Environmental and Resource Economics
https://doi.org/10.1007/s10640-024-00946-y

Optimal Point Source Abatement Technology Adoption: The Impact of Uncertainty in the Benefits of Abatement

Andrew L. Goodkind[1] · Jay S. Coggins[2] · Christopher W. Tessum[3] · Julian D. Marshall[4]

Accepted: 25 November 2024

Abstract

Reducing emissions from point sources may be justified by the large expected benefits of improved health. However, the optimal reduction in emissions is complicated by the large uncertainty regarding the magnitude of these benefits. In particular, there is uncertainty in the size of the impact of pollution on increased premature mortality, and in the monetary valuation of reducing risks of mortality. We calculate the optimal emission reductions from abatement technology adoption at most point sources of SO_2, NO_X, and primary $PM_{2.5}$ in the United States across a wide range of uncertainty in the parameters used to estimate benefits of reductions. The results demonstrate that although the range of uncertainty in benefits is very wide, as long as the benefits are not at the low end of the distribution, the optimal abatement from sources is in a relatively narrow range. It is when benefits of reducing pollution are well below their mean estimates that the optimal reduction in emissions varies substantially. Resolving the likelihood of very low benefits of abatement could potentially reduce the uncertainty regarding optimal abatement policy.

Keywords Air pollution · Value of a statistical life · Concentration-response · Uncertainty · Environmental economics · Abatement costs

✉ Andrew L. Goodkind
agoodkind@unm.edu

Jay S. Coggins
jcoggins@umn.edu

Christopher W. Tessum
ctessum@illinois.edu

Julian D. Marshall
jdmarsh@uw.edu

[1] Department of Economics, University of New Mexico, Albuquerque, USA

[2] Department of Applied Economics, University of Minnesota, St. Paul, USA

[3] Department of Civil and Environmental Engineering, University of Illinois, Champaign, USA

[4] Department of Civil and Environmental Engineering, University of Washington, Seattle, USA

Published online: 22 January 2025

1 Introduction

Fine particulate matter air pollution ($PM_{2.5}$) is a major health threat around the world. Two features of the science in relation to $PM_{2.5}$ are well understood. One feature is that the health threat is more serious than was known even a few decades ago. A state-of-the-science study places the estimate at 8.9 million premature deaths around the world each year (Burnett et al. 2018). Another feature, closely related, is that reducing air pollution saves lives, even in locations that might appear to be clean, and where $PM_{2.5}$ concentrations are well below the United States (US) EPA's national ambient air-quality standards (Thakrar et al. 2020). Particulate pollution is decidedly not a problem only in the developing world, and therefore it continues to present important policy challenges in the US.

Because $PM_{2.5}$ is so dangerous, estimates show that reducing emissions of fine particles and their precursors confers large health benefits in lives saved. A number of sophisticated reduced-form tools have been developed in recent years to estimate the benefits of reductions and where they occur across the US (e.g., ISRM (Goodkind et al. 2019), AP3 (Muller 2014), EASIUR (Heo et al. 2016)). A common finding in this literature, confirmed in a number of EPA regulatory impact analyses (US EPA 2011, 2013, 2015), is that even as $PM_{2.5}$ concentrations drop, the economic benefits of further reductions remain high, in the billions of $US annually. For example, according to EPA estimates, the benefits associated with reductions of particle pollution that result from the Mercury and Air Toxins Standards exceed the costs by a factor that ranges from 3.4 to 9.4 (US EPA 2011). The lion's share of the monetary benefits flowing from further abatement are due to avoided premature mortality from exposure to $PM_{2.5}$. Shapiro and Walker (2020) find that for most regions in the US, marginal benefits of abatement exceed marginal abatement costs by more than tenfold.

An important policy question, then, is this: Just how low should we drive $PM_{2.5}$ levels to achieve the best outcome for society? In order to answer this question using a cost-benefit analysis framework, one needs to compare marginal benefits of further reductions to the marginal costs of further reductions, and one need to make this comparison in a way that is spatially sensitive. Another important policy question is how to make sound policy decisions in the face of remaining uncertainty around the various health and economic estimates associated with particulate pollution.

The purpose of the present paper is to address these two challenges. We conduct an empirical exercise that explicitly compares the cost and benefits of abatement technology adoption at most point sources of emissions in the US, accounting for the spatial distribution of health impacts from each source. Our model takes advantage of cost estimates from the EPA's Control Strategy Tool (CoST) model (US EPA 2016), which provides granular information on abatement and costs for a variety of control technologies at each source. Here our objective is to take one step in the direction of estimating the optimal level of $PM_{2.5}$ reductions, by which we mean the emissions levels, by source, at which the marginal cost of further abatement equals (or is greater than) the marginal benefits of further abatement.

A natural question arises in connection with much of the work in this area: How confident can one be in the leading estimates of health impacts, and of the economic valuation of those impacts? A close look at the main elements of the epidemiological and economic literature on which policy analyses are based suggests that the uncertainty around the most prominent benefits estimates is quite large. We introduce uncertainty into our model, with the aim of quantifying the range of possibly optimal policy outcomes. We focus on three main sources of uncertainty: (*i*) uncertainty in the concentration-response (C-R) functions that govern estimates of lives lost as a function of $PM_{2.5}$ concentrations; (*ii*) uncertainty in the economic estimates of the value of a statistical life (VSL); and (*iii*) the functional form of the C-R function, which relates to the curvature of that relationship.

The C-R functions we employ are from Krewski et al. (2009), the most recent reanalysis of the American Cancer Society (ACS) cohort study of air pollution and premature mortality. The Krewski study is attractive because it offers statistical estimates of the C-R in two functional forms, log-linear and log-log. It also has a lower estimate of the $PM_{2.5}$ health risk than does Lepeule et al. (2012), the other major study on which EPA estimates rely, which means our results lie at the conservative end of the range. Finally, the log-log results from Krewski closely resemble the C-R estimates from Burnett et al. (2018). Our approach to incorporating uncertainty in the C-R function is to exploit the published confidence intervals in Krewski et al. (2009).

Our treatment of uncertainty in the VSL is based upon the Weibull distribution that the EPA has fitted statistically to the sample of 26 studies that form the foundation of the agency's guidance on the use of VSL in policy analysis (US EPA 2010). Finally, although the comparison is not strictly statistical in nature, we also conduct our entire empirical exercise using both the log-linear and the log-log C-R functions from Krewski et al. (2009), incorporating their respective error structure.

The approach we have developed, in broad outlines, is to draw a set of values from the VSL and C-R distributions and then, for each of the C-R functional forms, to compute the optimal control strategies to be applied at each point source of emissions in the CoST dataset. A control strategy is optimal for a given source when the next most aggressive, or costly, strategy is no longer advantageous at the margin, with source-specific marginal benefits of abatement computed according to the InMAP Source-Receptor Matrix (ISRM) (Goodkind et al. 2019; Tessum et al. 2017). The resulting vector of optimal controls becomes a high-resolution spatial distribution of $PM_{2.5}$ precursor emissions, and resulting $PM_{2.5}$ concentrations across the continental US. The map of concentrations generates health impacts in the form of avoided premature mortality, which are converted to economic benefits by multiplying avoided statistical deaths by the relevant VSL draw.

Our results, summarized briefly, show that across the range of uncertainty in benefits of abatement, substantial emission reductions are optimal for point sources: ~24–42% reductions in SO_2, 31–43% reductions in NO_X, 30–47% reductions in primary $PM_{2.5}$.[1] When marginal benefits lie within the interquartile range of outcomes, the corresponding range of optimal reduction in damages from emissions is relatively narrow (32–46% reduction in total damages under the log-linear C-R function). The functional form of C-R makes a smaller difference in this case than one might expect as well, with a 44–53% reduction in total damages if log-log.

The gap between the optimal level of emissions reduction under the log-linear and log-log functional forms is small, according to our results, when the background concentrations are at their status quo levels. As background concentration are ratcheted down in our model, perhaps as a result of emissions reductions in other sectors, or of an increase in renewables penetration in electricity generation, the difference due to functional form grows substantially because the marginal benefits of a ton of abatement also grow substantially with the log-log C-R as $PM_{2.5}$ concentrations decline. This is one of the perplexing implications of a

[1] The ranges reflect the interquartile range of reductions from the uncertainty analysis.

log-log C-R function, explored by among others Pope et al. (2011), Marshall et al. (2015), and Pope et al. (2015). The results in this paper provide new insights into this question, regarding the role of uncertainty in the health and economic gradient between the log-linear and log-log specifications.

The CoST model is a useful tool for the analysis presented in this study. It provides a detailed set of control technologies, and their costs and abatement levels, for many point sources in the country, and for different pollutants at each source. That model, though, does not allow us to explore a number of questions that might also have a significant effect on pollution, health, and monetary benefits. In particular, the model covers point sources of emissions, and only the technologies that can be applied to the existing fleet of emission sources. As coal plants are closed, and then perhaps as renewable generation sources like wind and solar begin to cause natural gas plants to close as well, $PM_{2.5}$ concentrations in the country will drop even further. An extension of this study to account for the health benefits of a renewable transition in electricity generation might add an important element to the discussion of climate action.

2 Methods

In this paper we estimate the cost of emission reductions at most point sources of emissions in the US and compare that with the benefits of reducing those emissions. Evaluating many possible abatement technologies for each point source, we calculate the optimal abatement technology adoption given the benefits. We then apply this analysis across a wide range of uncertainty in the benefits of abatement to determine how the optimal abatement is influenced by the range of uncertainty in benefits.

2.1 Abatement Costs

To calculate the costs of abatement for point sources of emissions, we use the Control Strategy Tool (CoST) model. This model is used frequently by the EPA to evaluate the cost of meeting emission reduction or pollution concentration targets. The model includes an inventory of the available and existing controls on each unit. The model is designed to calculate the discounted cost of implementing a new emission control on a unit over the lifetime of the control and compare that with the quantity of emissions reduced. The key output we sought with this model is the total annualized cost, which is the sum of the annualized capital cost (discounted at a 3% rate) and the annual variable and fixed operating and maintenance costs. We compare the total annualized cost of the control with the emission reductions to obtain a cost per unit of emission reduction.

The CoST model is run from a baseline emission inventory which includes the current emissions and the existing controls. We used the 2017 National Emissions Inventory (NEI), the most recent inventory that is complete (US EPA 2021). We included 36,397 individual point sources from the inventory, and excluded many thousands of other point sources (those with the fewest emissions) that cumulatively accounted for 17% of emissions. In total, the included sources emitted 1.96 million tons of SO_2, 1.93 million tons of NO_X, and 0.33 million tons of primary $PM_{2.5}$ in 2017.

We needed to extract from the CoST model every possible control strategy for each unit. This information was necessary because for a particular unit the optimal control strategy depends on the benefits of the emission reductions; therefore, we need all possible controls to see which one is best depending on the value of those emission reductions. To obtain this information, we exploited a function of the CoST model that calculates the least cost set of controls to achieve a nation-wide emission reduction goal. The least cost function identifies the sources and the controls employed to meet the specified target by choosing those controls, among all the possibilities, that have the lowest cost per unit of emissions reduced. We iteratively run the model, increasing the emission reduction goal from 0 to 100% of existing emissions. With each iteration, we save the control strategy (i.e., which control is used at which source), the control cost, and the emission reduction.

Note that for a particular source, as the emission reduction goal gets higher, the least cost control may (and often does) change. For example, if one particular control at a source reduces 100 tons of SO_2 emissions for a cost of $1,000,000 (or $10,00 per ton), this may be part of a least cost strategy with a low goal. When the goal gets progressively higher, there may be another control strategy for this source which is more expensive per unit of emissions reduced but also reduces many more emissions. For example, another strategy may cost $10 million and reduce 500 tons (or $20,000 per ton). This control would potentially be part of the least cost strategy that requires reducing a larger share of emissions. And, importantly, this control would supersede the previous control (i.e., they would not work additively, rather, the second control would be used in place of the original control).

We ran the CoST model across all possible emission reductions (i.e., 0 to 100%) separately for three pollutants: SO_2, NO_X, and primary $PM_{2.5}$. These runs produced a dataset of all potentially least-cost control strategies for most point sources and for the three pollutants. For some of the controls, multiple pollutants were reduced—specifically, some controls reduced both SO_2 and primary $PM_{2.5}$. We calculate the combined benefits of a control from all pollutants that are reduced to compare with the costs, as we explain further down in our description of the model. Before we discuss how the different control technologies are chosen for the optimal abatement we first must discuss the benefits of abatement.

2.2 Benefits of Abatement

For the benefits of abatement, we focus solely on the reduction in adult premature mortality. There are other categories of benefits, but based on our understanding of the literature, this category comprises the overwhelming majority of the benefits (Muller and Mendelsohn 2007). Still, our results are excluding several other categories of benefits and should be considered a lower bound. The benefits in our model are from reductions in exposure to $PM_{2.5}$. The key damage from the emissions in our model (SO_2, NO_X and primary $PM_{2.5}$) are in their contribution to $PM_{2.5}$ concentrations that people are exposed to. Primary $PM_{2.5}$ directly contributes to $PM_{2.5}$ concentrations, whereas, SO_2 and NO_X are precursor pollutants that form into secondary $PM_{2.5}$ in the atmosphere. $PM_{2.5}$ concentrations are the combination of primary $PM_{2.5}$ and secondary $PM_{2.5}$.

The benefits we calculate are source specific—that is, we estimate the contribution to human exposure separately for each source of emissions. This is done using the InMAP Source-Receptor Matrix (ISRM) (Goodkind et al. 2019; Tessum et al. 2017). The ISRM isolates the impact of emissions for each source to provide the marginal benefits—or marginal change in the exposure to $PM_{2.5}$—from reductions of emissions at a particular source.

Then we focus on three key factors of uncertainty in the calculation of benefits: (*i*) the VSL, (*ii*) the C-R function between premature mortality and exposure to $PM_{2.5}$, and (*iii*) the functional form of the C-R function. Our goal is to understand how uncertainty in each of these factors impacts the optimal emission reductions at each source. Our application of the uncertainty is meant to demonstrate and represent the broad range from each facet, rather than be an exhaustive accounting and estimation of the uncertainty. Certainly, others could quibble with how we present the uncertainty, and we would not object, but our goal was to illustrate the uncertainty in a tractable way for the question we are looking to explore.

VSL

The frequently employed method of calculating the benefits of reductions in $PM_{2.5}$ is to multiply the change in premature mortality by a single number, the VSL, representing the value of those lives saved. There are several potential issues with this method (e.g., whether the value should be applied uniformly across all people, regardless of age or other factors) which we do not attempt to engage with. We only investigate the magnitude of the VSL and the uncertainty around this estimate.

There are several meta-analyses of the VSL (see Viscusi and Aldy (2003) and Kochi et al. (2006)) which tend to find relatively similar estimates, but we focus specifically on the estimate produced by the EPA (US EPA 2010), which is used to estimate the benefits of federal rules which impact $PM_{2.5}$ concentrations. The EPA VSL estimate is based on 26 revealed-preference and contingent-valuation studies. In their original analysis, a Weibull distribution was fit to the estimates of these studies. Generally, only the mean value of this distribution is used by the EPA. We make use of the entire distribution, which has a mean of \$9.6 million and an interval from the 2.5th to the 97.5th percentile of \$0.9 million to \$28.8 million. This enormous range of the VSL suggests that depending on which value is used from the distribution, it will dramatically alter the estimated benefits of abatement. Our goal here is to ask the question: How do different values in the range of uncertainty in benefits impact the optimal emission reduction? We draw values from the VSL distribution, recalculate the benefits of emission reductions, and then calculate the optimal emission reduction from each source.

C-R function

In addition to the uncertainty regarding how to value the lives saved from emission reductions, there is also the question of how many people are actually saved from reducing $PM_{2.5}$ concentrations. As with the VSL, there are several estimates of this relationship, including two key estimates that are used by the EPA in evaluating the benefits of concentration reductions. One, which we will focus on here, is from Krewski et al. (2009), which is the most recent results from a series of estimates based on a longitudinal study by the ACS. This study finds that for every 10 μgm^{-3} increase in $PM_{2.5}$ concentration, the premature mortality rate is 6% higher. The second, from the Harvard Six Cities (H6C) study, by Lepeule et al. (2012), finds much larger impacts of $PM_{2.5}$ than the ACS study (approximately double the impact—every 10 μgm^{-3} is associated with a 14% increase in mortality). The EPA employs both studies, using the mean results from each study and estimating the benefits separately to produce a high and low estimate.

We focus solely on Krewski et al. (2009), and use their estimated error to evaluate the range of uncertainty in C-R function. We believe this better illustrates the uncertainty compared with the EPA method, because their low range ignores the real possibility that impacts could be lower than the mean from Krewski et al. (2009). We do not combine the results from the two studies, leaving out the Lepeule et al. (2012) estimates entirely, which limits our estimate of the high-side uncertainty in impacts.

A comparison of the marginal benefits of reducing $PM_{2.5}$ using both the ACS and H6C estimates is shown in the middle panel of Fig. 1. Our use of Krewski et al. (2009) is based on three reasons. First, often the criticism of the benefits of reducing $PM_{2.5}$ is that the estimates are far too high, therefore, we feel it is prudent to use the conservative estimate of the C-R function. Second, a more recent meta-analysis of the C-R function using US and international data from Burnett et al. (2018) aligns closely in magnitude with Krewski et al. (2009) as shown in the right panel of Fig. 1.[2] Third, the Krewski et al. (2009) estimates presents the results with two different functional forms of the C-R function. The functional form introduces another area of uncertainty that we wish to examine. The Krewski et al. (2009) estimates provide a clean representation of the impact of the functional form.

C-R Functional Form

The two functional forms we examine are referred to as log-linear and log-log. The log-linear is the default form used in most of the analyses in the epidemiological literature. This functional form is nearly linear over the relevant range of $PM_{2.5}$ concentrations, and therefore, shows that a reduction in $PM_{2.5}$ concentrations will have similar benefits regardless of the initial concentration. The log-log C-R functional form, which was helpfully estimated by Krewski et al. (2009), is concave (or supralinear as it is referred to in this literature), and shows that the benefits of reducing $PM_{2.5}$ concentrations are larger when the initial concentration is low compared to with a high initial concentration—several other studies also found evidence of a supralinear C-R function between premature mortality and $PM_{2.5}$

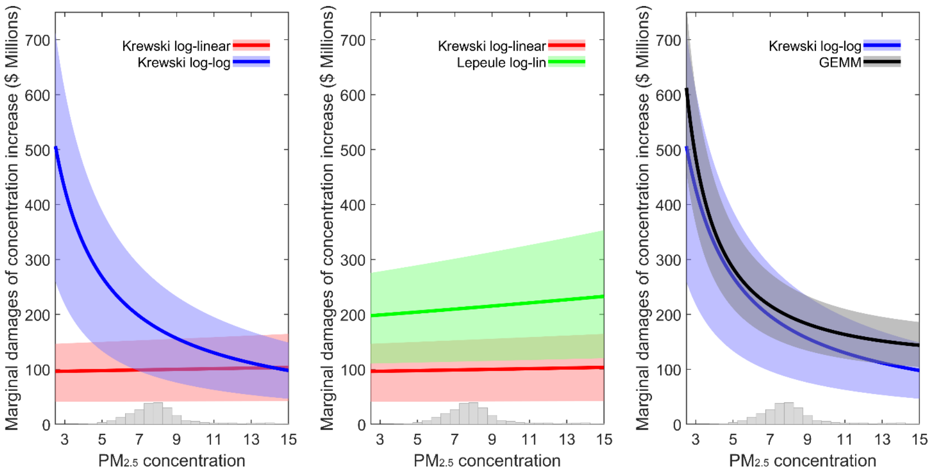

Fig. 1 Comparison of marginal benefits of reducing $PM_{2.5}$ concentrations in a hypothetical city of 250,000 people across different estimates of the C-R relationship. Marginal benefits are calculated using a VSL of \$8.6 million. Left panel, Krewski et al. (2009) log-linear versus log-log. Middle panel, Krewski et al. (2009) log-linear versus Lepeule et al. (2012) log-linear. Right panel, Krewski et al. (2009) log-log versus Burnett et al. (2018) which is based on a flexible functional form. The population-weighted $PM_{2.5}$ concentration distribution for the US is shown by the bars at the bottom of each panel

[2] However, the Burnett et al. (2018) estimates are partially influenced by Krewski et al. (2009)

concentrations (Pope et al. 2011; Crouse et al. 2012; Burnett et al. 2014, 2018; Vodonos et al. 2018, Miller et al. 2021). The concave log-log C-R function leads to an interesting increasing-returns-to-scale feature of pollution abatement, such that each additional unit reduction in concentration leads to a higher marginal benefit for additional abatement. The impact of the different C-R functional forms on policy showed that lower concentrations are optimal with log-log (Goodkind et al. 2014) and raises some ethical issues of inequality if policy was directed to cleaning up the already clean areas first (Marshall et al. 2015).

Here, we examine how the difference between the functional forms impacts optimal abatement technology adoption using real-world estimates of the costs of abatement. Figure 1, shows the marginal benefits of reducing $PM_{2.5}$ concentrations in a hypothetical city depending on the C-R function employed. In particular, this figure shows the benefits of a 1 μgm^{-3} reduction in $PM_{2.5}$ concentrations, with the dark lines showing the mean estimate, and the light area showing the error bands. The left panel of Fig. 1 compares the Krewski et al. (2009) log-linear and log-log C-R functions. The majority of the $PM_{2.5}$ concentrations in Krewski et al. (2009) are from locations with $PM_{2.5}$ levels above 10 μgm^{-3}, and the two estimates are similar in this range. The divergence is when concentrations get much lower. Here the benefits of cleaning up to very low concentrations are enormous if the C-R relationship is log-log, but are much more moderate if the C-R is log-linear.

The middle panel of Fig. 1 compares Krewski et al. (2009) log-linear with Lepeule et al. (2012) log-linear (i.e., the two C-R used in EPA analyses), showing that the shape is similar but the magnitude of the latter is much greater. The right panel of Fig. 1 compares Krewski et al. (2009) log-log with Burnett et al. (2018). These estimates are very similar in shape and magnitude. The Burnett et al. (2018) estimate is a state-of-the-science estimate of this relationship, especially for use in an international context, and we believe given the similarity, justifies the use of Krewski et al. (2009) log-log as a relevant measure of the impact of pollution at low concentrations.

Marginal benefit equation and index

So far, we have explained the general concepts of the uncertainty of benefits of abatement that we examine, and here we show specifically how we employ these estimates in our model. First, we show the marginal benefit equation, and then demonstrate how we incorporate uncertainty in the parameters of this equation. The marginal benefit in a location (or city/town) i, of a decrease in $PM_{2.5}$ concentration (C_i) is

$$MB_i^{\mathrm{lin}}(C_i) = V \cdot P_i \cdot \lambda_i^0 \cdot \gamma^{\mathrm{lin}} \cdot RR_i^{\mathrm{lin}}(C_i) \tag{1}$$

and

$$MB_i^{\mathrm{log}}(C_i) = V \cdot P_i \cdot \lambda_i^0 \cdot \gamma^{\mathrm{log}} \cdot \frac{RR_i^{\mathrm{log}}(C_i)}{C_i}, \tag{2}$$

with the log-linear (lin) and log-log (log) C-R functional forms. In Eqs. (1) and (2), V is the VSL, P_i is the population of location i, λ_i^0 is the baseline mortality rate in location i, RR_i^{lin} and RR_i^{log} are the relative risk of mortality equations for log-linear and log-log, respectively, of being exposed to a $PM_{2.5}$ concentration level, C_i, different from the baseline concentration, C_i^0. The γ parameters are from the relative risk equations:

$$RR_i^{\text{lin}}\left(C_i\right)=\exp\left\{\gamma^{\text{ lin}}\left(C_i-C_i^0\right)\right\}$$

and

$$RR_i^{\log}\left(C_i\right)=\left(\frac{C_i}{C_i^0}\right)^{\gamma^{\text{ log}}}.$$

Equations (1) and (2) show the marginal benefits of a concentration reduction in a location (e.g., these are the equations demonstrated in Fig. 1), but we need the marginal benefits of a reduction in emissions at a source. Define the increase in $PM_{2.5}$ concentrations at location i from emissions of pollutant p at source j as π_{ij}^{p}. These coefficients are estimated by the ISRM in Goodkind et al. (2019). Then the marginal benefits of reducing emissions of pollutant p at source j is just the sum of marginal benefits of a concentration reduction in all n downwind locations, weighted by the impact on concentrations in i from emissions at j:

$$MB_j^p=\sum\nolimits_{i=1}^{n}MB_i\left(C_i\right)\cdot\pi_{ij}^{p}. \tag{3}$$

Equation (3) can apply for either log-linear or log-log by plugging in Eq. (1) or (2), respectively.

With Eq. (3) we have our mean estimates of the marginal benefits of abating emissions from any point source for each of our three pollutants. These estimates come essentially directly from Goodkind et al. (2019), but we have adjusted the numbers to 2017 dollars and income levels (and used a different C-R function for log-log). These are our baseline marginal benefits, and then we adjust these numbers to account for specific draws from the VSL and C-R function distributions. We then calculate it separately for the two functional forms of the C-R we use. Changes to the VSL and C-R parameters in the marginal benefit calculation, almost (but not exactly), scale the values up and down proportionally. Also, we believe that the uncertainty in these two parameters are independent of each other, and therefore, the joint distribution of the uncertainty can be reflected by combining random draws from each distribution separately.

To illustrate this joint distribution, rather than show both dimensions, we combine it into one; and then scale it so that it is equal to one when both distributions are at their mean. We do this so that any draw from the joint distribution can be evaluated in how it scales up or down the marginal benefits. We call this the marginal benefit index (I_{MB}), which is simply the product of the ratio of each draw (k) to the mean value for each parameter:

$$I_{MB,k}=\left(\frac{RR_k}{\overline{RR}}\right)\cdot\left(\frac{VSL_k}{\overline{VSL}}\right)$$

where RR_k is the kth draw of the C-R distribution (relative-risk function), and $\overline{RR}$ is the mean; and VSL_k is the kth draw of the VSL distribution, and $\overline{VSL}$ is the mean. At the mean values of each distribution I_{MB} is equal to 1. For values below 1, marginal benefits are lower than the mean, and for values above 1, marginal benefits are above the mean. The marginal benefit index distribution is shown in Table 1. The table shows, for instance, that at the first percentile of the joint distribution of marginal benefits accounting for uncertainty

Table 1 Marginal benefits index distribution for log-linear and log-log functional forms

Percentile	I_{MB} log-linear	I_{MB} log-log
1	0.061	0.062
5	0.137	0.159
10	0.206	0.275
25	0.437	0.461
Median	0.792	0.784
75	1.313	1.264
90	1.916	1.942
95	2.402	2.337
99	4.827	5.063

in the log-linear C-R point estimate and the VSL, marginal benefits are 6% of the mean marginal benefit calculation. On the other extreme, the 99th percentile from this distribution shows that marginal benefits are 483% of the mean marginal benefit calculation. Put another way, the marginal benefits of abatement are 80-fold greater at the 99th percentile than the 1st percentile. It is perhaps not surprising that the range of uncertainty in marginal benefits is large, but estimates of the impacts of air pollution on mortality are rarely presented with an acknowledgement of the magnitude of the uncertainty.

Characterizing the uncertainty in marginal benefits, while necessary, is not the main focus of this analysis. We are interested in how the uncertainty in marginal benefits impacts optimal levels of abatement of emissions. To answer this question, we take a draw from the marginal benefit index joint distribution and apply this value to the mean marginal benefit estimate at each source—Equation (3), with the appropriate form of the C-R function. This provides the benefits, for a source, of their emission reductions. We then calculate the optimal emissions based on the available abatement technologies and their associated costs, described in detail in the section below.

This is a cumbersome process, so we needed to make decisions about the limited number of values to draw from the joint marginal benefit index. We decided on using a 3×3 matrix of values, using three values of the VSL distribution and three values of the C-R distribution, and every combination thereof. With this setup we can evaluate the impact on optimal abatement given low, medium, and high values from each distribution (VSL and C-R), rather than just knowing the value of the opaquer marginal-benefit index. The actual values chosen are the 10th percentile, mean, and 90th percentile of the VSL distribution; and the 2.5th percentile, mean, and 97.5th percentile of the C-R distribution. The resulting matrices (one for the log-linear and one for the log-log C-R) of the marginal benefit index are in Table 2. The specific percentiles chosen from each distribution are somewhat arbitrary, but the choice is not especially important—as is demonstrated in the results—because we can fit a curve through these nine points and provide a close approximation for any value of the marginal benefit index.

2.3 Calculating Optimal Abatement for Each Source

Now that we have each piece of the model, we have to put it together to find the optimal emission reductions at each point source. Each source is evaluated individually, and for each of the nine scenarios in the 3×3 matrix in Table 2—and done separately for the log-linear and log-log C-R functions. For a given scenario (value from the matrix) and source, we have

Table 2 3 × 3 matrix of marginal-benefit index for log-linear (top) and log-log (bottom) C-R functions. Columns are the relative-risk for a 10 μgm^{-3} change in $PM_{2.5}$. Rows are the VSL in millions

Log-linear C-R		2.5th=1.024	Relative risk Mean=1.060	97.5th=1.097
	10th=2.270	0.096	0.236	0.375
VSL	Mean=9.618	0.407	1.000	1.589
	90th=17.508	0.741	1.820	2.892

Log-log C-R		2.5th=1.044	Relative risk Mean=1.060	97.5th=1.148
	10th=2.270	0.112	0.236	.0359
VSL	Mean=9.618	0.474	1.000	1.521
	90th=17.508	0.864	1.820	2.768

the benefits of emission reductions, and all possible controls from the CoST model, and the reduction in emissions and the discounted annualized cost of each control.

Each control is compared with each other and the status quo. In the status quo, there are no benefits or costs, therefore, zero net benefits. If no available control has positive net benefits, then the model chooses the status quo, and there is no abatement. For each control, we calculate the benefits of abatement, which is simply the marginal benefits per ton of abatement times the tons of emission reduced.[3] For some controls, both SO_2 and primary $PM_{2.5}$ emissions are reduced, so the model calculates the benefits of each and adds them together. For controls of NO_X, there are no other pollutants reduced. Therefore, the model may choose to implement a SO_2/primary $PM_{2.5}$ control, or not; and may also choose to implement a NO_X control, or not. Note, that some controls reduce primary $PM_{2.5}$ only or SO_2 only, they are not necessarily both reduced, but we only allow the model to choose one control that reduces SO_2, primary $PM_{2.5}$, or both pollutants.

The model identifies the chosen control(s) that produce the greatest net benefits, calculates the emission reductions of each pollutant, the resulting benefits, and the annualized cost of the controls. By choosing the control for each pollutant that has the greatest net benefits, the model is picking the control at which the marginal benefits of abatement intersect with the marginal abatement cost function. Because of the discrete nature of the abatement controls, this intersection occurs where the marginal benefits per ton of abatement are equal to or greater than the marginal cost per ton of the abatement technology (this is illustrated in Fig. 2 below).

As mentioned above, the marginal benefit curves for any specific source are almost exactly constant across all quantities of abatement, but the curvature of the log-log C-R function greatly impacts the marginal benefit curve when the aggregate $PM_{2.5}$ concentrations change—with log-log the effect is that the marginal benefit curve shifts up substantially when $PM_{2.5}$ concentrations are lower; with log-linear the effect is that the marginal

[3] The marginal benefit curve for any particular source is very nearly flat because the impact on $PM_{2.5}$ concentrations from any particular source on any particular location is relatively small. Therefore, any curvature in the C-R function is not relevant for a specific source. Our assumption of a constant marginal benefit curve is a helpful simplification and very marginally impacts the results. The curvature in the C-R function is important when the aggregate emissions and therefore concentrations of $PM_{2.5}$ change (especially with log-log) and we flesh out that effect separately.

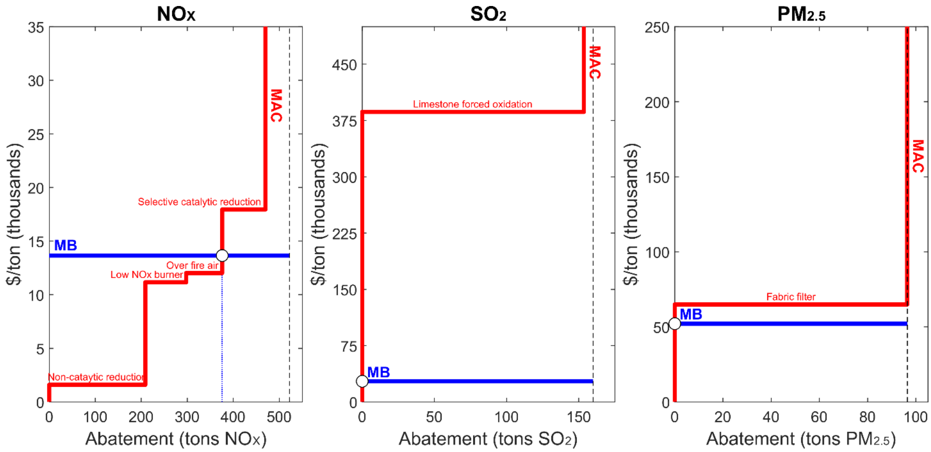

Fig. 2 Marginal abatement cost and marginal benefit functions from unit 1 from ALCOA Power Plant, Warrick, IN. Step functions, in red, are marginal abatement cost functions for each pollutant with the name of each abatement option listed on each step. The blue lines show the marginal benefits of abatement from this unit, at the mean marginal benefits ($I_{MB} = 1$)

benefit curve shifts down very little when $PM_{2.5}$ concentrations are lower. To evaluate the impact of this curvature we ask: What is the optimal abatement if emissions from all sources (not just point sources) were reduced by some amount? In addition to the baseline, we evaluate a scenario in which $PM_{2.5}$ concentrations everywhere are 50% of the baseline, and see how the optimal abatement technology adoption at point sources change.

3 Results

We start by illustrating the results from one particular point source of emissions: Unit 1 from the ALCOA Power Plant in Warrick, IN. This source is not a particularly large emitter, but is chosen to represent how the model functions. This unit emitted 522 tons of NO_X, 160 tons of SO_2 and 97 tons of primary $PM_{2.5}$ in 2017. Figure 2 shows the marginal abatement cost function in red for this unit, for each of the three pollutants. We see the stepwise feature of these functions, in particular for NO_X, in which there are progressively more costly abatement options which provides increasingly greater abatement. The blue line shows the marginal benefits of abatement from this unit. These are the mean marginal benefits (i.e., when the marginal benefits index is equal to one). At this point on the marginal benefit index, the marginal benefit of a unit of reduction of NO_X emissions is \$13,650—the marginal benefits for SO_2 and primary $PM_{2.5}$ reductions are \$27,500 and \$52,100,respectively.

The intersection of the marginal benefits and the marginal abatement costs leads to the level of abatement chosen by the model. In this example, a substantial share of the NO_X emissions is abated (72% or 376 tons), but neither of the abatement options for SO_2 nor primary $PM_{2.5}$ are selected because they are too costly. A different draw from the marginal benefit index would shift up or down the marginal benefits, and can lead to different intersections, and therefore, different levels of abatement.

3.1 Optimal Abatement Technology Adoption

The same process as in the above example is carried out for all point sources of emissions in the model. A couple of general themes emerge from our results: Under many circumstances, substantial reductions in emissions from point sources are optimal, and the largest variation in the quantity of abatement of emissions is when marginal benefits are drawn from the low end of the marginal benefit index.

We start by showing a comparison across our nine scenarios (each draw from the 3×3 matrix) of the share of initial damages that remain after optimal emission reductions of all pollutants from all point sources. Note that the remaining damages is just a restatement of the benefits of abatement; however, these are presented as the share of the initial damages remaining, such that a smaller number represents greater benefits.[4] These are illustrated in Fig. 3 with log-linear on the left and log-log on the right. The horizontal axis in Fig. 3 shows the marginal benefit index, and the box and whisker plot inside the figure shows the distribution of that index (i.e., the distribution described in Table 1): The mean (green line) is at 1; the median (black line) is lower; the box shows the interquartile range (IQR), and the whiskers extend to the 5th and 95th percentiles. The nine scenarios are the white dots, and a best fit line shows an approximation of the optimal remaining damages for any possible

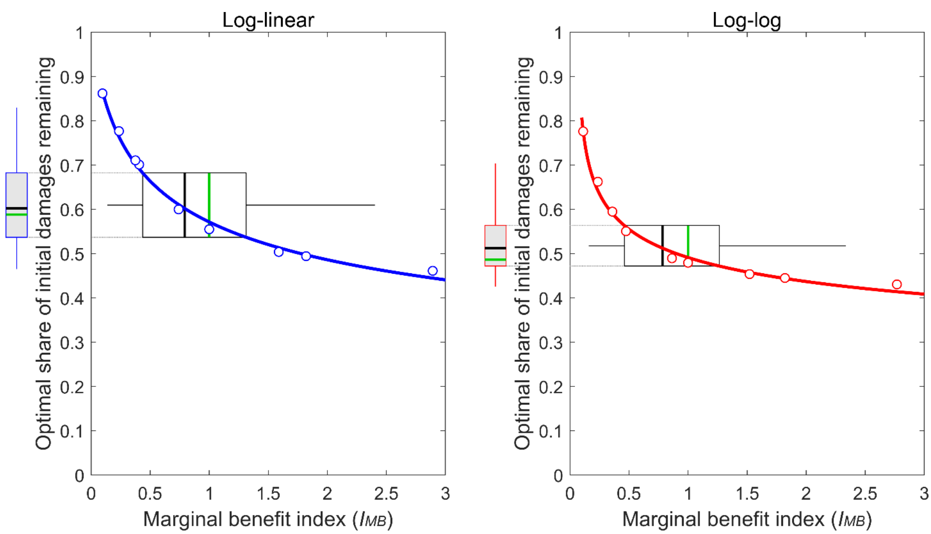

Fig. 3 Share of initial damages remaining after optimal control technologies implemented for each of nine scenarios (white dots). Krewski log-linear (left) and Krewski log-log (right) functional forms for C-R. The fitted lines demonstrate a close approximation of the optimal damages for any value of the marginal benefit index. The box and whisker plot, inside the graph, shows the distribution of the marginal benefit index across draws. The box and whisker plot, outside the axes, shows the distribution of optimal damages corresponding to the marginal benefit index. The green line is the mean, the black line is the median, the box represents the interquartile range, and the whiskers extend to the 5th and 95th percentiles

[4] We illustrate our results as the share of initial damages remaining, rather than the actual level, because at different values along the marginal benefit index, the initial damages are also different, and this way we provide a comparable baseline. Additionally, we use damages instead of emissions of each pollutant to see the broadest picture of the change in air pollution impacts from a single metric.

value of the marginal benefit index. The close fit to the data of the best-fit line suggests that the decision regarding which values to include in the 3×3 matrix is not pivotal because a close approximation for any value of the marginal benefit index is obtainable.

The figure demonstrates that across the very wide range of marginal benefits the optimal abatement—as measured by share of remaining damages—varies substantially. However, over the IQR of the marginal benefit index, the range of optimal abatement is relatively narrow. The IQR of the marginal benefit index for log-linear extends from 0.44 to 1.31—approximately a three-fold difference—yet the range of optimal abatement goes from a 32% reduction in damages to a 46% reduction (as shown in the box to the left of each axis). With log-log, the optimal abatement in the IQR of the marginal benefit index is from 44% reduction in damages to a 53% reduction.

The results from this model also suggest that if marginal damages are on the high end of the distribution, optimal abatement does not substantially increase beyond the 75th percentile results, illustrated by the flatter portion of the best-fit line. Even at the 95th percentile of the marginal benefit index, optimal damage reduction is 53% (for log-linear) and 57% (for log-log). At the low end of the marginal benefit index the optimal abatement is much lower, illustrated by the steeper portion of the best-fit line. At the 5th percentile, the optimal abatement is only 17% (for log-linear) and 30% (for log-log). The box-and-whisker plots outside and left of each vertical axis demonstrates these aspects, showing the distribution of optimal reduction in damages that maps from the marginal benefit index.

The short length of the bottom whisker and narrow range of the IQR of these distributions shows that over a wide range of the marginal benefit index, optimal abatement is in a narrow range. Thus, the interesting uncertainty, in terms of having a large impact on optimal abatement technology adoption, is on the low end of the marginal benefit index. The reason for this is because the marginal abatement cost curves become very steep at high levels of abatement, making further abatement a net negative for almost any higher value of the marginal benefits of abatement. At low levels of abatement, the marginal abatement cost curves are relatively flat, meaning that small changes in the marginal benefits can lead to large changes in the optimal amount of abatement. Our model shows that over a wide range of the marginal benefit index, most of the lowest-cost abatement options are a net positive. This suggests that devoting resources to resolving the likelihood of the low-end marginal benefits of reducing $PM_{2.5}$ air pollution could achieve the biggest payoff from reduced uncertainty regarding optimal policy.

Table 3 shows the results across the nine scenarios in the 3×3 matrix, where the values are the share of initial damages remaining. This helps compare the importance of the uncertainty in the C-R versus the VSL, in determining the optimal reduction in damages. Notice that for a relatively small amount of abatement to be optimal, it is almost required that the VSL be low. Even at the low end of the C-R function (2.5th percentile) the optimal share of initial damages is 70% for log-linear and 55% for log-log at the mean VSL. To get one-quarter or less of initial damages reduced, both the VSL and C-R must be towards the low end of their respective distributions. Because the distributions of the VSL and C-R functions are independent, the probability of draws from both distributions being towards the low extreme is very small—for example, the probability of draws at or below the 10th percentile of the VSL and at or below the 2.5th percentile of the C-R function is 0.25%.

Figure 4 is similar to Fig. 3, but illustrates the optimal abatement across the three pollutants. In each panel, the blue dots and best-fit line represent the log-linear C-R function, and

Table 3 Share of initial damages remaining with optimal abatement for 3 × 3 matrix of marginal benefit index for log-linear (top) and log-log (bottom) C-R functions

		Log-linear C-R		
		2.5th	Mean	97.5th
VSL	10th	0.86	0.78	0.71
	Mean	0.70	0.56	0.50
	90th	0.60	0.49	0.46

		Log-log C-R		
		2.5th	Mean	97.5th
VSL	10th	0.78	0.66	0.60
	Mean	0.55	0.48	0.45
	90th	0.49	0.45	0.43

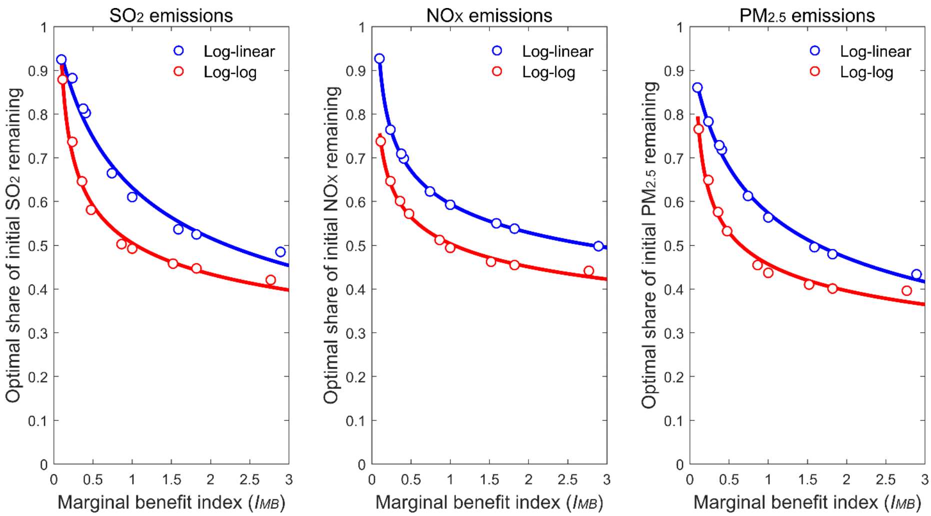

Fig. 4 Share of initial emissions under optimal control technology across nine scenarios, for SO_2, NO_X and $PM_{2.5}$ emissions. Red and blue lines are fitted to the data points (white circles)

the red dots and best-fit line represent the log-log C-R function. Reductions are similar for SO_2, NO_X, and primary $PM_{2.5}$. In each case, the reduction is greater with log-log compared with log-linear, but the difference is relatively small. The initial quantity of SO_2 emissions (1.96 million tons) are greatest, by mass, compared with NO_X (1.93 million tons) and primary $PM_{2.5}$ (0.33 million tons). Similarly, the initial damages of SO_2 make up a larger share of total damages (51%), than NO_X (21%) and primary $PM_{2.5}$ (28%). In the scenario with both distributions at the mean (i.e., the marginal-benefit index equals 1) with log-linear, the optimal damages are 56% of the initial damages, with 54% of the remaining damages from SO_2 emissions, 22% from NO_X, and 24% from primary $PM_{2.5}$.

Abatement costs to achieve the optimal emission reductions increases sharply as we move from the low-end to the high-end of the marginal benefit index. With log-linear, in the scenario with the lowest value of the marginal benefit index—the 2.5th percentile of C-R and the 10th percentile of VSL—the abatement costs are \$0.5 billion, or \$0.03 billion per percentage point reduction in initial damages. At the mean values of each distribution, the abatement costs increase to \$12.7 billion, or \$0.29 billion per percentage point reduction in damages. Finally, in the scenario with the highest marginal benefit index—the 97.5th percentile of C-R and the 90th percentile of VSL—the abatement costs are \$26.8 billion, or \$0.50 billion per percentage point reduction in damages. Plotting the abatement costs and percentage point decrease in initial damages, we construct the abatement cost and marginal abatement cost functions. The equation for the aggregate marginal abatement cost (MAC) function is

$$MAC\,(x) = 20{,}885{,}480e^{0.0934x}$$

where x is the percentage point decrease in initial damages.

The optimal emission reductions are not uniformly distributed across the US. We find that the sources with the largest initial damages, and the sources that, on average, reduce their initial damages the most, are concentrated in the Midwest/Rust Belt, with some other key reductions coming from the south and southeast regions. There are some sources in the west that are large emitters, but the impact of those emissions is relatively small, and they tend to not reduce their emissions/damages appreciably in the model.

Of the point sources included in the model, the largest sector of initial damages is from the combustion of coal—for electricity generation and industrial boilers—representing 51% of total initial damages. These sources reduce their damages (in the scenario in which the marginal benefit parameters are at their means) by 63%. Emissions from industrial processes, that make up one-third of initial damages, reduce their damages in this scenario by 45%. All other sources reduce their initial damages by 33%.

3.2 Lower Background $PM_{2.5}$ Concentrations

Up to this point, our results show that optimal damages are consistently similar but lower with the log-log C-R function compared with log-linear. Next, we show how these results change when background $PM_{2.5}$ concentrations are lower than the baseline. As was mentioned earlier, when $PM_{2.5}$ concentrations decrease, the log-linear marginal benefit function shifts down very slightly, but the log-log marginal benefit function shifts up appreciably. Mechanically, for log-log this is explained by Eq. (2), such that the marginal benefits are divided by the $PM_{2.5}$ concentration—marginal benefits increase with reductions in concentration. The same is not true of the log-linear marginal benefits in Eq. (1). With log-log, for a given source, when concentrations are decreased in all downwind locations, the benefit of reducing each unit of emissions is greater, thus causing the source's marginal benefit function to shift up. When concentrations decline, and the marginal benefits shift up for each source, this leads some sources to abate more emissions, and other sources to start abating emissions.

Figure 5 illustrates these effects, where we reduce background concentrations from 100% of the baseline, to 50% of the baseline. With both scenarios we hold constant the initial emissions from any given source, and then calculate the optimal emission reductions, and resulting share of initial damages remaining. In Fig. 5, the solid lines (and dots) are from

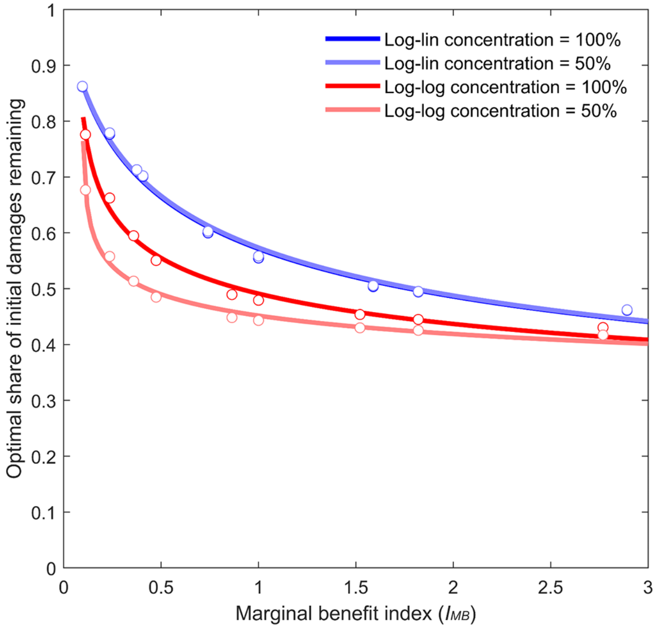

Fig. 5 Change in optimal damages given different background $PM_{2.5}$ concentrations for log-linear (blue) and log-log (red). The lines are fitted to the data points (white circles). The solid lines and data points represent the results when background concentrations are at baseline levels (100%). The lighter lines and data points represent the results when background concentrations are at 50% of the baseline

the baseline concentrations, and the lighter lines (and dots) show the optimal damages with lower background concentration.

The lines for log-linear (blue) are nearly overlapping, showing that marginal benefits stay nearly constant regardless of the background $PM_{2.5}$ concentration. The optimal damages for log-log (red), on the other hand, get smaller with lower concentrations. The effect is quite small at large values of the marginal benefit index, showing that we are reaching the limit of the available controls for reducing emissions. At the low end of the marginal benefit index, the drop in remaining damages across background concentration is substantial. For the 2nd lowest scenario of the marginal benefit index (I_{MB} = 0.236), the optimal damage reduction for log-log is 34% with 100% background concentrations, and this reduction increases to 44% when the background concentrations are 50%. This translates to a much larger gap between log-linear and log-log when the background concentration is lower. In the above example, the gap with 100% background concentrations is 11 percentage points (i.e., log-log decreases initial damages 11 percentage points more than log-linear), but this gap increases to 22 percentage points when background concentrations are at 50%. The key takeaway is that functional form matters much more when background concentrations are substantially lower than current concentrations and at the low end of the marginal benefit index.

In addition, when the background concentrations are lower (50% of current levels), the range of optimal abatement with log-log is narrow. For example, when I_{MB} = 0.236 this leads to a 44% reduction in damages, and when I_{MB} = 2.77 this leads to a 58% reduction in damages. In other words, a ~ 12-fold increase in marginal benefits (2.77 divided by 0.236) only increases damage reductions by 14 percentage points (58% minus 44%).

3.3 Plant Shutdown and Optimal Abatement

An important feature of pollution management, heretofore not included in the model, is the potential for the complete shutdown of a pollution source. In these situations, the socially optimal situation is for the facility to cease operation and eliminate all emissions. This can occur when the damages from the source are large (or the benefits of abatement are large), and the available abatement technologies are too expensive for the facility to adopt. This could potentially allow for substantially greater levels of abatement, especially when the marginal benefits of abatement are large. The information on the costs of such shutdowns, at each of the over 36,000 point sources modeled, is not available in the CoST model, and so cannot be directly simulated. Instead, we run a sensitivity analysis in which all emissions may be eliminated at a fixed dollar-per-ton shutdown cost. Then we vary this cost to see how this parameter affects optimal abatement decisions. The shutdown condition is included in the model in additional to the suite of available abatement technologies. The model then chooses either the status quo, an abatement technology, or shutdown, depending on which option provides the greatest social net benefits.

Given the diversity of sources in the model, the shutdown costs per ton of emissions reduced could vary enormously between facilities. We decided to choose values that represent the potential cost of providing emissions-free alternative electricity generating units to replace coal and natural gas power plants. Assuming a range between \$35 and \$60 per megawatt-hour (MWh) for renewable energy (EIA 2022), the cost of emissions reduced from coal power plants could be as low as \$30,000 per ton; and the cost of emissions reduction from natural gas power plants could be as high as \$200,000 per ton.[5] For facilities with many emissions, the shutdown costs per ton of emissions could be relatively low given the potentially large benefits of eliminating all of the harmful emissions. On the other hand, facilities with few emissions, the shutdown costs may be very high, because there are relatively limited health benefits of removing these emissions.

We run our model with three possible shutdown costs, using \$50,000, \$100,000, and \$150,000 per ton of emissions reduced. The results are presented in Fig. 6. The red line and circles represent the results presented earlier in Fig. 3 for the log-log C-R function. The dotted lines and corresponding circles represent optimal damage reductions when shutdown is available, again using the log-log C-R function. The general shape of the functions with shutdown is similar to without, but the drop off in damages (or increase in abatement) is substantial. Providing the option to shut down production becomes increasing attractive for the range when marginal benefits of abatement are approximately half their mean values ($I_{MB} = 0.5$) to 50% above their mean values ($I_{MB} = 1.5$). Over this range, the pollution damages drop substantially below the base-case levels, across all three of the shutdown cost scenarios.

While Fig. 6 can only provide a general approximation of how optimal abatement would change with the option to shut down production, it does illustrate the importance of understanding facility closures, and provides an avenue for future research.

[5] We base these calculations using the average emission rates of SO_2 and NO_X per MWh of electricity produced from coal and natural gas power plants. Coal power plants release, on average, approximately 2 pounds of SO_2 and 1 pound of NO_X per MWh of electricity (EPA 2020). Natural gas power plants release, on average, approximately 0.5 pounds of NO_X per MWh of electricity (EPA 2020).

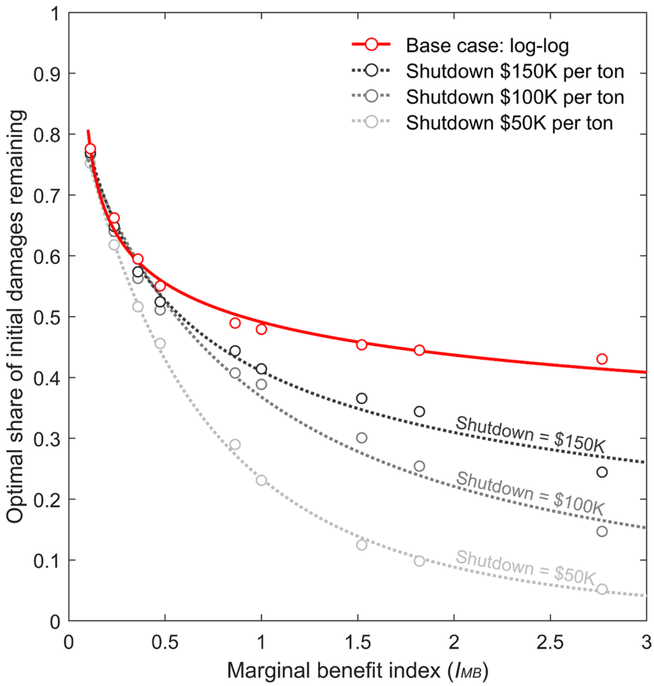

Fig. 6 Share of initial damages remaining after optimal control technologies or shutdown implemented for each of nine scenarios (white dots) and best-fit lines. The red line and circles represent the base case with no shutdown possible. The darker dotted line and circles represents abatement with the option to shutdown operators and eliminate all emissions for $150,000 per ton. The lighter dotted lines and circles represent abatement when shutdown costs are $100,000 and $50,000 per ton

4 Conclusion

The stakes in achieving the optimal level of air pollution across the landscape are high. Even at the relatively low concentrations observed in the US, ambient pollution of fine particles remains a significant health threat. Analysis of policies to improve air quality is bedeviled by a number of analytical challenges, including both the correct shape of the C-R relationship and the role of uncertainty around the main components of the problem.

We provided what we hope is a step forward on both of these dimensions. Using the CoST model, we compared the cost and benefits of abatement technology adoption at most point sources of air pollution in the US, accounting for the spatial distribution of health impacts from each source. Our results provide insights into the optimal level of pollution reduction, balancing the marginal cost of further reductions at each source against the marginal benefits of that reduction. Uncertainty concerning the C-R function as well as the value of a statistical life shows that the range of optimal reduction can be quite large. The shape of the C-R function, whether log-linear or log-log, has its own large impact on the results, with the divergence growing as air quality improves. The payoff to further refinements of our understanding of both uncertainty and functional form would appear to be quite high.

Our analysis is limited by a few factors. First, we rely on the CoST model for our estimates of abatement costs; however, Shapiro and Walker (2020) showed that these engineering costs may deviate substantially from marginal abatement costs inferred from pollution markets. In particular, Shapiro and Walker found the CoST model tends to underestimate costs of NOx abatement and overestimate costs of VOC abatement. In addition, the CoST model employed here cannot incorporate other important factors that impact emissions and abatement technology adoption. For instance, natural gas prices may influence facilities' fuel input decisions, leading to more or less coal and oil consumption and altering the emission profile for a source. These types of production decisions are beyond the abatement technology adoption decisions we can model. For the US-wide source-specific estimates of abatement costs that we sought for this analysis—to compare with the source-specific marginal benefits—the CoST model was an appropriate choice.

Second, the scale of the emission reductions in many of the scenarios outlined in the paper could have substantial impacts on local and regional air quality. This could potentially lead to population sorting towards areas with improved air quality, and have general equilibrium effects that are not accounted for here.

Third, we do not include the possibility of plants shutting down in our main analysis. To evaluate the actual costs of this option requires knowing the alternative production (and emissions) that would be adopted to replace each source. Given the scale of the sources evaluated here, this was not a feasible option. However, our sensitivity analysis with a relatively crude shutdown option illustrates the importance of incorporating this decision into pollution abatement models. Further research along the path we have laid down could profitably address the effects of a move away from fossil energy production. A renewable, electrified energy system for transportation and buildings and industry would likely drive a much larger reduction in particulate concentrations that we have considered here. That is a fruitful area for further research.

Acknowledgements This publication was developed as part of the Center for Air, Climate, and Energy Solutions (CACES), which was supported under Assistance Agreement No. R835873 awarded by the US EPA. It has not been formally reviewed by EPA. The views expressed in this document are solely those of the authors and do not necessarily reflect those of the Agency. EPA does not endorse any products or commercial services mentioned in this publication. The authors declare no financial or non-financial conflicts of interest.

References

Burnett R, Pope CA III, Ezzati M et al (2014) An integrated risk function for estimating the global burden of disease attributable to ambient fine particulate matter exposure. Environ Health Persp 122(4):397–403. https://doi.org/10.1289/ehp.1307049

Burnett R, Chen H, Syzszkowicz M et al (2018) Global estimates of mortality associated with long-term exposure to outdoor fine particulate matter. P Natl Acad Sci 115(38):9592–9597. https://doi.org/10.1073/pnas.1803222115

Crouse DL, Peters PA, van Donkelaar A et al (2012) Risk of nonaccidental and cardiovascular mortality in relation to long-term exposure to low concentrations of fine particulate matter: a Canadian national-level cohort study. Environ Health Persp 120(5):708–714. https://doi.org/10.1289/ehp.1104049

Environmental Protection Agency US (2021) 2017 National emissions Inventory: January 2021 Update Release, Technical Support Document. Office Air Qual Plann Stand Res Triangle Park EPA-454/R-21-001. https://www.epa.gov/air-emissions-inventories/national-emissions-inventory-nei

Goodkind AL, Coggins JS, Marshall JD (2014) A spatial model of air pollution: the impact of the concentration-response function. J Assoc Environ Reso 1(4):451–479. https://doi.org/10.1086/678985

Goodkind AL, Tessum CW, Coggins JS, Hill JD, Marshall JD (2019) Fine-scale damage estimates of particulate matter air pollution reveal opportunities for location-specific mitigation of emissions. P Natl Acad Sci 116(18):8775–8780. https://doi.org/10.1073/pnas.1816102116

Heo J, Adams PJ, Gao HO (2016) Public health costs of primary $PM_{2.5}$ and inorganic $PM_{2.5}$ precursor emissions in the United States. Environ Sci Technol 50(11):6061–6070. https://doi.org/10.1021/acs.est.5b06125

Kochi I, Hubbell B, Kramer R (2006) An empirical Bayes approach to combining and comparing estimates of the value of a statistical life for environmental policy analysis. Environ Resour Econ 34(3):385–406. https://doi.org/10.1007/s10640-006-9000-8

Krewski D, Jerrett M, Burnett RT et al (2009) Extended Follow-Up and Spatial Analysis of the American Cancer Society Study Linking Particulate Air Pollution and Mortality. Health Effects Institute, Boston 140:5–114

Lepeule J, Laden F, Dockery D, Schwartz J (2012) Chronic exposure to fine particulates and mortality: an extended follow-up of the Harvard Six Cities Study from 1974 to 2009. Environ Health Persp 120(7):965–970. https://doi.org/10.1289/ehp.1104660

Marshall JD, Apte JS, Coggins JS, Goodkind AL (2015) Blue skies bluer? Environ Sci Technol 49(24):13929–13936. https://doi.org/10.1021/acs.est.5b03154

Miller N, Molitor D, Zou E (2021) A causal concentration-response function for air pollution: Evidence from wildfire smoke. Working Paper. https://conference.nber.org/conf_papers/f161570.pdf

Muller NZ (2014) Using index numbers for deflation in environmental accounting. Environ Dev Econ 19(4):466–486. https://doi.org/10.1017/S1355770X1300048X

Muller NZ, Mendelsohn R (2007) Measuring the damages of air pollution in the United States. J Environ Econ Manag 54(1):1–14. https://doi.org/10.1016/j.jeem.2006.12.002

Pope CA III, Burnett RT, Turner MC (2011) Lung cancer and cardiovascular disease mortality associated with ambient air pollution and cigarette smoke: shape of the exposure-response relationships. Environ Health Persp 119(11):1616–1621. https://doi.org/10.1289/ehp.1103639

Pope CA III, Cropper M, Coggins JS, Cohen A (2015) Health benefits of air pollution policy: role of the shape of the concentration-response function. J Air Waste Manage 65(5):516–522. https://doi.org/10.1080/10962247.2014.993004

Shapiro JS, Walker R (2020) Is air pollution regulation too stringent? NBER Working Paper No. w28199. https://ssrn.com/abstract=3768251

Tessum CW, Hill JD, Marshall JD (2017) InMAP: a model for air pollution interventions. PLoS ONE 12(4):e0176131. https://doi.org/10.1371/journal.pone.0176131

Thakrar SK, Balasubramanian S, Adams PJ et al (2020) Reducing mortality from air pollution in the United States by targeting specific emission sources. Environ Sci Tech Let 7(9):639–645. https://doi.org/10.1021/acs.estlett.0c00424

US Environmental Protection Agency (2020) Emissions & Generation Resource Integrated Database (eGRID). Historical eGRID Data. https://www.epa.gov/egrid/historical-egrid-data

US Environmental Protection Agency (2011) Regulatory Impact Analysis for the Final Mercury and Air Toxics Standards. Office of Air Quality Planning and Standards, Research Triangle Park EPA-452/R-11-011. https://www3.epa.gov/ttnecas1/regdata/RIAs/matsriafinal.pdf

US Environmental Protection Agency (2013) Regulatory Impact Analysis for the Final Revisions to the National Ambient Air Quality Standards for Particulate Matter. Office of Air Quality Planning and Standards, Research Triangle Park EPA-452/R-12-005. https://www3.epa.gov/ttnecas1/regdata/RIAs/finalria.pdf

US Environmental Protection Agency (2015) Regulatory Impact Analysis for the Clean Power Plan Final Rule. Office of Air Quality Planning and Standards, Research Triangle Park EPA-452/R-15-003. https://www3.epa.gov/ttnecas1/docs/ria/utilities_ria_final-clean-power-plan-existing-units_2015-08.pdf

US Environmental Protection Agency (2010) Guidelines for Preparing Economic Analyses. National Center for Environmental Economics. https://www.epa.gov/environmental-economics/guidelines-preparing-economic-analyses

US Environmental Protection Agency (2016) Control Strategy Tool (CoST) development documentation. Office of Air Quality Planning and Standards, Research Triangle Park

US Energy Information Administration (2022) Levelized Costs of New Generation Resources in the Annual Energy Outlook.https://www.eia.gov/outlooks/aeo/pdf/electricity_generation.pdf

Viscusi WK, Aldy JE (2003) The value of a statistical life: a critical review of market estimates throughout the world. J Risk Uncertain 27(1):5–76. https://doi.org/10.1023/A:1025598106257

Vodonos A, Awad YA, Schwartz J (2018) The concentration-response between long-term $PM_{2.5}$ exposure and mortality; a meta-regression approach. Environ Res 166:677–689. https://doi.org/10.1016/j.envres.2018.06.021

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.