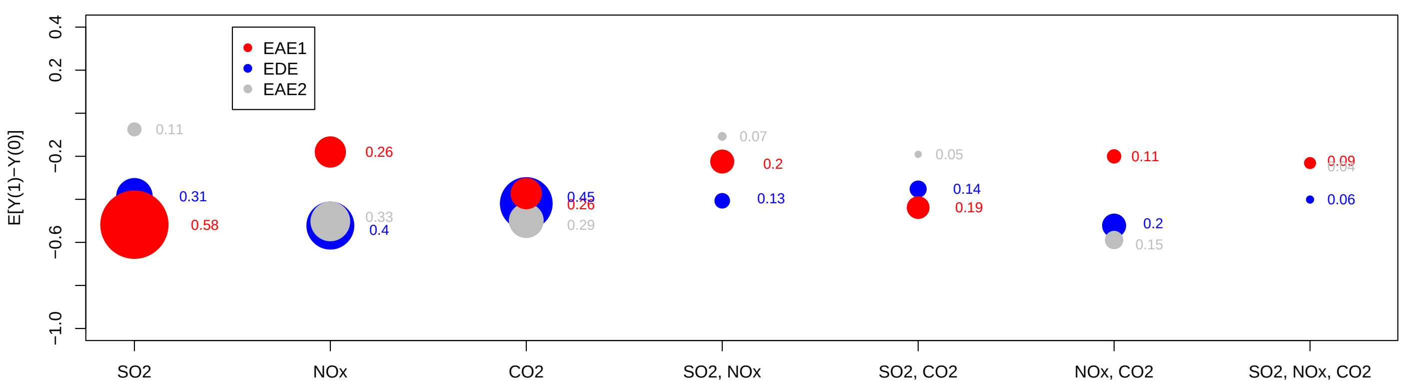

Figure D.1: Posterior mean estimates of average associative and dissociative effects of SO2 controls when power plants are linked to monitors within a 75km radius. Size of circle is proportional to the percent of observations estimated to belong in the corresponding strata, and number listed is posterior mean proportion.

Table D.4: Posterior mean and 95% probability intervals for total, direct, and indirect effects of SO_2 scrubbers when power plants are linked to monitors within a 75km radius.

	TE	$JNIE_{123}$	NDE
Mean	-0.413	-0.264	-0.149
95% C.I.	(-1.428, 0.696)	(-0.387, -0.139)	(-1.153, 0.981)
	NIE_1	NIE_2	NIE_3
Mean	-0.228	-0.021	-0.015
95% C.I.	(-0.331, -0.136)	(-0.069, 0.023)	(-0.073, 0.040)
	$JNIE_{12}$	$JNIE_{23}$	$JNIE_{13}$
Mean	-0.249	-0.035	-0.243
95% C.I.	(-0.358, -0.140)	(-0.117, 0.036)	(-0.365, -0.133)

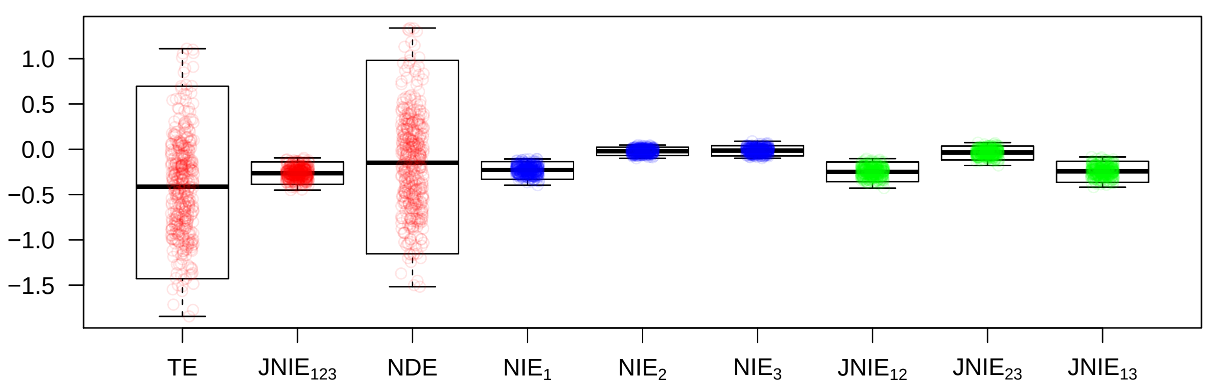

Figure D.2: Posterior distributions of direct and indirect effects in the analysis of SO_2 controls when power plants are linked to monitors within a 75 km radius.

Appendix E. Preliminary Extension of the Power-Plant Case Study to Health Outcomes

E.1 Power-Plant Centered Data Linkage

The first data linkage strategy we employ for a health outcomes analysis, which we call "power-plant centered linkage," is a straightforward analog to data linkage in the analysis in the Case Study 2: Accountability Assessment of Power Plant Emissions Controls Section of the main text to include Medicare health outcomes. Considering each coal-fired power plant as the unit of analysis, we link ambient monitoring locations and Medicare health outcomes as follows. For a given coal-fired power plant, we draw a circle with a 150 km radius and find all $PM_{2.5}$ monitoring stations located within the circle.

Here, if a monitoring station belongs to more than one 150km radius circle, we uniquely assign it to the closest power plant. Next, for these monitoring stations located within the circle, we draw 6 mile radius circle around each monitoring station and find all zip code centroids contained in these circles. Zip codes with centroids that are within 6 miles of more than one monitor are assigned to the closest monitoring location. Medicare outcomes among these zip codes are aggregated to represent health outcomes among Medicare beneficiaries residing within 6 miles of a monitoring location that is at most 150 km from a power plant. This power-plant centered linkage is illustrated in Figure E.1. In this linkage, we assume a coal-fired power plant is treated if all its EGUs have installed SO_2 scrubbers and otherwise consider it untreated. Also note that, for the health-outcomes analysis, we include data on possible confounders from the U.S. Census. These data are available at the zip code level, and are linked to monitoring locations in the same way as the zip codes of Medicare beneficiaries; characteristics are aggregated to represent population demographics of the general population of zip codes within 6 miles of a monitoring location that is at most 150 km from a power plant. A summary of the data for the power-plant centered linkage appears in Table E.1.

E.2 Monitor-Centered Data Linkage

An alternative to the power-plant centered analysis above is the "monitor-centered" linkage, where the monitoring location (not the power plant) is the unit of analysis. Rather than link monitors to power plants based on an arbitrary distance radius, the monitor-centered linkage calculates, for a given monitoring location, the distance between that location and every power plant in the data. Power plant characteristics are then aggregated to the level of the ambient monitor by calculating a weighted average of all the power plants in the data, weighted by $1/d_{ij}^3$, where d_{ij} is the distance between the i^{th} monitor and the j^{th} power plant. This allows every power plant to contribute information to every monitor, with the closest power plants contributing the most and very distant power plants contributing virtually no information. After aggregating power-plant information to the monitor level, we then link monitors with Medicare and census data by enumerating all zip code centroids within a 6 mile radius of the monitoring location. Zip codes that are within 6 miles of more than one monitor are uniquely linked to the closest monitor. The final monitor-centered data set contains a record for each ambient monitor containing ambient $PM_{2.5}$ measures, distance-weighted averages of power plant characteristics, and aggregated population demographics and Medicare health outcomes among all zip codes within a 6 mile radius. Figure E.2 depicts this linkage schematically.

While the monitor-centered linkage permits every power plant to contribute information to every monitoring location, it presents an added complication in defining the accountability question and the causal effect of interest. Monitors are, of course, not "treated" with scrubbers, which complicates definition of an intervention and subsequent potential outcomes. For the monitor-centered analysis, we take a monitoring location to be "treated" with a scrubber if the distance-weighted average of SO_2 scrubber installation statuses among all power plants is larger than 0.5 and "untreated" if the distance-weighted average of SO_2 scrubber installation statuses is less than 0.05. The causal effects are then defined as the effect of having distance-weighted average of scrubber presence of greater than 0.50 versus having a distance-weighted average of scrubber presence of less than 0.05. The monitor-centered data set is described in Table E.2.

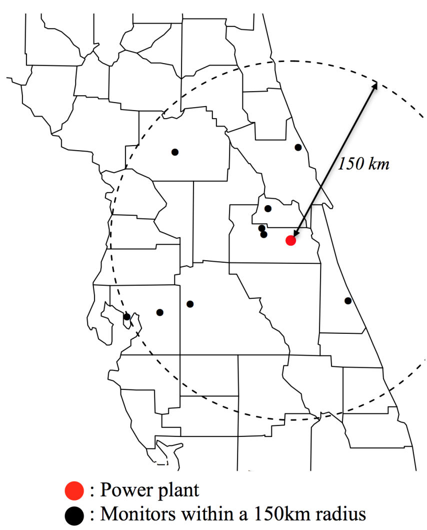

(a) Link monitors to power plant

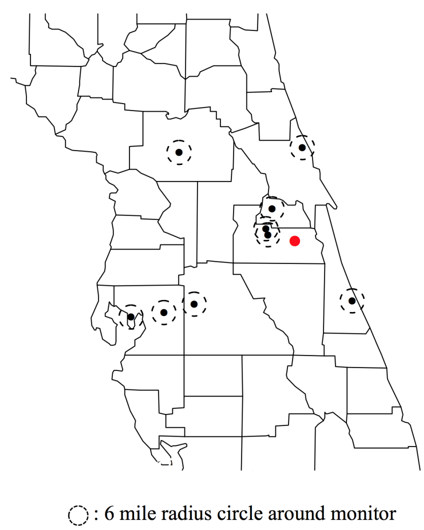

(b) Link medicare to power plant

Figure E.1: Power plant-centered linkage

Table E.1: Summary statistics for covariates and outcomes from the power plant-centered dataset

	Plants with scrubbers (n=39)		Plants w/o scrubbers (n=198)	
	Mean	SD	Mean	SD
Total Beneficiary	38728.90	63939.70	39935.70	87548.60
Person-Year	27768.10	45070.30	33462.50	68771.50
All CVD	2825.40	4779.60	3567.40	7271.60
Respiratory	872.40	1308.00	1105.90	2105.20
Total Death	2484.70	4166.00	2525.80	5082.50
$PM_{2.5}$	11.60	4.10	13.30	2.50
Temperature	14.00	4.60	13.90	3.40
Barometric Pressure	717.40	53.90	741.30	24.10
NO_x Scrub Indicator	0.90	0.20	0.80	0.40
Phase 1 Indicator	0.20	0.40	0.20	0.40
Coal Sulfur Content	1.50	1.10	0.80	0.60
NO_x (tons)	436.60	335.80	289.10	269.70
SO_2 (tons)	383.00	392.20	947.50	919.50
CO_2 (tons)	253869.10	176101.10	171834.60	140856.80
Smoking Rate	0.30	0.00	0.30	0.00
Mean Age	75.00	1.30	75.20	0.80
Female Rate	0.60	0.10	0.60	0.00
Black Rate	11.50	13.90	13.30	13.40
Age18-64 Rate	61.50	5.20	62.50	4.00
Urban Rate	75.80	25.80	73.00	27.60
No College Rate	52.00	11.90	51.50	10.90
Median Household Income	38702.60	10464.80	39450.70	10462.20
English Only Rate	87.50	14.10	91.00	8.00

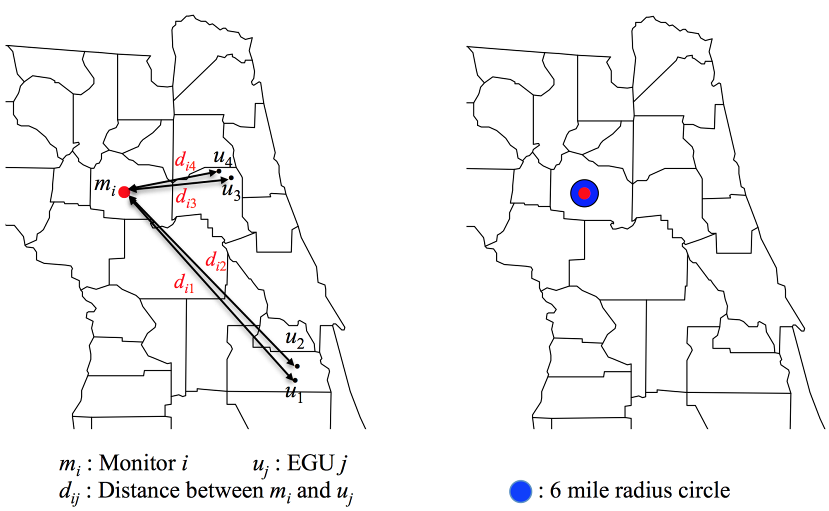

(a) Weight Power Plants to Monitors by Distance

(b) Link Medicare and Census to Monitors

Figure E.2: Monitor-centered linkage

Table E.2: Summary statistics for covariates and outcomes from the monitor-centered data set

	Plants with scrubbers (n=97)		Plants w/o scrubbers (n=256)	
	Mean	SD	Mean	SD
Total Beneficiary	19306.80	18727.10	13990.20	12805.90
Person-Year	13585.00	12259.00	11818.40	10679.40
All CVD	1396.10	1363.10	1250.30	1181.20
Respiratory	414.20	348.20	400.50	352.10
Total Death	1235.00	1208.30	891.20	812.40
$PM_{2.5}$	10.90	3.80	13.80	2.60
Temperature	15.00	5.30	13.40	3.10
Barometric Pressure	724.90	48.50	744.20	14.70
NO_x Scrub Indicator	0.90	0.10	0.70	0.40
Phase1 Indicator	0.20	0.30	0.20	0.30
Coal Sulfur Content	1.40	0.80	0.80	0.50
NO_x (tons)	417.00	182.60	199.60	160.20
SO_2 (tons)	539.30	358.40	691.90	585.40
CO_2 (tons)	239852.90	103860.90	124332.20	90499.50
Smoking Rate	0.30	0.00	0.30	0.00
Mean Age	75.30	1.10	75.30	0.90
Female Rate	0.60	0.00	0.60	0.00
Black Rate	9.90	13.30	14.60	16.60
Age18-64 Rate	61.40	4.20	62.80	4.70
Urban Rate	84.70	22.90	79.60	28.90
No College Rate	49.10	11.20	49.90	12.40
Median Household Income	39514.90	8996.30	41530.50	12973.40
English Only Rate	85.50	14.90	89.70	8.90

References

Banerjee, S., Gelfand, A. E., Finley, A. O., and Sang, H. (2008). Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **70,** 825–848.

Escobar, M. D. and West, M. (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association* **90,** 577–588.

Finley, A. O., Banerjee, S., and Carlin, B. P. (2007). spBayes: an R package for univariate and multivariate hierarchical point-referenced spatial models. *Journal of Statistical Software* **19,** 1–24.

Guo, J., Betancourt, M., Brubaker, M., Carpenter, B., Goodrich, B., Hoffman, M., Lee, D., Li, P., Morris, M., Trangucci, R., and Gelman, A. (2015). Rstan: The R Interface to Stan, version 2.4.

Jara, A., Hanson, T. E., Quintana, F. A., Mller, P., and Rosner, G. L. (2011). DPpackage: Bayesian Non and Semiparametric Modeling in R.

Jin, H. and Rubin, D. B. (2008). Principal stratification for causal inference with extended partial compliance. *Journal of the American Statistical Association* **103,** 101–111.

Müller, P., Erkanli, A., and West, M. (1996). Bayesian curve fitting using multivariate normal mixtures. *Biometrika* **83,** 67–79.

Nelsen, R. B. (1999). *An introduction to copulas*, volume 139. Springer Science & Business Media.

Rasmussen, C. E. (1999). The infinite Gaussian mixture model. In *NIPS*, volume 12, pages 554–560.

Rubin, D. B. (1980). Discussion of randomization analysis of experimental data: the Fisher randomization test by D. Basu. *Journal of the American Statistical Association* **75,** 591–593.

Sethuraman, J. (1994). A Constructive Definition of Dirichlet Priors. *Statistica Sinica* **4,** 639–650.

Sobel, M. E. (2006). What Do Randomized Studies of Housing Mobility Demonstrate? *Journal of the American Statistical Association* **101**, 1398-1407.

Taddy, M. A. (2008). Bayesian nonparametric analysis of conditional distributions and inference for Poisson point processes. ProQuest.

Wackernagel, H. (2003). Multivariate geostatistics: an introduction with applications. Springer.

Zigler, C. M., Dominici, F., and Wang, Y. (2012). Estimating causal effects of air quality regulations using principal stratification for spatially correlated multivariate intermediate outcomes. Biostatistics (Oxford, England) 13, 289–302.