

The air pollution tradeoff in India: Saving more lives versus reducing the inequality of exposure

Ashwini Sankar (✉ sanka010@umn.edu)

Macalester College <https://orcid.org/0000-0003-2873-165X>

Andrew Goodkind

University of New Mexico <https://orcid.org/0000-0003-4461-5026>

Jay Coggins

University of Minnesota

Article

Keywords: Air pollution, India, Pollution inequality, Environmental health

Posted Date: January 13th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1053774/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Abstract

Chronic exposure to ambient fine particulate matter (PM_{2.5}) represents one of the largest global public health risks, leading to millions of premature deaths annually. For a country facing high and spatially variable exposures, prioritizing where to reduce PM_{2.5} concentrations leads to an inherent tradeoff between saving the most lives and reducing inequality of exposure. This tradeoff results from the shape of the concentration-response function between exposure to PM_{2.5} and mortality, which indicates that the additional lives saved per unit reduction in PM_{2.5} declines as concentrations increase. We estimate this concentration-response function for urban areas of India, finding that a 10 unit reduction in PM_{2.5} in already-clean locations will reduce the mortality rate substantially (4.2% for a reduction from 30 to 20 $\mu\text{g}\text{m}^{-3}$), while a 10 unit reduction in the dirtiest locations will reduce mortality only modestly (1.2% for a reduction from 90 to 80 $\mu\text{g}\text{m}^{-3}$). We explore the implications of this PM_{2.5}/mortality relationship by considering a thought experiment. If India had a fixed amount of resources to devote to PM_{2.5} concentration reductions across urban areas, what is the lives saved/inequality of exposure tradeoff from three different methods of employing those resources? Across our three scenarios—1) which reduces exposures for the dirtiest districts, 2) which reduces exposures everywhere equally, and 3) which reduces exposures to save the most lives—scenario 1 saves 18,000 lives per year while reducing the inequality of exposure by 65%, while scenario 3 saves 126,000 lives per year, but increases inequality by 19%.

Significance Statement

Designing policies to reduce exposure to PM_{2.5} is complicated by the apparent supralinear relationship between PM_{2.5} and premature mortality. This relationship—which we estimate for Indian urban areas—indicates that more lives can be saved by reducing exposures in the already-clean locations than in the dirtiest locations. Thus, policymakers face a troubling tradeoff, between maximizing lives saved and reducing the inequality of exposure to pollution. Many air policies impose an upper limit on exposure, thereby cleaning the dirtiest locations while reducing exposure inequality. We illustrate the tradeoff between lives saved and inequality, highlighting the challenge facing policy makers charged with protecting human health equitably.

Introduction

India faces some of the worst air pollution levels in the world from a growing number of coal-fired power plants¹, household air pollution², and numerous other sources such as open burning of agricultural residue³. Fifteen of the 20 most polluted cities in the world are in India, leading to significant numbers of premature deaths each year^{4–8}. Reducing air pollution, specifically, fine particulate matter (PM_{2.5}), would reap benefits for India in terms of improved health^{9,10}.

An important question is how and where air pollution mitigation efforts should be focused. In other words, what are the benefits, in terms of lives saved, of improving air quality in dirty or clean areas of India? A strategy that appeals to our sense of environmental justice, and is perhaps the most intuitive,

involves reducing air pollution in those places facing the highest concentrations. Indeed, that is exactly the directive of the Indian National Ambient Air Quality Standard (NAAQS), with a goal of reducing PM_{2.5} concentrations below 40 µgm⁻³ in all locations^{3,11}, a level exceeded by half of the population in India¹².

Here we scrutinize this recommendation by illustrating an inherent tradeoff for PM_{2.5} pollution policy for high-concentration countries: more lives can be saved by reducing PM_{2.5} in locations with already low concentrations than in locations with high concentrations. The concept itself is somewhat paradoxical to conventional notions of environmental improvements (i.e., that the greatest incremental benefit of eliminating pollution ought to be to those facing the worst pollution), and it is antithetical to most people's moral imperative to reduce inequality (be it pollution exposure, income, access to healthcare, etc.). We explore these ideas to make explicit these tradeoffs that are often not included in discussions of air pollution policy.

First, why are more lives saved by reducing PM_{2.5} from low- than from high-concentration locations? The answer is found in the shape of the concentration-response (C-R) function between PM_{2.5} concentration exposure and premature adult mortality. Growing evidence suggests that the C-R function is supralinear, which means that the incremental (or marginal) increase in the risk of mortality declines as PM_{2.5} concentrations rise^{6,13-17}. To be sure, this marginal risk is always positive (i.e., higher concentrations always pose a greater risk of death), but the incremental risk becomes smaller as concentrations rise. The consequence of this shape is that as we improve air quality, each additional unit of reduction in PM_{2.5} at already-low concentrations reduces mortality risk by more than each additional unit of reduction at high concentrations.

The evidence for a supralinear shape is abundant and increasingly accepted as likely for relatively low-concentration locations (in particular in the United States and Europe)^{6,13,17}. The implication of this shape in relatively clean countries is that there are large benefits to reducing PM_{2.5} concentrations to very low levels¹⁸⁻²⁰.

For countries that experience much higher concentrations, the evidence is less clear, and so are the implications. For these countries, if the C-R remains supralinear then substantial reductions in PM_{2.5} concentrations would be required before there are large reductions in mortality risk. The IER¹⁶ and GEMM¹⁵ are two C-R functions that have been estimated to understand the potential improvements in global health from reductions in PM_{2.5} concentrations. These functions combine several studies of this relationship from relatively low-concentration countries and fit the relationship across the range of exposures worldwide using either studies of high exposure to PM_{2.5} from smoking (IER), or using one study of the health effects of PM_{2.5} in China (GEMM), to estimate the health risks in high-concentration countries.

The shape and magnitude of the C-R function at high concentration levels is enormously important and not currently well understood. Here we provide an estimate of the C-R function for Indian urban districts

(henceforth referred to as districts), and assess the implications for Indian air quality improvement efforts. We find clear evidence of the supralinear shape of the C-R function and note a pronounced flattening of the curve at high concentrations. Our estimate displays a similar general shape to GEMM, though it shows substantially lower mortality risk at middle and high concentrations than does GEMM. We acknowledge that our estimate, comparing the mortality rates across Indian districts by PM_{2.5} concentrations from 1998 – 2015, is not a state-of-the-science approach using a long-term cohort survival study. Even so, our analysis serves as an initial estimate of this relationship at high concentrations, and also allows us to explore the difficult and counterintuitive policy options that accompany India's efforts to improve air quality.

We pose the question: given a set amount of resources to reduce PM_{2.5} in India, how many lives can be saved by directing those resources to different locations? Specifically, we introduce the concept of a person-PM_{2.5} unit, defined as a reduction in PM_{2.5} exposure of 1 μgm^{-3} to a single individual. Reducing exposure by 10 μgm^{-3} for the residents of an urban district with population of 100,000, for example, would reduce one million person-PM_{2.5} units. Using our dataset of Indian districts, and our estimate of the C-R function, we propose a thought experiment in which we imagine having a certain budget of person-PM_{2.5} units to “spend” to improve air quality in India. And then we compare the outcomes of three possible scenarios of how to spend this budget. Importantly, in each scenario, the same number of person-PM_{2.5} units are used.

- *Equal reduction* (hereafter “*Equal*”): all districts receive the same reduction in exposure to PM_{2.5} (in our main results all districts improve their air quality by 10 μgm^{-3} —equivalent to a roughly 20% reduction in PM_{2.5}, on average).
- *Uniform standard* (hereafter “*Standard*”): no district is permitted to be exposed to PM_{2.5} above some limit (i.e., set some limit, reduce exposure in all districts above the limit, and then lower the limit incrementally until the budget is exhausted).
- *Optimized lives saved* (hereafter “*Optimized*”): prioritize exposure reductions to those locations in which each person-PM_{2.5} unit has the largest reduction in the mortality rate (in our preferred log-log model estimates, these are the locations with the lowest initial concentration and the highest initial mortality rate).

With these three scenarios (each of which produces the same population-weighted average exposure across the districts under study) we compare the number of lives saved, and the inequality in pollution exposure across districts. Inequality in exposure is measured by the Gini coefficient (generally used as a measure of wealth or income inequality), which captures the degree to which exposure to pollution varies across locations. To illustrate the importance of the shape of the C-R function, we produce our results using three estimates of the mortality/PM_{2.5} relationship—two from our dataset, referred to as log-log and log-linear, and one from the existing literature: GEMM.

Results

Our preferred estimate of the relationship between PM_{2.5} exposure and the all-cause mortality rate in Indian urban areas, denoted “log-log,” expresses the natural log of mortality as a function of the natural log of the PM_{2.5} concentration. This estimate produces the supralinear shape of the C-R at higher concentrations (see orange line on left panel of Figure 1). The relative risk between exposure at the mean (49.7 μgm⁻³) and 10 units below the mean is 1.024 (95% CI: 1.006 to 1.042, p<0.01). With log-log, a one-unit reduction in PM_{2.5} at an already low concentration location has a much larger reduction in mortality compared with a similar reduction at a high-concentration location—e.g., the reduction in the relative risk is 4.3 times greater for a reduction from 20 to 19 μgm⁻³ than for a reduction from 100 to 99 μgm⁻³ (see orange line in right panel of Figure 1).

Our alternative estimate for comparison, denoted “log-linear,” expresses the natural log of mortality as a function of the PM_{2.5} concentration. The log-linear form is common in the epidemiological literature of health effects from air pollution. The relative risk between exposure at the mean and 10 units below the mean is 1.014 (95% CI: 1.0004 to 1.028, p<0.05). With the log-linear C-R, every unit reduction in the PM_{2.5} exposure has a similar change in the relative risk (see gray line in right panel of Figure 1). Log-log is our preferred estimate over log-linear (and several other specifications) as it improves the goodness of fit of the model to the data according to the AIC and BIC measures (see S.2 in Supplemental Information for discussion of model selection and goodness of fit). As a comparison to our estimated C-R relationships, we also calculate the results using GEMM. The GEMM function is steeper than our estimates (i.e., more lives are saved for each unit of PM_{2.5} reduced), but has less curvature than our log-log estimate (i.e., the difference in shape between low- and high-exposure locations is not as large) (see black dotted line on left panel of Figure 1).

Figure 2 illustrates the key tradeoff between lives saved and inequality in air pollution exposure that we explore across the three scenarios and for the three functional forms of the C-R relationship. The three scenarios all utilize the same budget of person-PM_{2.5} unit reductions, but employ them very differently. The size of the budget is set by assuming a 10-unit reduction across all districts. This is the *Equal* scenario. Instead of reducing all districts by 10 units, the *Standard* scenario reduces pollution in the dirtiest locations first, and exhausts its budget by ensuring no district is exposed to more than 53 μgm⁻³. The *Optimized* scenario picks those locations where a one-unit reduction in PM_{2.5} will have the largest effect on the mortality rate, and continues until the budget is exhausted—we limit reductions in any district to no less than 5 μgm⁻³, roughly equivalent to a very clean location in the United States.

The y-axis of Figure 2 shows the percent reduction in the Gini coefficient compared with the status quo pollution exposures—where positive values represent greater equality in exposures across districts, and negative values a widening of inequality. For each functional form, the *Standard* scenario provides the greatest gains in pollution exposure equality, but the fewest lives saved; the *Equal* scenario delivers more lives saved than the *Standard* scenario and a small reduction in inequality; and the *Optimized* scenario

increases inequality of exposure but saves the greatest number of lives. Our focus is on the steepness of the lives saved/inequality tradeoff between the scenarios.

For our alternative log-linear, the tradeoff is very steep, suggesting that a relatively small number of additional lives are saved (13,600 lives) in the *Optimized* scenario compared with the *Standard* scenario, but at a huge cost in inequality of exposure (65% reduction in inequality to 9% increase in inequality).

For our preferred log-log, the tradeoff is much flatter, which means that although exposure inequality declines sharply as we move from the *Optimized* scenario to *Equal*, and then to the *Standard* scenario, the number of lives saved declines by over 80%. The *Standard* scenario, which reduces inequality by 65%, saves 28,000 fewer lives than reducing pollution by 10 units in all locations (*Equal* scenario), which reduces inequality by only 14%. Compared with the *Equal* scenario, the *Optimized* scenario saves an additional 79,400 lives, but comes with an increase in the inequality of exposure (19% increase compared with the initial situation).

For GEMM, the tradeoff is similar to log-log but shifted to the right—more lives are saved in all scenarios compared with log-log.

Figure 3 examines the number of lives saved in each scenario (by the height of the bars) and also shows the initial concentration of the districts where the lives are being saved (by the colors). Across the three scenarios, it is not just that the number of lives saved varies, but the lives are saved from very different locations. The *Standard* scenario protects those who are initially exposed to the highest concentrations, whereas the *Optimized* scenario generally saves lives for those with relatively low initial concentrations. If the relationship is log-linear, the gains in lives saved are relatively small in the *Optimized* and *Equal* scenarios over the *Standard* scenario, but the composition of those lives is very different.

With log-log, large gains are achieved under the *Optimized* scenario over the *Standard* scenario, in total lives saved (18,300 versus 125,700). However, the *Standard* scenario saves 14,000 lives for those exposed to over $75 \mu\text{gm}^{-3}$, whereas the *Optimized* scenario saves zero lives from these groups. The vast majority of the lives saved in the *Optimized* scenario (91%), are those of people living in below average $\text{PM}_{2.5}$ concentration locations. Again, GEMM illustrates a similar tradeoff in terms of lives saved and inequality of exposure to log-log, with larger overall magnitudes of lives saved, but lesser ratios of lives saved between the *Optimized* and *Standard* scenarios.

In aggregate, the outcomes from our three scenarios show vast differences. Examining how the scenarios direct their budget of person- $\text{PM}_{2.5}$ units to different Indian districts, leading to these differences, is instructive for understanding how the scenarios are implemented. Figure 4 plots four outcomes (panels A, B, C and D) for each Indian district, across the three scenarios, using our preferred log-log estimate. The districts are organized across the x-axis according to the initial $\text{PM}_{2.5}$ concentration (each bubble represents a district, and the bubble size is proportional to the population). In panel A, we see the concentration reduction for each district. In the *Equal* scenario (blue), all districts are reduced by 10 units. The *Standard* scenario (red) takes the 27 districts with the highest concentrations and reduces them each

down to $53 \mu\text{gm}^{-3}$. With this standard, 82% of the person- $\text{PM}_{2.5}$ units are spent on reducing the eight districts with initial concentrations above $80 \mu\text{gm}^{-3}$. The *Optimized* scenario utilizes its budget very differently, mostly reducing exposure for the cleanest districts. This scenario reduces pollution for 50 districts, 48 of which are in below average concentration locations. The pollution in each of these districts is reduced to a pristine level of $5 \mu\text{gm}^{-3}$.

In panel B, we see how the concentration reductions change the relative risk of mortality. This panel appears similar to A, but here we see the effect of the supralinear shape of the C-R function. The *Equal* scenario reduces the risk of mortality substantially more for the initially low-concentration districts than the initially high-concentration districts, even though their concentration change is the same. This flattening is also seen in the other scenarios, most obviously for the *Standard* scenario such that despite very large reductions in concentrations, there are less substantial reductions in risks for the initially high-concentration districts.

Panel C is subtly different from panel B, in looking at the reduction in the mortality rate, rather than the change in the relative risk. All else equal, a given unit of reduction in pollution in a location with a high initial mortality rate will reduce the absolute mortality rate more than that same reduction in a location with a low initial mortality rate. The *Optimized* scenario directs its pollution reductions to those locations with relatively low initial concentrations—to take advantage of the steepest portion of the C-R function—and to those locations with the highest initial mortality rates. Panel C highlights the enormous improvements in conditions the *Optimized* scenario makes to the selected districts that are included.

Finally, the number of lives saved for each scenario and each district is shown in panel D. In the *Equal* scenario, despite reducing all locations by $10 \mu\text{gm}^{-3}$, 87% of the lives saved are to the half of the population living in the initially cleanest locations. The clearest distinction between the scenarios is illustrated by the large red bubble in the upper right of panel D, representing Delhi. The *Standard* scenario saves 7,400 lives in Delhi (40% of the total across all districts for the scenario, expending 42% of the exposure-reduction budget), and reduces the concentration from 106 to $53 \mu\text{gm}^{-3}$. The *Equal* scenario saves 1,070 lives in Delhi, representing 2.3% of the lives saved in the scenario but 8% of the person- $\text{PM}_{2.5}$ units in the budget. The *Optimized* scenario spends no resources on improving pollution from Delhi. In comparison, the *Optimized* scenario spends 9.4% of its budget of person- $\text{PM}_{2.5}$ units in the district of Bangalore (one-fourth as much of the budget as the *Standard* scenario spent on Delhi), reducing the concentration from 26.9 to $5 \mu\text{gm}^{-3}$ and saving 8,200 lives.

Discussion

In this paper we present an estimate of the relationship between all-cause premature mortality and $\text{PM}_{2.5}$ concentrations in India. For high-concentration countries, there is both a lack of empirical evidence on this relationship, and an under-appreciation of the difficult tradeoffs that exist for air pollution policies. Those tradeoffs must strike a balance between saving the most lives and achieving goals of more equal

pollution exposures across the population. The results of our thought experiment, comparing hypothetical policy scenarios, is designed to provide some clarity regarding both points.

Our empirical estimate of the C-R function shows clear evidence in support of the supralinear shape (see orange line in Figure 1), which, for medium-low concentration locations, is largely consistent with the available literature. Below $40 \mu\text{gm}^{-3}$, the curvature of our estimate is similar to GEMM, but the magnitude is slightly lower. For medium-high concentration locations we see a sharper bending down of the relationship than GEMM, and a substantially lower magnitude.

To compute the results of our scenarios (in particular the *Optimized* scenario), we found it necessary to reduce concentrations below the lowest observed concentrations in our sample, thereby taking our calculations outside the range of our observed data. As a sensitivity analysis, we reran our scenarios combining our log-log C-R function for concentrations in our sample ($\text{PM}_{2.5} \geq 17.1$) with the GEMM C-R function for concentrations below 17.1. The results are very similar, with modestly fewer lives saved in our *Optimized* scenario (see S.5 in Supplemental Information).

Our estimate indicates that in order to achieve large reductions in mortality, concentrations in the dirtiest locations require substantial reductions in $\text{PM}_{2.5}$. This finding is consistent with results from a study showing that the avoided deaths from not building new coal power plants in India is lower in places with high pollution levels¹. Our panel D in Figure 4 highlights this point, where the number of lives saved in Delhi approaches the number saved in Bangalore, which is much cleaner, only when a much greater level of cleanup is achieved. Under our *Optimized* scenario, scarce resources are devoted to relative clean Bangalore, and none on Delhi, though it is one of the dirtiest cities in the world.

The supralinear curvature is the essence of these difficult and perplexing tradeoffs. How should decision makers weigh the greater number of lives that could be saved by directing more resources to the already cleanest locations, against the unfairness of leaving behind those now suffering the effects of the worst pollution?

Our exercise is a substantial simplification of reality, and none of our three scenarios could be strictly implemented. First, our analysis assumes that the policy maker can focus on each location in isolation, but $\text{PM}_{2.5}$ pollution reductions are not confined to any one location, due to emissions dispersing up to thousands of kilometers from their source (21). Second, the concept of person- $\text{PM}_{2.5}$ units is not directly analogous to any air pollution reduction policies, and probably overstates the resources required to reduce pollution in high-population locations. Still, we believe this is a useful concept for thinking about abatement resource allocation, and our results may offer a counsel of hope. The potential to improve human health through improvements in air quality, sizable even in the rich countries of the world, is especially large and striking in India. Further study of the health and policy implications of our findings might deliver real dividends for people.

Methods And Materials

Econometric estimates

Annual mortality and PM_{2.5} data was collected for 110 Indian urban districts for the years 1998 – 2015 (for additional details see S.1 in Supplemental Information). Our dataset is at the urban district level, which is an administrative division within a state that is a conglomeration of towns and cities. We run the following model, with the natural log of the mortality rate as the dependent variable and control for linear and quadratic district-specific time trends, year fixed effects, and urban literacy rates ⁹.

$$\ln(AMR_{it}) = \alpha + \beta X_{it} + \gamma \ln(PM_{2.5,it}) + \delta_i + \tau_t + \delta_i \times t + \delta_i \times t^2 + \epsilon_{it},$$

1

where $\ln(AMR_{it})$ is the log of adjusted mortality rate in district i and year t , $\ln(PM_{2.5})$ is the log of PM_{2.5} for each district-year, and X is the control variable we use, i.e., urban literacy rate. δ_i represents the district level fixed effects and τ_t represents the year fixed effects. We have also included the district-specific linear and quadratic trends (denoted by $\delta_i \times t$ and $\delta_i \times t^2$ respectively) to control for any systematic trends specific to each district. Our preferred model, shown in Equation (1), is a log-log model in which we take the natural log of the PM_{2.5} concentration. Our alternative estimate is a log-linear model which uses the PM_{2.5} level (non-logged) (see S.2 in Supplemental Information for alternative estimated models).

The mortality data for India includes many high and low outliers. To provide a more precise estimate of the mortality-PM_{2.5} relationship, we use the Huber M-estimator ²², which corrects for outliers by adjusting by Cook's distance. In addition, we follow a rigorous identification strategy that addresses issues of reverse causality, unobserved heterogeneity, cumulative exposure, measurement error, violation of Stable Unit Treatment Value Assumption (SUTVA), and omitted variable bias that may confound our main parameter of estimate of interest (see S.3 in Supplemental Information for the details of our identification strategy).

Pollution reduction scenarios

Unlike the estimated C-R functions, which utilizes a panel of data over 18 years, we use a single data point for each district in our concentration-reduction scenarios. For each district, we take the average from 2011 – 2015 of the population, PM_{2.5} concentration, and the mortality rate. To construct our scenarios for reductions, we start with a budget of person-PM_{2.5} units of resources available for each scenario. For our main results, we calculate the budget based on number of person-PM_{2.5} units required to reduce the pollution level of all districts by 10 μgm^{-3} —the budget is equal the sum of the population of each district times 10 units of reduction. This also represents our *Equal* scenario. For the *Standard* scenario, we start with a pollution level (starting with a high pollution level), and lower the PM_{2.5} concentration of all districts above this level to this point, and calculate the person-PM_{2.5} units required for this reduction. As long as the total person-PM_{2.5} units are less than the budget, we lower the level and run the process again, and continue until the budget is exhausted. For the *Optimized* scenario we

calculate the marginal change in the mortality rate for each district of reducing the PM_{2.5} concentration by one unit, and then reduce the concentration in the district with the highest marginal change in the mortality rate. This process continues until the budget is exhausted (we set a minimum limit of 5 µgm⁻³ which no district can drop below). See S.4 and S.5 in Supplemental Information for additional details on the calculation of lives saved and the concentration-reduction scenarios.

Inequality of exposure metric

The Gini coefficient (G) for pollution exposure inequality is calculated for the initial concentration and for each reduction scenario. Define $p_i = \text{pop}_i / \sum_{i=1}^N \text{pop}_i$ as the share of the population in district i , and $s_i = (q_i \cdot p_i) / \sum_{i=1}^N (q_i \cdot p_i)$ as the share of air quality “wealth” in district i , where the air quality is $q_i = \text{PM}^{\text{max}} - \text{PM}_i$, PM^{max} is the highest PM_{2.5} concentration for any district in the scenarios, equal to 112.1, and PM_i is the PM_{2.5} concentration for a scenario in district i . The equation for the Gini coefficient is $G = A / (A + B)$ where B is the area under the Lorenz curve²³, and A is the area between the line of parity and the Lorenz curve. G can be rewritten as $G = 1 - 2B$ because $A + B = 0.5$, and

$$B = \sum_{k=1}^N p_k \left(\frac{s_k}{2} + \sum_{i=0}^{k-1} s_i \right)$$

where the s_i and p_i are first sorted in ascending order according to the air quality q_i , and $s_0 = 0$. This produces a value between 0 and 1, with 0 representing perfect equality, and 1 representing perfect inequality. We calculate the change in inequality as a percentage change from the initial concentration Gini coefficient.

References

1. Cropper, M. *et al.* The mortality impacts of current and planned coal-fired power plants in India. *Proc. Natl. Acad. Sci. U. S. A.* **118**, (2021).
2. Chowdhury, S. *et al.* Indian annual ambient air quality standard is achievable by completely mitigating emissions from household sources. *Proc. Natl. Acad. Sci.* **116**, 10711–10716 (2019).
3. Venkataraman, C. *et al.* Source influence on emission pathways and ambient PM_{2.5} pollution over India (2015–2050). *Atmos. Chem. Phys.* 8017–8039 (2018) doi:https://doi.org/10.5194/acp-18-8017-2018, 2018.
4. Balakrishnan, K. *et al.* The impact of air pollution on deaths, disease burden, and life expectancy across the states of India: the Global Burden of Disease Study 2017. *Lancet* **3**, E26–E39 (2019).
5. Dominici, F., Daniels, M., Zeger, S. L. & Samet, J. M. Air Pollution and Mortality: Estimating Regional and National Dose-Response Relationships. *J. Am. Stat. Assoc.* **97**, 100–111 (2002).
6. Krewski, D. *et al.* Reanalysis of the Harvard Six Cities Study and the American Cancer Society Study of Particulate Air Pollution and Mortality. *Spec. Rep. to Heal. Eff. Inst.* (2009).

7. Lepeule, J., Laden, F., Dockery, D. & Schwartz, J. Chronic Exposure to Fine Particles and Mortality: An Extended Follow-up of the Harvard Six Cities Study from 1974 to 2009. *Environ. Health Perspect.* **120**, 965–970 (2012).
8. Pope III, A. C. *et al.* Particulate Air Pollution as a Predictor of Mortality in a Prospective Study of U.S. Adults. *Am. J. Respir. Crit. Care Med.* **151**, 669–674 (1995).
9. Sankar, A., Coggins, J. S. & Goodkind, A. L. Effectiveness of air pollution standards in reducing mortality in India. *Resour. Energy Econ.* **62**, (2020).
10. Chowdhury, S., Dey, S. & Smith, K. R. Ambient PM_{2.5} exposure and expected premature mortality to 2100 in India under climate change scenarios. *Nat. Commun.* **9**, (2018).
11. CPCB. *National Ambient Air Quality Standards, Central Pollution Control Board.* (2009).
12. Pant, P. *et al.* Monitoring particulate matter in India: recent trends and future outlook. *Air Qual. Atmos. Heal.* **12**, 45–58 (2019).
13. Crouse, D. L. *et al.* A New Method to Jointly Estimate the Mortality Risk of Long-Term Exposure to Fine Particulate Matter and its Components. *Sci. Rep.* **6**, (2016).
14. Pope III, A. C. *et al.* Lung Cancer and Cardiovascular Disease Mortality Associated with Ambient Air Pollution and Cigarette Smoke: Shape of the Exposure–Response Relationships. *Environ. Health Perspect.* **119**, (2011).
15. Burnett, R. *et al.* Global estimates of mortality associated with longterm exposure to outdoor fine particulate matter. *Proc. Natl. Acad. Sci.* **115**, 9592–9597 (2018).
16. Burnett, R. T. *et al.* An integrated risk function for estimating the global burden of disease attributable to ambient fine particulate matter exposure. *Environ. Health Perspect.* **122**, (2014).
17. Vodonos, A., Awad, Y. A. & Schwartz, J. The concentration-response between long-term PM_{2.5} exposure and mortality; A meta-regression approach. *Environ. Res.* **166**, 677–689 (2018).
18. Marshall, J. D., Apte, J. S., Coggins, J. S. & Goodkind, A. L. Blue Skies Bluer? *Environ. Sci. Technol.* **49**, (2015).
19. Goodkind, A. L., Coggins, J. S. & Marshall, J. D. A Spatial Model of Air Pollution: The Impact of the Concentration-Response Function. *J. Assoc. Environ. Resour. Econ.* **1**, 451–479 (2014).
20. Pope III, A. C., Cropper, M., Coggins, J. S. & Cohen, A. Health benefits of air pollution abatement policy: Role of the shape of the concentration–response function. *J. Air Waste Manag. Assoc.* **65**, 516–522 (2015).
21. Goodkind, A. L., Tessum, C. W., Coggins, J. S., Hill, J. D. & Marshall, J. D. Fine-scale damage estimates of particulate matter air pollution reveal opportunities for location-specific mitigation of emissions. *Proc. Natl. Acad. Sci.* **116**, 8775–8780 (2019).
22. Huber, P. J. Robust regression: asymptotics, conjectures, and Monte Carlo. *Ann. Stat.* **1**, 799–821 (1973).
23. Lorenz, M. O. Methods of Measuring the Concentration of Wealth. *Methods Meas. Conc. Wealth* **9**, 209–219 (1905).

Figures

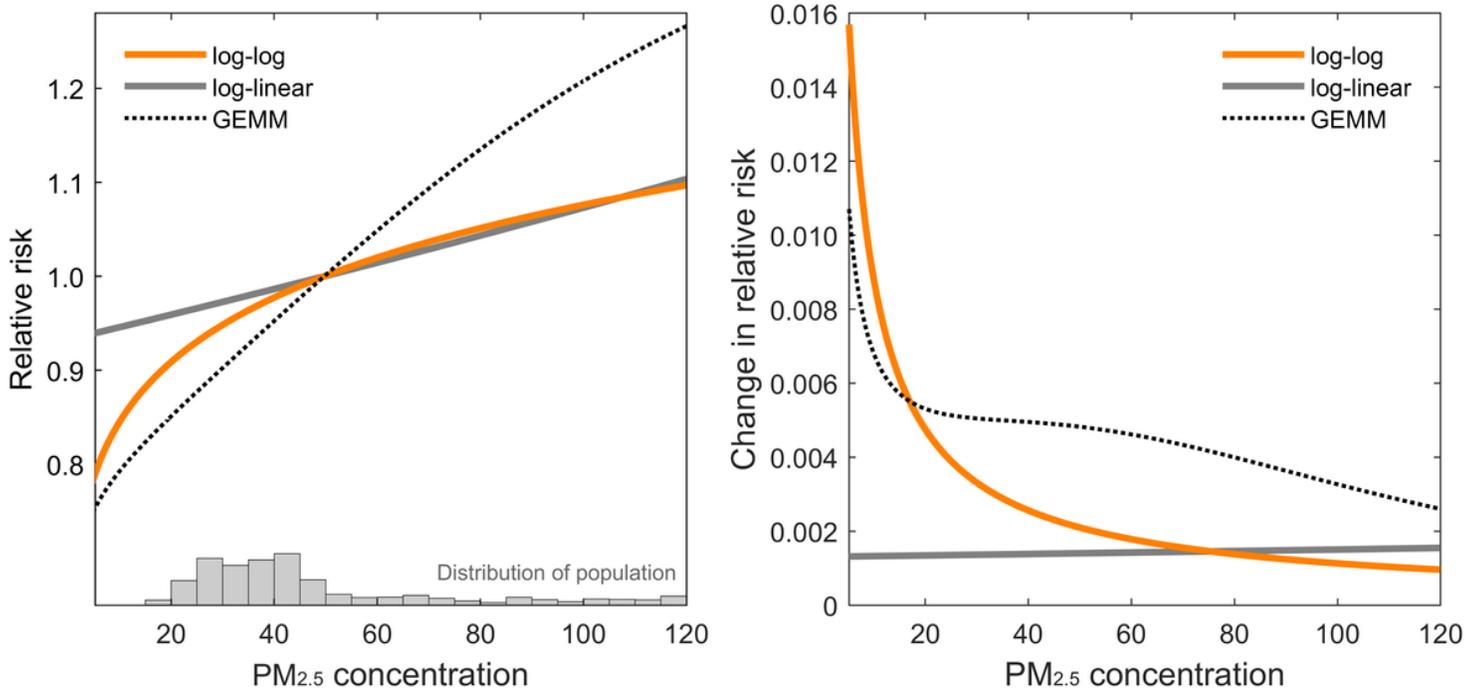


Figure 1

Relationship between all-cause mortality rate and PM_{2.5} concentration. Left panel: relative risk of mortality as a function of PM_{2.5} concentration (relative risk equal to 1 at the mean: 49.7 μg/m³). The lowest level of PM_{2.5} is 17 μg/m³ in this dataset. The distribution of urban population exposures from 1998-2005 is displayed along the bottom. Right panel: the change in the relative risk of mortality from a one μg/m³ unit increase in PM_{2.5} concentration. Log-log and log-linear estimated functions from our Indian urban dataset; GEMM from Burnett et al. (2018) for comparison.

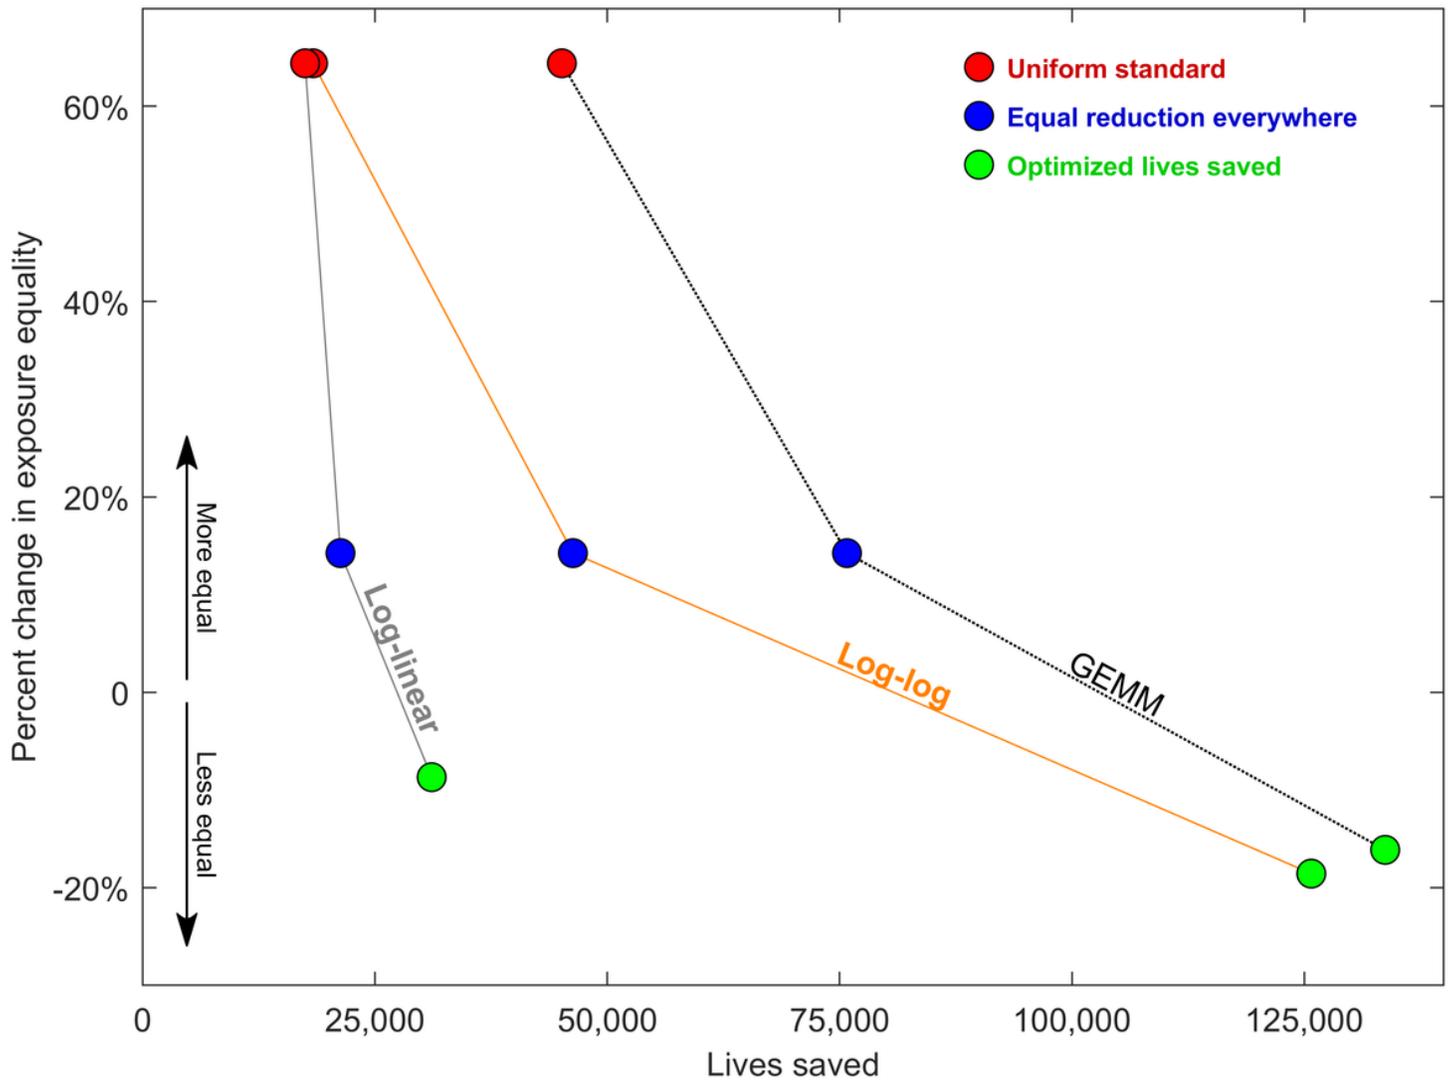


Figure 2

Tradeoff between lives saved and reducing inequality of PM_{2.5} exposure. Every dot represents the lives saved (x-axis) and the percentage reduction (or increase) in the Gini coefficient (y-axis) for a pollution reduction scenario compared with the status quo. Each of the three scenarios are displayed for three C-R functions: log-log, log-linear, and GEMM. The lines connecting the three scenarios for each respective C-R function are used only to illustrate the grouping of scenarios.

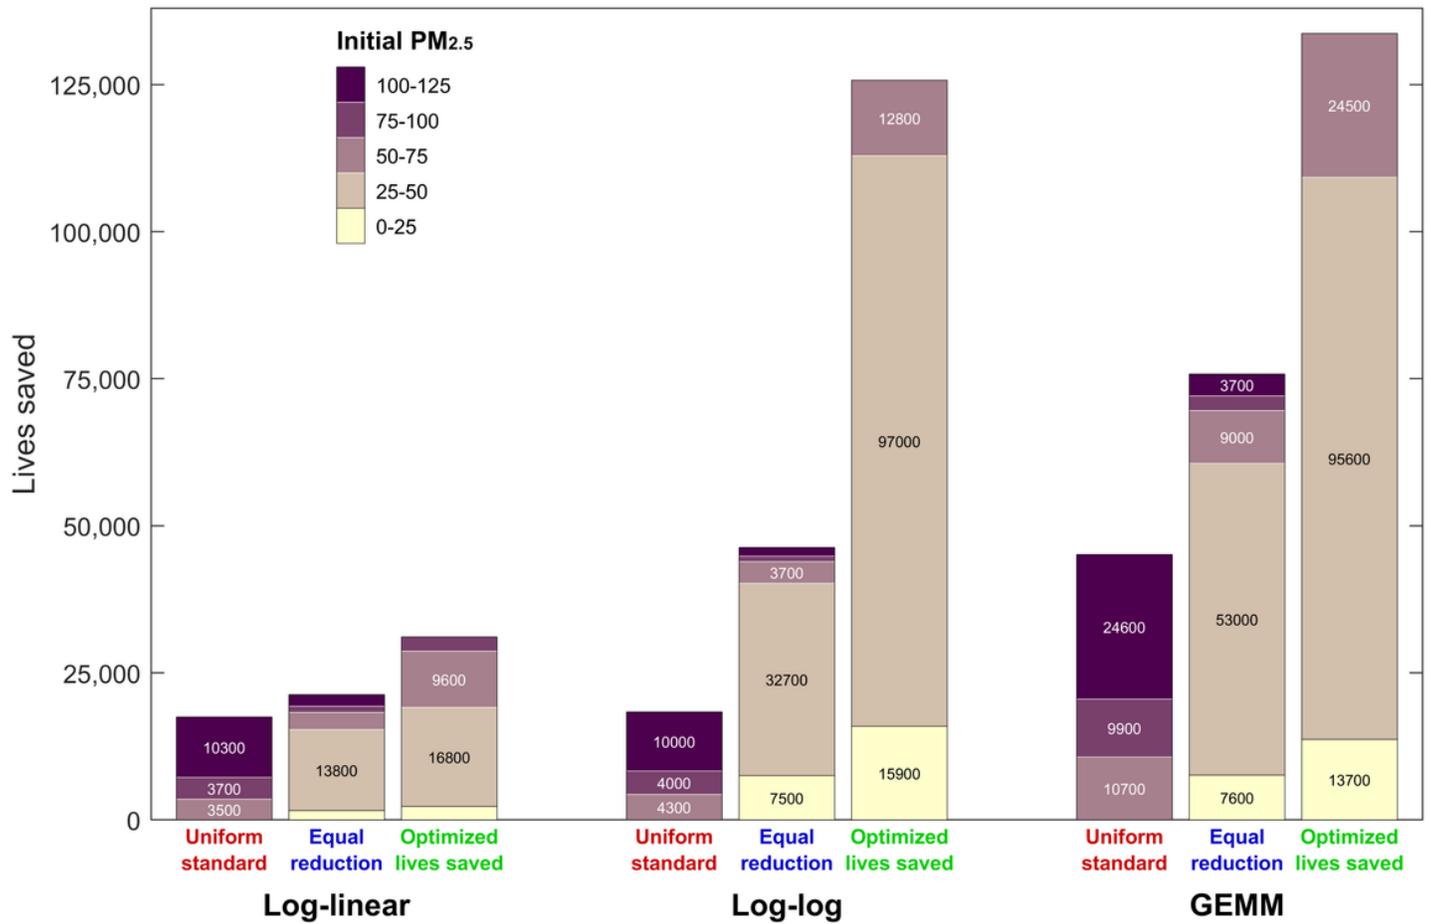


Figure 3

Composition of lives saved under each scenario by initial PM_{2.5} concentration. The bar height represent the number of lives saved, and the bar colors represent the initial PM_{2.5} concentration of the districts affected by the scenarios. The numbers are rounded to the nearest 100, and values below 2,000 are omitted.

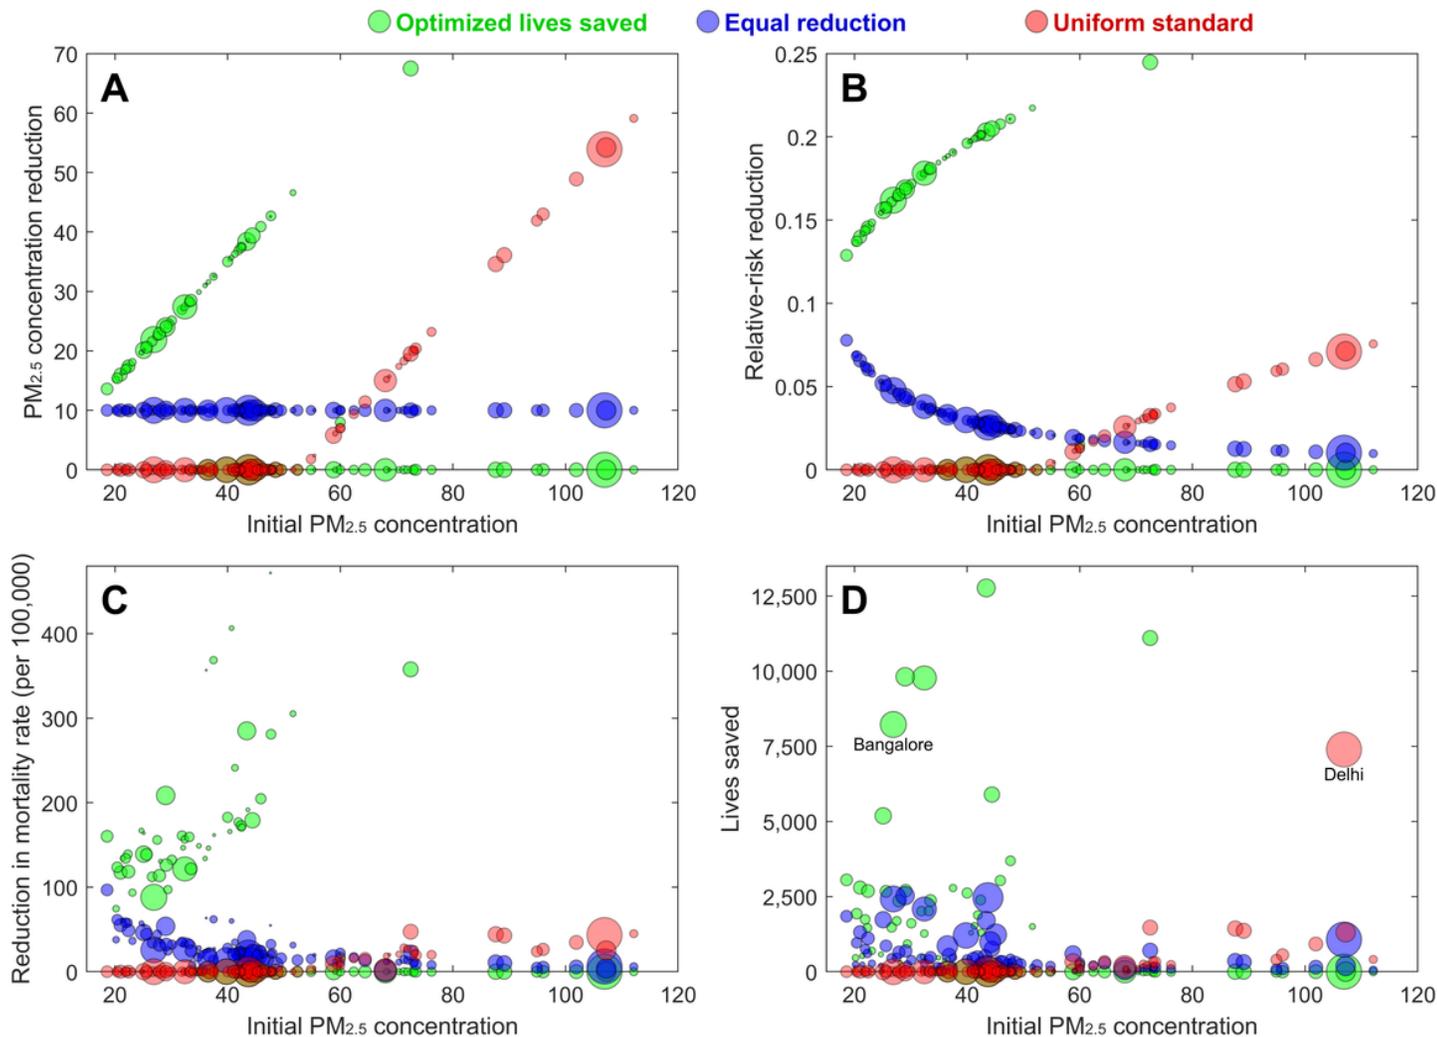


Figure 4

Outputs by district using our preferred log-log model estimation. Districts are organized by initial PM_{2.5} concentration (*x*-axis), bubbles are proportional to district population, and colors represent the three scenarios (each district is represented by three bubbles in each panel, one for each scenario). Panel A illustrates how much concentrations are reduced in each district. Panel B shows the reduction in the relative-risk of mortality due to the concentration reduction (*e.g.*, 0.1 represents a 10% reduction in the risk of mortality). Panel C shows the reduction in the mortality rate, which differs from panel B according to differences by district in the status quo mortality rate. Panel D shows the lives saved in each district resulting from the concentration reduction.

Supplementary Files

This is a list of supplementary files associated with this preprint. [Click to download.](#)

- [sianon.docx](#)