Published in final edited form as:

J R Stat Soc Ser A Stat Soc. 2015 February ; 178(2): 425–443. doi:10.1111/rssa.12073.

Estimating the health benefit of reducing indoor air pollution in a randomized environmental intervention

Roger D. Peng¹, Arlene M. Butz², Amber J. Hackstadt¹, D'Ann L. Williams³, Gregory B. Diette⁴, Patrick N. Breyse³, and Elizabeth C. Matsui⁵

¹Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health

²Division of General Pediatrics, Johns Hopkins School of Medicine

³Department of Environmental Health Sciences, Johns Hopkins Bloomberg School of Public Health

⁴Division of Pulmonary and Critical Care Medicine, Johns Hopkins School of Medicine

⁵Division of Pediatric Allergy and Immunology, Johns Hopkins School of Medicine

Abstract

Recent intervention studies targeted at reducing indoor air pollution have demonstrated both the ability to improve respiratory health outcomes and to reduce particulate matter (PM) levels in the home. However, these studies generally do not address whether it is the reduction of PM levels specifically that improves respiratory health. In this paper we apply the method of principal stratification to data from a randomized air cleaner intervention designed to reduce indoor PM in homes of children with asthma. We estimate the health benefit of the intervention amongst study subjects who would experience a substantial reduction in PM in response to the intervention. For those subjects we find an increase in symptom-free days that is almost three times as large as the overall intention-to-treat effect. We also explore the presence of treatment effects amongst those subjects whose PM levels would not respond to the air cleaner. This analysis demonstrates the usefulness of principal stratification for environmental intervention trials and its potential for much broader application in this area.

Keywords

asthma; Bayesian; clinical trial; particulate matter; principal stratification

1 Introduction

Numerous studies to date have indicated strong evidence of an association between ambient particulate matter (PM) and mortality and morbidity. Recent work has largely focused on the fine fraction of PM (PM_{2.5} microns in aerodynamic diameter) because of its ability to penetrate deep into the lungs. Fine PM has been shown to be associated with a range of both

cardiovascular and respiratory health outcomes (Brook et al., 2010; Dominici et al., 2006). Research has shown evidence that mortality and morbidity are associated with both daily changes in $PM_{2.5}$, so-called acute effects, as well as differences in long-term average levels across communities, or chronic effects (Pope and Dockery, 2006). Children have long been recognized by the U.S. Environmental Protection Agency as a susceptible population (Bateson and Schwartz, 2008). A recent large time series study of $PM_{2.5}$ and children in Atlanta, GA showed $PM_{2.5}$ and other traffic-related pollutants to be associated with emergency department visits for asthma and wheeze (Strickland et al., 2010).

In addition to ambient air pollution, pollution inside the home can pose a health risk, especially to children with asthma. Indoor PM has a number of sources, both natural and human-made. Sources include cigarette smoking, biomass burning in stoves and fireplaces, cooking, cleaning activities that resuspend dust, and outdoor PM that penetrates inside the home (Breysse et al., 2010). Although indoor PM differs from outdoor PM in source, chemical composition, and concentration, studies of indoor PM and asthma in children have shown an association with decreased lung function, slower lung growth, increases in airway reactivity, and increases in asthma symptoms and exacerbations (Koenig et al., 2005; McCormack et al., 2011; Ahluwalia and Matsui, 2011).

Interventions designed to lower outdoor PM levels typically must be implemented at the state or national level via regulations such as the U.S. National Ambient Air Quality Standards. These regulations are costly, complex in their implementation, and may take a long time to have an effect on either PM or population health (Zigler et al., 2012). By comparison, interventions to lower indoor PM are relatively inexpensive. However, evidence of their effectiveness at improving asthma outcomes is only still accumulating. For example, two randomized intervention studies found that the introduction of high-efficiency particulate air (HEPA) cleaners in the home reduced $PM_{2.5}$ and improved asthma symptoms in children (Eggleston et al., 2005; Butz et al., 2011). Other studies, focused on non-pollutant asthma triggers (i.e. allergens), have found that indoor environmental interventions can lead to improvements in health (Morgan et al., 2004).

Although in these studies we know that randomization to the active intervention group was associated with asthma improvements and was associated with reductions in $PM_{2.5}$, we are limited in our capacity to draw conclusions about whether it was the reduction in $PM_{2.5}$ that actually resulted in the improvements in asthma. The conventional approach in such trials is to see if there are correlations between the degree of reduction in $PM_{2.5}$ and the degree of improvement in asthma. However, because the change in $PM_{2.5}$ occurs post-randomization and is not under the investigator's control, such an analysis does not benefit from the randomized assignment and is susceptible to the kinds of potential confounding normally found in observational studies. Specifically, we cannot expect that the subjects who experienced a decrease in $PM_{2.5}$ and the subjects who did not are balanced with respect to their various baseline covariates as we can when comparing those assigned to intervention to those assigned to control. In particular for this setting, finding a correlation between $PM_{2.5}$ and asthma outcomes doesn't exclude the possibility that the true cause of the improvement was due to concomitant reduction in another environmental factor (or perhaps improvement in medication adherence) that is correlated with the reduction in $PM_{2.5}$.

Developing a better understanding of the factors that mediate the effects of an environmental intervention on asthma is important for identifying a causal role for indoor pollutants in asthma morbidity. Providing a strong body of evidence supporting a causal relationship between indoor exposure to pollutants and worse asthma would help support shifts in clinical guidelines for managing asthma and even public health policy. Furthermore, although we know that a comprehensive environmental intervention can improve asthma, we do not know what aspects of indoor exposure are most important to target, or more specifically, in the presence of scarce resources, which exposures should be prioritized for reduction. Current recommendations in the asthma guidelines simply advocate for an individually tailored multifaceted intervention. Ultimately, obtaining better evidence of the effectiveness of reducing indoor exposures can help us to optimize future environmental interventions so we can achieve the most positive health effect with the most efficient use of health care resources.

In this paper we examine the effect of a randomized air cleaner intervention on asthma symptoms and focus specifically on the subset of individuals for whom the intervention is most effective at reducing $PM_{2.5}$ levels. We use the approach of principal stratification (Frangakis and Rubin, 2002) to divide the study subjects into latent classes defined by their potential $PM_{2.5}$ levels under different treatment assignments and to estimate the effect of the treatment assignment within these principal strata. Our goal is to estimate the treatment effect on those study subjects whose $PM_{2.5}$ level is responsive to the treatment assignment. Our approach further allows us to explore possible effects of treatment amongst those subjects for whom $PM_{2.5}$ levels are constant as treatment assignment varies. Large effects in these other principal strata provide some evidence that the intervention works potentially via some other mechanism (i.e. other than lowering $PM_{2.5}$) that would require further exploration.

1.1 The PREACH Study

The analysis in this paper is based on data from the Particulate Reduction and Education in City Homes (PREACH) study, which was a randomized 3-arm controlled environmental intervention to assess the effect of introducing air cleaners into childrens' homes on improving asthma symptoms (Butz et al., 2011). A requirement for being included in the study was the child had to live with a smoker who smoked at least 5 cigarettes per day and lived there at least 4 days a week. The three arms of the trial were:

- control: received asthma education from a nurse during four home visits
- air cleaner: received two air cleaners (one for the child's bedroom and one for the family/living room) in addition to the asthma education
- air cleaner + health coach: received the two air cleaners, asthma education, and a behavioral intervention aimed at supporting the primary caregiver in changing smoking behaviors to reduce the child's second hand smoke exposure.

The outcome that we focus on here is symptom-free days, which is defined as the number of days in the past two weeks that were free of asthma symptoms. The pollutant of interest here

is $PM_{2.5}$. While there were other pollutants to consider in this study, the intervention was primarily designed to reduce $PM_{2.5}$ levels in the home. In the original study the addition of the health coach was not associated with either a reduction in $PM_{2.5}$ or an improvement in asthma symptoms. Therefore, for the purposes of this analysis, we have combined the two air cleaner groups into a single intervention group.

The original analysis of Butz et al. (2011) compared the change in symptom-free days from baseline to followup across treatment groups and found that subjects in the air cleaner group experienced a larger increase in symptom-free days between baseline and followup than the control group. They also found that the assignment to the air cleaner group was associated with larger decreases in $PM_{2.5}$ from baseline to followup compared to the control group. Hence there was some evidence that the treatment was operating through the hypothesized mechanism of lowering $PM_{2.5}$ but no formal analysis of this question was conducted.

2 Methods

Our goal in this study was to develop evidence for the effect of lowering pollution in a child's home on the child's symptom-free days in a randomized air cleaner intervention. Because the change in the $PM_{2.5}$ level in this study is not directly controlled, as is the case in all indoor environmental control trials in asthma, we cannot easily assess any causal effects of lowering $PM_{2.5}$ by comparing outcomes with observed $PM_{2.5}$ levels. To address this key question we used the method of principal stratification (Frangakis and Rubin, 2002) to investigate the effect of the air cleaner intervention on those children whose $PM_{2.5}$ levels were responsive to the introduction of the air cleaner. Following the approach of Imbens and Rubin (1997), we take a Bayesian formulation of the problem and estimate the posterior distributions of the parameters of interest.

For each subject i ($i = 1, \dots, n$) we observe a baseline level of symptom-free days SFD_{0i} and a followup (post-randomization) level of symptom-free days measured 6-months later, SFD_{1i} . We also observe their baseline level of in-home $PM_{2.5}$, which we denote as PM_{0i} and their followup level PM_{1i} . Each subject is also assigned to an intervention group indicated by Z_i where $Z_i = 0$ indicates the control group and $Z_i = 1$ indicates the air cleaner (treatment) group. Finally, we observe a set of baseline covariates X_i .

Our primary outcome is the difference in symptom-free days between the baseline measurement and the followup measurement $Y_i = SFD_{1i} - SFD_{0i}$. We also calculate the difference in $PM_{2.5}$ levels between baseline and followup as $S_i^* = PM_{1i} - PM_{0i}$. (Throughout this analysis we will use the terms “symptom-free days” and “ $PM_{2.5}$ levels” to indicate the changes of these values from baseline.) The original study compared Y_i between the two treatment groups and found that on average, assignment to the air cleaner group resulted in an increase of 1.36 symptom-free days, a significant increase both statistically and clinically. However, the original analysis ignored the value of S_i^* for each subject when evaluating the benefit of the intervention. Here, we are interested in estimating the change in symptom-free days for those subjects where the air cleaner had a beneficial impact on $PM_{2.5}$ levels.

Although the value of S_i^* (the change in $PM_{2.5}$ level from baseline to followup) is technically continuous, we dichotomize its value and denote $S_i = 1$ if $S_i^* < -20$ and $S_i = 0$ otherwise. In other words, if we observe a decrease in $PM_{2.5}$ of at least $20 \mu g/m^3$ then we set $S_i = 1$. The reasons for using this cutoff were (1) it is a clinically relevant decrease in $PM_{2.5}$ that is achievable through this type of intervention, as we have seen in previous studies (Eggleston et al., 2005); (2) previous studies of ambient air pollution had found changes of this magnitude to be associated with meaningful changes in health outcomes, and (3) it was approximately the mean decrease in $PM_{2.5}$ observed in the air cleaner group. We explore the sensitivity of our findings to the choice of the cutoff in the Results below.

2.1 Potential Outcomes Framework and Principal Stratification

Let random variables $Y(0)$ and $Y(1)$ be the potential values of the outcome (symptom-free days) for a subject under the control or the air cleaner scenarios, respectively. For a given subject i , the effect of the treatment (air cleaner) can be written as $\delta_i = Y_i(1) - Y_i(0)$.

Although we never observe both potential outcomes for any subject, because of the randomization of the subjects to the intervention groups in this study, we can estimate the average effect (Holland, 1986) of the air cleaner intervention,

$\delta = E[Y(1) - Y(0)] = E[Y|Z=1] - E[Y|Z=0]$ using observed data with

$$\hat{\delta} = \bar{Y}_1 - \bar{Y}_0 = \frac{1}{n} \sum_{i=1}^n Y_i Z_i - \frac{1}{n} \sum_{i=1}^n Y_i (1 - Z_i).$$

Let $S_i(0)$ and $S_i(1)$ be the potential values of $PM_{2.5}$ for subject i under either the control or the air cleaner scenarios, respectively. Each subject is presumed to have potential values for each of these conditions, of which we will only observe one in the course of the study (either the control or the air cleaner value). Because S_i is binary, there are four possible combinations of the pairs $S_i(0)$, $S_i(1)$ that an individual could potentially experience. These combinations are called principal strata and are shown in Table 1.

Let C_i indicate the principal stratum membership of each subject i . C_i takes values A , N , C , or D for always-reduced, never-reduced, cleaner-under-treatment, and dirtier-under-treatment. In principle, every subject can be categorized into one of the four principal strata shown in Table 1. However, because we do not observe both of the potential outcomes $S_i(0)$ and $S_i(1)$ for any subject, the value of C_i is missing for all subjects.

From Table 1 we highlight the “cleaner-under-treatment” principal stratum as one of particular interest. In this group of subjects, $PM_{2.5}$ would only be lowered if they were assigned to the air cleaner group and would not be lowered if they were assigned to the control group. This group of subjects essentially follows the intent of the study design and the treatment effect on these subjects provides some information as to whether any change in asthma symptoms is due to the lowering of $PM_{2.5}$.

While the treatment effect in the cleaner-under-treatment group is of primary interest, we are also interested in possible treatment effects in either the always-reduced or never-reduced groups. In both of these groups, assignment to either the control or air cleaner group has no effect on the level of $PM_{2.5}$, so a large treatment effect in these subjects would not be explained by a lowering of $PM_{2.5}$. It is possible, for example, that the intervention modified

some other harmful pollutant in the air or perhaps modified a subject's behavior in a manner that affected symptoms.

If the principal stratum membership for each subject were somehow known, we could simply use the observed data to estimate the average causal effect in the cleaner-under-treatment stratum by using the difference of mean response between subjects in the air cleaner and control groups within that stratum (i.e. the average effect within stratum). However, because the stratum membership C_i is unknown, a goal of analysis is to infer the value of C_i so that we can estimate the average effect within each stratum.

Table 2 shows how the observed data (Z_i and S_i) relate to the principal stratum categorizations. For each combination of treatment group Z_i and $PM_{2.5}$ level S_i the subjects represent a mixture of people in two different principal strata. For example, subjects in the air cleaner group where $PM_{2.5}$ levels were observed to be lowered are either in the cleaner-under-treatment or always-reduced group. Subjects in the control group who had a reduction in PM are in the always-reduced or dirtier-under-treatment group.

2.2 Assumptions

We make a number of assumptions to allow for estimation of relevant parameters and for simplifying the computation. The first is the stable unit treatment value assumption (Rubin, 2005), which assumes that (a) the values of $Y_i(0)$ and $Y_i(1)$ are not affected by the randomization assignment of any other subject $j \neq i$ and (b) no matter how subject i receives the treatment (or control), the outcome that would be observed is $Y_i(1)$ (or $Y_i(0)$). The second assumption is that the treatment assignment is ignorable, which is satisfied in this case because the treatment assignment was randomized.

A key assumption that we make is that there are no members of the dirtier-under-treatment group, i.e. there are no subjects for whom $PM_{2.5}$ would be lowered under the control scenario but would *not* be lowered with the air cleaner. This assumption is commonly referred to as monotonicity (Angrist et al., 1996) because we assume $S_i(1) \geq S_i(0)$ in Table 1. In the context of this intervention, it seemed implausible that a subject would actively lower his/her in-home $PM_{2.5}$ level under the control situation but somehow attempt to thwart the air cleaner by maintaining or increasing $PM_{2.5}$ levels under the air cleaner scenario. Assuming no one in the dirtier-under-treatment group simplifies Table 2 because it allows us to immediately identify the principal strata of some of the subjects: Subjects who received the air cleaner but did not lower their $PM_{2.5}$ are all in the never-reduced group and subjects who received the control but did lower their $PM_{2.5}$ are all in the always-reduced group.

Exclusion restrictions, which are assumptions that there is no effect of the intervention in the always-reduced or never-reduced strata, are commonly made if we hypothesize that the effect of the intervention can only go through the intermediate variable $PM_{2.5}$ (Angrist et al., 1996; Imbens and Rubin, 1997). However, because we are interested in examining whether the intervention was beneficial without lowering $PM_{2.5}$, we do not use any exclusion restrictions here. We explore the sensitivity of the model to the use of exclusion restrictions in the Results.

2.3 Model

We model the observed response Y_i as following a Normal distribution with mean μ_{pz} and standard deviation σ_p^2 , where $p \in \{c, a, n\}$ represents the principal stratum (recall no one is in the dirtier-under-treatment group) and $z \in \{0, 1\}$ indicates the treatment group. The use of separate variances within each principal stratum mirrors an approach taken by Gallop et al. (2009). Note that the outcome Y_i represents a difference in symptom-free days between followup and baseline. Because symptom-free days is recorded for the previous 14 days, the observed response Y_i can only range between -14 and 14 . We nevertheless chose to model the response with a Normal distribution because the observed variability in symptom-free days was modest and was reasonably captured by the Normal. Furthermore, the use of the Normal distribution greatly simplified the computation in model fitting.

Let C_i indicate the unobserved principal stratum membership for subject i and X_i be a vector of observed baseline covariates. We model C_i as a multinomial random variable so that $C_i \in \{c, a, n\}$. C_i then has the density

$$\gamma(C_i|X_i) \propto \pi_{ic}^{1\{C_i=c\}} \pi_{ia}^{1\{C_i=a\}} \pi_{in}^{1\{C_i=n\}}$$

where for each stratum p , we have

$$\pi_{ip} = \frac{\exp(X_i' \beta_p)}{\sum_{j \in \{c, a, n\}} \exp(X_i' \beta_j)} \quad (1)$$

We normalize these probabilities by setting β_n equal to a vector of zeros (Hirano et al., 2000). Given the stratum indicators $\mathbf{C} = (C_1, \dots, C_n)$, the information found in Table 2, and the assumption of no one in the dirtier-under-treatment group, the complete data posterior distribution of the model parameters $\mu = (\mu_{c0}, \mu_{c1}, \mu_{a0}, \mu_{a1}, \mu_{n0}, \mu_{n1})$, $\sigma^2 = (\sigma_c^2, \sigma_a^2, \sigma_n^2)$, and $\beta = (\beta_c, \beta_a, \beta_n)$ is

$$f(\mu, \sigma^2, \beta | \mathbf{C}, \mathbf{Y}, \mathbf{Z}, \mathbf{S}) \propto \left\{ \prod_{z \in \{0,1\}} \prod_{p \in \{c,a,n\}} \prod_{Z_i=z, C_i=p} \pi_{ip} \varphi(Y_i | \mu_{pz}, \sigma_p^2) \right\} p(\mu) p(\sigma^2) p(\beta)$$

where φ is the density of the Normal distribution.

We sampled from the posterior distribution of the model parameters and the stratum membership indicator C_i using a hybrid Gibbs sampler. The full conditional distributions for each of the parameters are shown in the Appendix. Sampling was conducted for 20,000 iterations and convergence of the sampler was judged by examining standard errors for the posterior means of the parameters estimated via batch means (Jones et al., 2006). A Metropolis-Hastings rejection step was used to sample the full conditional distribution of β_p conditional on the observed data and current parameter values. The remaining full conditional distributions could be sampled directly. The primary targets of inference in this

analysis were the differences in the stratum-specific Normal means, $\theta_p = \mu_{p1} - \mu_{p0}$ for $p \in \{c, a, n\}$.

2.4 Prior distributions

For each of the μ_{pz} parameters we use a Normal prior distribution with mean 0 and variance 100. For the σ_p^2 parameters we use inverse gamma $IG(2, 3)$ priors, and for the β_p parameters we use a Normal distribution with mean 0 and variance 1. We chose the priors for σ_p^2 and β_p to be somewhat informative given the modest number of subjects included in the analysis and the constraints on the parameters imposed by the problem. For the inverse gamma priors on σ_p^2 , the parameters chosen focus 99% of the mass in the range of 0 to 20.

3 Results

There were a total of 126 subjects in the original PREACH study. Of the 126 original subjects, 48 were missing $PM_{2.5}$ or symptom data at either the baseline or the followup visit. Of the remaining subjects, 3 were missing information on baseline covariates, leaving 75 subjects to be included in this analysis. At baseline, the mean symptom-free days over a 2-week period was 11.9 days with a between-subject standard deviation of 2.7 days. The median $PM_{2.5}$ level at baseline was $30.0 \mu\text{g}/\text{m}^3$ with an IQR of $38.8 \mu\text{g}/\text{m}^3$. Although we took a subset of the original study for this analysis, the treatment and control groups were generally balanced with respect to key baseline covariates. When comparing treatment to the control group, subjects in the treatment group had somewhat higher baseline $PM_{2.5}$ levels (average of $42 \mu\text{g}/\text{m}^3$ vs. $38 \mu\text{g}/\text{m}^3$), lower nitrogen dioxide (NO_2) levels (average of 31 ppb vs. 34 ppb), and were on average 2 months younger. None of these differences were statistically significant and there were no significant differences in asthma severity between the two groups. Greater detail about the characteristics of the study population can be found in Butz et al. (2011).

Figure 1 shows the analysis from the original study applied to the subset of individuals used here. Figure 1(a) shows the change in symptom-free days from the baseline to 6-month followup in the control and air cleaner groups. Here we see that the air cleaner group, on average, saw an increase of about 2 symptom-free days whereas the control group, on average, saw no change. Figure 1(b) shows the change in the in-home level of $PM_{2.5}$ in the control and air cleaner groups. Here we see that the control group generally experienced a slight increase in $PM_{2.5}$ whereas the air cleaner group generally saw a decrease in $PM_{2.5}$ levels. The average change in $PM_{2.5}$ in the air cleaner group was a decrease of $18.4 \mu\text{g}/\text{m}^3$ while the average change in the control group was an increase of $2.2 \mu\text{g}/\text{m}^3$. The difference in the change in $PM_{2.5}$ over the 6-month followup between the air cleaner and control groups was $20.6 \mu\text{g}/\text{m}^3$.

The boxplots in Figure 1 suggest that the air cleaner intervention improved symptoms and was effective at decreasing $PM_{2.5}$. However, we are interested in whether the assignment to the air cleaner group improved symptom outcomes in those subjects where $PM_{2.5}$ was decreased significantly. Here we define a significant decrease as $20 \mu\text{g}/\text{m}^3$.

We applied the model described in Section 2.3. After conducting exploratory analysis, the baseline covariates chosen for the model in (1) for predicting the principal stratum membership were age, baseline levels of $\text{PM}_{2.5}$ and NO_2 , and baseline asthma severity. Based on the posterior distribution of the principal stratum membership, we estimate that there were 22 subjects in the cleaner-under-treatment group, 16 in the always-reduced group, and 37 in the never-reduced group. Figure 2 shows the posterior means and 95% posterior intervals for the change in symptom-free days by intervention group and principal stratum. Specifically, Figure 2 shows posterior means for μ_{c0} , μ_{c1} , μ_{a0} , μ_{a1} , μ_{n0} , and μ_{n1} . Here we can see that for the cleaner-under-treatment stratum, the subjects in the control group saw no change in symptom-free days from baseline to followup whereas the air cleaner group saw an increase of over 6 symptom-free days. For the always-reduced stratum there was little difference in the outcomes for the control and air cleaner groups as both saw modest increases in symptom-free days. For the never-reduced group, the control group had a large decrease in symptom-free days while the air cleaner group saw a modest increase, the net effect being an increase in symptom-free days.

The first row of Table 3 shows the estimate of the overall treatment effect on symptom-free days applying the same analysis used in the original study. On average, the air cleaner group experienced a change in symptom-free days from baseline to 6-months that was 2.1 (95% CI: 0.4, 3.8) days greater than the control group. Rows 2–4 of Table 3 show the posterior means (95% posterior intervals) for the treatment effects in the cleaner-under-treatment, always-reduced, and never-reduced principal strata under models with different restrictions.

Recall that the treatment effects $\hat{\theta}_c$, $\hat{\theta}_a$ and $\hat{\theta}_n$ are the differences in the within-stratum means shown in Figure 2. Row 2 shows our Full Bayesian Model with none of the exclusion restrictions described in Section 2.2. Here, the treatment effect in the cleaner-under-treatment group was an increase in 6.2 (95% PI: 2.7, 11.6) symptom-free days. The treatment effect in the cleaner-under-treatment group was approximately three times as large as the overall treatment effect. The treatment effect estimated in the never-reduced group was 5.0 (95% PI: 0.3, 10.2) symptom-free days, similar in magnitude to the cleaner-under-treatment effect. From Figure 2 we can see that the effect is primarily due to a large decrease in symptom-free days in the control group. The treatment effect in the always-reduced group was essentially zero as both control and air cleaner groups saw a modest increase in symptom-free days over the course of the study.

For the multinomial regression model predicting stratum membership we found that baseline $\text{PM}_{2.5}$, NO_2 , age, and asthma severity had modest associations with membership in the different strata. In particular, children with more severe asthma were somewhat less likely to be in the cleaner-under-treatment stratum while those with higher baseline $\text{PM}_{2.5}$ or NO_2 were more likely to be in the cleaner-under-treatment or always-reduced group.

3.1 Sensitivity analysis

We checked the sensitivity of the model fitting to the use of exclusion restrictions, the cutoff used to define the intermediate variable, and the presence of outliers. We also compared the prior and posterior distributions for key parameters.

To test the sensitivity of the model to exclusion restrictions we fit two additional models: one which assumed no treatment effect in the always-reduced group ($\theta_a = 0$) and one which assumed no treatment effect in both the always-reduced and never-reduced group ($\theta_a = \theta_n = 0$). The latter analysis is similar to an instrumental variables type of analysis in that it assumes that the effect of the treatment assignment can only go through a change in the $PM_{2.5}$ variable (Angrist et al., 1996). In the model that excludes an effect for the always-reduced group, we see that both the never-reduced and the cleaner-under-treatment effects were largely unchanged (Table 3, row 3). In the model with both exclusions, the cleaner-under-treatment effect was larger than the effect found in the previous two models (Table 3, row 4). However, given the relatively large effect found for the never-reduced group in the Full Bayesian Model, the assumption that $\theta_n = 0$ seems implausible.

In Figure 3 we compare the simulated posterior distribution of each of the treatment effect parameters with their respective prior distributions. The priors chosen were relatively flat and one can see from Figure 3 that the posterior is far more concentrated than the prior. This pattern is reassuring in that it indicates that there is some information in the dataset to inform the treatment effect parameters, even in the Full Bayesian Model which does not contain any exclusion restrictions.

We conducted the same comparison of the posterior and prior for the σ_p^2 parameters and the results are shown in the Appendix in Figure A1. Here, we see that the posteriors for σ_c^2 (cleaner-under-treatment) and σ_n^2 (never-reduced) appear to deviate substantially from their priors. The posterior for σ_a^2 (always-reduced) only deviates modestly from its prior suggesting that the amount of information in the dataset for estimating σ_a^2 was not substantial.

Because the model that we fit to the data is not fully nonparametrically identified, the results may be sensitive to the specific probability distributions that we used to model the response within principal strata. Therefore, we conducted a sensitivity analysis where we modeled the within-stratum response with a t -distribution with 2.5, 3, or 8 degrees of freedom, to accommodate heavier tails (Figure S1). Overall, the results were similar; in particular the treatment effects for the never-reduced group were unchanged. Depending on how many degrees of freedom were used for the always-reduced and cleaner-under-treatment group, there appeared to be some lack of identifiability associated with their respective treatment effects. We believe that this issue is attributable not specifically to the use of the t -distribution here, but rather to the differing variances used to model the response. Because our original model does not fix variances, but rather integrates over a range of possible variances (i.e. σ_a^2 and σ_c^2), we believe that the uncertainty introduced by this partial lack of identifiability is incorporated into the posterior intervals for the treatment effects.

We also tested the sensitivity of the analysis approach to the cutoff used to dichotomize the change in $PM_{2.5}$ from baseline to followup. While the choice of the cutoff is not purely a statistical one and we rely on the clinical usefulness of a cutoff in this context, we re-ran the analysis using a different cutoff merely to check the sensitivity of the methodology itself to specific thresholds. For the sensitivity analysis we defined a significant change in $PM_{2.5}$ as a

10 $\mu\text{g}/\text{m}^3$ decrease from baseline to followup. The results are shown in the Appendix in Figure A2 and Table A1. One can see there that the overall patterns are similar to the results above but the effect sizes are somewhat smaller, as one might expect from using a smaller threshold for the decrease in $\text{PM}_{2.5}$.

Finally, we checked to see if the results were sensitive to possible outliers in the response variable. We re-ran the analysis excluding the three outliers shown in the air cleaner group in Figure 1(a). The results in this subset analysis showed the same patterns found in Figure 2 and Table 3.

3.2 Missing data

There were a large number of subjects from the original study who were excluded from this analysis, primarily because of missing $\text{PM}_{2.5}$ data. In addition to checking the balance of key baseline covariates between the treatment and control groups above, we explored the sensitivity of the principal stratum treatment effects to the missing data via multiple imputation. Here, we assumed a multivariate Normal distribution on key baseline covariates (NO_2 , age, baseline $\text{PM}_{2.5}$) as well as on the followup $\text{PM}_{2.5}$ and multiply imputed 1,000 complete datasets for the entire original cohort of 126 subjects. Using these imputed datasets, we fit our original model and combined the posterior means and variances in a manner that incorporates the extra uncertainty due to the imputation (Little and Rubin, 2002). With the multiple imputation, we found that the overall pattern of treatment effects between the three principal strata was preserved, but the uncertainty about the treatment effects in the cleaner-under-treatment and always-reduced groups was greatly increased as reflected in the increased width of the posterior intervals (Figure S2).

3.2.1 Comparison with full cohort—The missing data from the original study additionally raise a question as to whether the included subjects were somehow different from the larger group. This question goes to whether we may be able to apply the results found here to the larger group of children, and perhaps beyond. To explore this question we used a logistic regression model to predict membership in the subset of subjects included in this analysis (Stuart et al., 2011). Covariates included in this model were age, baseline $\text{PM}_{2.5}$, asthma severity (“intermittent”, “mild persistent”, “moderate persistent”, “severe persistent”), sex, and caregiver’s educational status. On average, older children were more likely to be included in the analysis subset (odds ratio of 1.3 [95% CI: 1.0, 1.6] for a 1 year difference) and children with more severe asthma were less likely to be included. For children with “severe persistent” asthma the odds ratio of being included in this analysis relative to children with “intermittent” asthma was 0.3 (95% CI: 0.07, 0.8).

4 Discussion

In this paper we applied the method of principal stratification to estimate the respiratory health benefit of reducing indoor PM air pollution in homes of children with asthma. We found that a subset of children for whom $\text{PM}_{2.5}$ would be reduced by at least 20 $\mu\text{g}/\text{m}^3$ experienced an increase of about 6 symptom-free days, an outcome that was nearly three times as large as the overall effect of the intervention. Our approach was applied to data from an intervention study where the target of the intervention ($\text{PM}_{2.5}$) was not directly

manipulated. Nevertheless, we show that, with some caveats noted below, it is possible to generate strong evidence about a causal role for reducing indoor PM and improving asthma symptoms. We took a Bayesian approach here which allowed us to impose some constraints on the estimation of the parameters of interest using relevant prior knowledge.

Our analysis found some evidence that the intervention had an effect in the never-reduced group, for whom $PM_{2.5}$ would not have been affected by the intervention (but not in the always-reduced group). This is somewhat surprising in that the intervention (air cleaners) was not specifically designed to reduce any other relevant environmental factor. However, it is possible, for example, that the use of the air filters could have removed airborne allergens residing on particles. In this asthmatic population, such a change in allergen levels could have had an effect. The definition of the cutoff for determining a reduction in $PM_{2.5}$ (a decrease of $20 \mu g/m^3$) may have affected this outcome as some of the never-reduced group may have in fact had a reduction in $PM_{2.5}$ that was less than $20 \mu g/m^3$ but still non-trivial. This is one of the trade-offs one must make in dichotomizing a continuous intermediate variable. While simplicity is gained by such an approach, there may be some ambiguity introduced into the interpretation of the principal stratum effects.

Of course, there is always the possibility that some other unobserved shared feature of the never-reduced group explains the treatment effect. Certainly, this relatively large positive effect in the never-reduced group warrants further exploration. Separately, in our sensitivity analysis, the estimates of the treatment effect in the cleaner-under-treatment stratum did not appear to be affected by the use of exclusion restrictions.

Our analysis contained a number of limitations. Although we found a strong effect of the intervention in the cleaner-under-treatment stratum that we hypothesize is attributable to a reduction in $PM_{2.5}$, it is possible that some other covariate explains the difference in treatment effect for this sub-group. For example, the cleaner-under-treatment stratum subjects had somewhat less severe asthma and had higher baseline $PM_{2.5}$ levels, but these associations were relatively modest. Our dichotomization of $PM_{2.5}$ was a choice that was made to accommodate the analysis and to simplify the interpretation of the results, but it is possible (with suitable additional assumptions) to treat $PM_{2.5}$ as a continuous intermediate variable (e.g. Zigler et al., 2012). In particular, as noted above, the use of a cutoff does not distinguish between subjects right at the boundary and those with much smaller decreases in $PM_{2.5}$. An attractive alternative would be the methodology of Schwartz et al. (2011), who develop a very flexible model for estimating principal stratification treatment effects without having to dichotomize the intermediate variable. However, in this context, the motivation behind dichotomizing $PM_{2.5}$ centered around the desire to identify a specific treatment or intervention (e.g. reduce $PM_{2.5}$ by at least $20 \mu g/m^3$). In addition, we did not believe that our modest dataset of 75 observations would allow for a much more complex model to produce useful results. However, we believe the use of continuous measures could be important in other contexts, such as evaluating costs and benefits of $PM_{2.5}$ reduction.

We took a Bayesian approach in this analysis in order to incorporate existing prior knowledge of the effects of environmental interventions on asthma symptoms. In general, there appeared to be a reasonable amount of information in the data to estimate the

parameters of interest, in particular the mean parameters for each stratum/treatment combination. However, the comparison of the posterior and prior for the stratum-specific variances indicated somewhat less information for estimating these parameters. It is possible that the modest size of the dataset limited our ability to make precise inferences about the variances and therefore the results are somewhat more driven by the prior distributions. In our investigation, using substantially less informative priors on the variances lead to difficulties in getting the MCMC procedure to converge properly. While the Bayesian approach certainly had its advantages here, it may be useful to explore alternative modeling frameworks in similar studies.

Missing data were an issue with this cohort as many did not have followup values of $PM_{2.5}$. As a result, only 75 of the 126 subjects were available for the principal stratification analysis. We believe in our case that the missing data did not affect the validity of the estimated treatment effect because the treatment and control groups remained balanced on key baseline covariates. However, we believe it could affect our ability to generalize the estimated treatment effect to a broader population of asthmatics. In particular, the subjects included in the analysis were older and had less severe asthma than the broader cohort. Hence, further investigation, such as sensitivity analysis (Stuart et al., 2011), may be useful in obtaining a better understanding of what treatment effects of air cleaners would be in a more representative population.

4.1 Causal inference for environmental interventions

In general, the use of principal stratification to investigate interventions that potentially operate through an intermediate variable is a subject of considerable recent debate in the literature (e.g. Pearl, 2011, and subsequent discussion). An alternative approach in this setting is the causal mediation analysis framework arising from the work of Robins and Greenland (1992) and others which involves estimation of direct and indirect effects (Imai et al., 2010). Here, the direct effect is the health effect of the air cleaner that is separate from its effect on lowering $PM_{2.5}$ while the indirect effect is the health effect that is attributable to lowering $PM_{2.5}$. This approach assumes that both the intervention and the intermediate factor are manipulable and define potential outcomes on pairs of indices rather than on a single index as we have done here. Whether one can hypothesize a mechanism for separately manipulating the intermediate variable will depend on the specific setting. For example, with an intermediate variable like indoor air pollution, one may consider directly manipulating the source of the pollutant, if it is known.

Our primary motivation for using principal stratification for this problem was to explore the possibility that there was a sub-population of children who experience significant health benefits from the introduction of air cleaners in the home. Previous trials in this area had only seen modest health benefits from introducing environmental interventions but none had seriously examined the question of whether that effect might be modulated by the change in the targeted environmental intermediate variable. Here, principal stratification allowed us to estimate treatment effects in sub-populations that were associated with a substantial reduction in $PM_{2.5}$ and to demonstrate that those treatment effects were of a clinically relevant size, deserving of continued investigation.

The fact that we found a relatively large treatment effect in the cleaner-under-treatment group is highly relevant, but further study would be required to obtain a better understanding of the causal mechanism of the treatment. As noted by VanderWeele (2011), the treatment effect estimated for the cleaner-under-treatment group includes possible direct and indirect effects, and so we cannot conclude from our results that the treatment is “mediated” through a reduction in $PM_{2.5}$ without additional assumptions (Mealli and Mattei, 2012). In this setting, we believe the causal mediation analysis framework could potentially be useful when attempting to attribute the effects of the treatment to various intermediaries (i.e. pollutants vs. allergens). Such an analysis could allow one to prioritize interventions on certain environmental irritants when resources are insufficient to target everything at once.

The use of causal inference techniques such as principal stratification in environmental intervention studies is thus far uncommon but we believe the application of these methodologies has much potential in this area. In particular, they provide a tool for investigating in greater detail the effectiveness of targeting certain indoor pollutants for improving asthma morbidity in children. Indoor environmental control strategies often target multiple asthma triggers such as particulate matter, gases, or allergens, and being able to tease apart the magnitudes of the effects of reducing each environmental exposure can provide valuable evidence in the development of guidelines, public health policy, and clinical recommendations.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

5 Acknowledgments

This work was supported by Awards R01ES019560, R21ES020152, T32012871, P50ES015903, P01ES018176, and P30ES03819 from the National Institute of Environmental Health Sciences. This work was also supported by Grant R832139 from the U.S. Environmental Protection Agency (EPA), Award Numbers R01AI070630 and U01AI083238 from the National Institute of Allergy and Infectious Diseases and by the Johns Hopkins University School of Medicine General Clinical Research Center grant number M01-RR00052, from the National Center for Research Resources. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute of Environmental Health Sciences, the National Institute of Allergy and Infectious Diseases, the U.S. EPA, or the National Institutes of Health.

Appendix

A Additional Sensitivity Analysis

Figure A1.

Comparison of the posterior (histogram) and prior (solid line) distributions for σ_c^2 , σ_a^2 and σ_n^2 under the Full Bayesian Model.

Figure A2.

Posterior mean (95% posterior interval) of change in symptom-free days from baseline to 6-month followup by principal stratum and intervention group in the PREACH study. The cutoff for defining a significant change in $\text{PM}_{2.5}$ was a decrease of $10 \mu\text{g}/\text{m}^3$.

Table A1

Posterior means (95% posterior intervals) of the association between treatment assignment and symptom-free days in the cleaner-under-treatment ($\hat{\theta}_c$), always-reduced ($\hat{\theta}_a$), and never-reduced ($\hat{\theta}_n$) principal strata, by model type. The cutoff for defining a significant change in PM_{2.5} was a decrease of 10 $\mu\text{g}/\text{m}^3$.

	$\hat{\theta}_c$	$\hat{\theta}_a$	$\hat{\theta}_n$	Overall
Original Analysis				2.1 _(0.4,3.8)
Full Bayesian Model	4.4 _(1.7,7.7)	-0.3 _(-1.3,0.6)	4.6 _(-0.3,10.3)	
Exclude Always-reduced	4.4 _(1.7,7.6)	0	4.6 _(-0.5,10.3)	
Exclude Always- and Never-reduced	5.5 _(2.0,9.8)	0	0	

B Details of Gibbs Sampler

We present the full conditional distributions for the hybrid Gibbs sampler used to fit the model described in Section 2.3. The sampler alternates between sampling from the full conditionals for the stratum indicator C_i , the stratum/treatment means μ_{pz} , stratum variances σ_p^2 , and stratum regression coefficients β_p .

B.1 Full conditional for C_i

There are two cases when C_i is always known. When $Z_i = 1$ and $S_i = 0$, we know that $C_i = n$ (never-reduced) because we assume there is no one in the "dirtier under treatment" group. Similarly, when $Z_i = 0$ and $S_i = 1$, we know that $C_i = a$ (always-reduced). In those cases we do not need to sample C_i . There are two cases where we sample C_i :

1. $Z_i = 1, S_i = 1$: In this case C_i is either c or a , so we sample from the binomial distribution where the conditional probability for class c is

$$P(C_i = c | Z_i = 1, S_i = 1, \mu, \sigma^2, \beta, \mathbf{Y}_i) = \frac{\pi_{ic} \varphi(Y_i | \mu_{c1}, \sigma_c^2)}{\pi_{ic} \varphi(Y_i | \mu_{c1}, \sigma_c^2) + \pi_{ia} \varphi(Y_i | \mu_{a1}, \sigma_a^2)}$$

2. $Z_i = 0, S_i = 0$: In this case C_i is either c or n , so we sample from the binomial distribution where the probability for class c is

$$P(C_i = c | Z_i = 0, S_i = 0, \mu, \sigma^2, \beta, \mathbf{Y}_i) = \frac{\pi_{ic} \varphi(Y_i | \mu_{c0}, \sigma_c^2)}{\pi_{ic} \varphi(Y_i | \mu_{c0}, \sigma_c^2) + \pi_{in} \varphi(Y_i | \mu_{n0}, \sigma_n^2)}$$

B.2 Full conditional for μ_{pz}

We assume a $N(0, B)$ prior distribution for each μ_{pz} for principal stratum p and treatment group z , so the full conditional distribution of μ_{pz} is Normal with mean

$$\mu_{pz}^* = \frac{B}{B + \sigma_p^2/n_{pz}} \bar{Y}_{pz}$$

and variance

$$\tau_{\star}^2 = \frac{B\sigma_p^2/n_{pz}}{B + \sigma_p^2/n_{pz}}$$

where

$$\bar{Y}_{pz} = \frac{1}{n_{pz}} \sum_{i=1}^n Y_i 1\{C_i=p, Z_i=z\}$$

$$\text{and } n_{pz} = \sum_{i=1}^n 1\{C_i=p, Z_i=z\}.$$

B.3 Full conditional for σ_p^2

We use a inverse gamma $IG(a, b)$ prior for the σ_p^2 parameter and so the full conditional is also inverse gamma with parameters $a^* = a + \frac{n_p}{2}$ and

$$b^* = b + \frac{1}{2} \left\{ \sum_{z \in \{0,1\}} \sum_{i: C_i=p} (Y_i - \mu_{pz})^2 \right\}$$

B.4 Full conditional for β

The prior distribution for the vector β_p is $\mathcal{N}(0, D)$ where D is a diagonal $J \times J$ matrix where J is the number of predictors included in the prediction model, including an intercept. Let X_i be the $J \times 1$ vector containing the predictor information and intercept column for subject i . Then the full conditional for β_p is proportional to

$$\left\{ \prod_{i=1}^n \pi_{ic}(\beta_c)^{1\{C_i=c\}} \pi_{ia}(\beta_a)^{1\{C_i=a\}} \pi_{in}(\beta_n)^{1\{C_i=n\}} \right\} \varphi(\beta_c | 0, D) \varphi(\beta_a | 0, D) \varphi(\beta_n | 0, D)$$

where

$$\pi_{ip}(\beta_p) = \frac{\exp(X_i' \beta_p)}{\sum_{j \in \{c,a,n\}} \exp(X_i' \beta_j)}$$

We sample from this full conditional distribution using a random walk Metropolis-Hastings rejection sampling algorithm.

References

- Ahluwalia SK, Matsui EC. The indoor environment and its effects on childhood asthma. *Curr Opin Allergy Clin Immunol*. 2011; 11:137–143. [PubMed: 21301330]
- Angrist JD, Imbens GW, Rubin DB. Identification of Causal Effects Using Instrumental Variables (Disc: P456-472). *Journal of the American Statistical Association*. 1996; 91:444–455.
- Bateson TF, Schwartz J. Children's response to air pollutants. *J. Toxicol. Environ. Health Part A*. 2008; 71:238–243. [PubMed: 18097949]
- Breyse PN, Diette GB, Matsui EC, Butz AM, Hansel NN, McCormack MC. Indoor Air Pollution and Asthma in Children. *Proceedings of the American Thoracic Society*. 2010; 7:102–106. [PubMed: 20427579]
- Brook RD, Rajagopalan S, Pope CA, Brook JR, Bhatnagar A, Diez-Roux AV, Holguin F, Hong Y, Luepker RV, Mittleman MA, Peters A, Siscovick D, Smith SC, Whitsel L, Kaufman JD. Particulate matter air pollution and cardiovascular disease: An update to the scientific statement from the American Heart Association. *Circulation*. 2010; 121:2331–2378. [PubMed: 20458016]
- Butz AM, Matsui EC, Breyse P, Curtin-Brosnan J, Eggleston P, Diette G, Williams D, Yuan J, Bernert JT, Rand C. A randomized trial of air cleaners and a health coach to improve indoor air quality for inner-city children with asthma and secondhand smoke exposure. *Arch Pediatr Adolesc Med*. 2011; 165:741–748. [PubMed: 21810636]
- Dominici F, Peng RD, Bell ML, Pham L, McDermott A, Zeger SL, Samet JM. Fine Particulate Air Pollution and Hospital Admission for Cardiovascular and Respiratory Diseases. *Journal of the American Medical Association*. 2006; 295:1127–1134. [PubMed: 16522832]
- Eggleston PA, Butz A, Rand C, Curtin-Brosnan J, Kanchanaraksa S, Swartz L, Breyse P, Buckley T, Diette G, Merriman B, Krishnan JA. Home environmental intervention in inner-city asthma: a randomized controlled clinical trial. *Ann. Allergy Asthma Immunol*. 2005; 95:518–524. [PubMed: 16400889]
- Frangakis CE, Rubin DB. Principal Stratification in Causal Inference. *Biometrics*. 2002; 58:21–29. [PubMed: 11890317]
- Gallop R, Small DS, Lin JY, Elliott MR, Joffe M, Ten Have TR. Mediation analysis with principal stratification. *Stat Med*. 2009; 28:1108–1130. [PubMed: 19184975]
- Hirano K, Imbens GW, Rubin DB, Zhou X-H. Assessing the effect of an influenza vaccine in an encouragement design. *Biostatistics*. 2000; 1:69–88. [PubMed: 12933526]
- Holland PW. Statistics and Causal Inference. *Journal of the American Statistical Association*. 1986; 81:945–960.
- Imai K, Keele L, Yamamoto T. Identification, inference, and sensitivity analysis for causal mediation effects. *Statistical Science*. 2010; 25:51–71.
- Imbens GW, Rubin DB. Bayesian inference for causal effects in randomized experiments with noncompliance. *Annals of Statistics*. 1997; 25:305–327.
- Jones GL, Haran M, Cao BS, Neath R. Fixed-Width Output Analysis for Markov Chain Monte Carlo. *Journal of the American Statistical Association*. 2006; 101:1537–1547.
- Koenig JQ, Mar TF, Allen RW, Jansen K, Lumley T, Sullivan JH, Trenga CA, Larson T, Liu LJ. Pulmonary effects of indoor- and outdoor-generated particles in children with asthma. *Environ. Health Perspect*. 2005; 113:499–503. [PubMed: 15811822]
- Little RJ, Rubin DB. Statistical analysis with missing data. 2002
- McCormack MC, Breyse PN, Matsui EC, Hansel NN, Peng RD, Curtin-Brosnan J, Williams DL, Wills-Karp M, Diette GB. Indoor particulate matter increases asthma morbidity in children with non-atopic and atopic asthma. *Ann. Allergy Asthma Immunol*. 2011; 106:308–315. [PubMed: 21457879]
- Mealli F, Mattei A. A Refreshing Account of Principal Stratification. *The International Journal of Biostatistics*. 2012; 8
- Morgan WJ, Crain EF, Gruchalla RS, O'Connor GT, Kattan M, Evans R, Stout J, Malindzak G, Smartt E, Plaut M, Walter M, Vaughn B, Mitchell H. Results of a home-based environmental intervention among urban children with asthma. *N. Engl. J. Med*. 2004; 351:1068–1080. [PubMed: 15356304]

- Pearl J. Principal stratification—a goal or a tool? *The International Journal of Biostatistics*. 2011; 7
- Pope CA, Dockery DW. Health effects of fine particulate air pollution: lines that connect. *Journal of the Air and Waste Management Association*. 2006; 56:709–742. [PubMed: 16805397]
- Robins JM, Greenland S. Identifiability and Exchangeability for Direct and Indirect Effects. *Epidemiology*. 1992; 3:143–155. [PubMed: 1576220]
- Rubin DB. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*. 2005; 100
- Schwartz SL, Li F, Mealli F. A Bayesian semiparametric approach to intermediate variables in causal inference. *Journal of the American Statistical Association*. 2011; 106:1331–1344.
- Strickland MJ, Darrow LA, Klein M, Flanders WD, Sarnat JA, Waller LA, Sarnat SE, Mulholland JA, Tolbert PE. Short-term associations between ambient air pollutants and pediatric asthma emergency department visits. *Am J Respir Crit Care Med*. 2010; 182:307–16. [PubMed: 20378732]
- Stuart EA, Cole SR, Bradshaw CP, Leaf PJ. The use of propensity scores to assess the generalizability of results from randomized trials. *Journal of the Royal Statistical Society, Series A*. 2011; 174:369–386.
- VanderWeele TJ. Principal Stratification—Uses and Limitations. *The International Journal of Biostatistics*. 2011; 7:1–14.
- Zigler CM, Dominici F, Wang Y. Estimating causal effects of air quality regulations using principal stratification for spatially correlated multivariate intermediate outcomes. *Biostatistics*. 2012; 13:289–302. [PubMed: 22267524]

Figure 1. Differences (followup minus baseline) in (a) the outcome and (b) the intermediate variable $PM_{2.5}$ by treatment group from the PREACH study.

Figure 2. Posterior mean (95% posterior interval) of change in symptom-free days from baseline to 6-month followup by principal stratum and intervention group in the PREACH study.

Figure 3. Comparison of the posterior (histogram) and prior (dashed line) distributions for θ_c , θ_a , and θ_n under the Full Bayesian Model.

Table 1

Definitions of principal strata in an air cleaner vs. control intervention study.

$S_i(0)$	$S_i(1)$	Description	Nickname
1	1	PM _{2.5} is always lowered (regardless of treatment group)	Always-reduced (A)
0	0	PM _{2.5} is never lowered (regardless of treatment group)	Never-reduced (N)
0	1	PM _{2.5} is lowered under air cleaner <i>only</i>	Cleaner-under-treatment (C)
1	0	PM _{2.5} is lowered under control <i>only</i>	Dirtier-under-treatment (D)

Table 2

Principal strata [cleaner-under-treatment (C), always-reduced (A), never-reduced (N), dirtier-under-treatment (D)] and their relationship to the observed treatment group (Z_i) and $\text{PM}_{2.5}$ level (S_i).

Z_i	S_i	Observed Result for Subject	Principal Stratum
1	1	$\text{PM}_{2.5}$ lowered with an air cleaner	(C) or (A)
1	0	$\text{PM}_{2.5}$ not lowered with an air cleaner	(N) or (D)
0	1	$\text{PM}_{2.5}$ lowered under control	(A) or (D)
0	0	$\text{PM}_{2.5}$ not lowered under control	(N) or (C)

Table 3

Posterior means (95% posterior intervals) of the association between treatment assignment and symptom-free days in the cleaner-under-treatment ($\hat{\theta}_c$), always-reduced ($\hat{\theta}_a$), and never-reduced ($\hat{\theta}_n$) principal strata, by model type.

	$\hat{\theta}_c$	$\hat{\theta}_a$	$\hat{\theta}_n$	Overall
Original Analysis				2.1 _(0.3,3.8)
Full Bayesian Model	6.2 _(2.7,11.6)	-0.2 _(-1.2,1.0)	5.0 _(0.3,10.2)	
Exclude Always-reduced	6.0 _(2.7,9.3)	0	5.0 _(0.4,10.1)	
Exclude Always- and Never-reduced	7.1 _(2.8,12.1)	0	0	