

Robust relationship between air quality and infant mortality in Africa

Sam Heft-Neal¹, Jennifer Burney², Eran Bendavid³ & Marshall Burke^{1,4,5*}

Poor air quality is thought to be an important mortality risk factor globally^{1–3}, but there is little direct evidence from the developing world on how mortality risk varies with changing exposure to ambient particulate matter. Current global estimates apply exposure–response relationships that have been derived mostly from wealthy, mid-latitude countries to spatial population data⁴, and these estimates remain unvalidated across large portions of the globe. Here we combine household survey-based information on the location and timing of nearly 1 million births across sub-Saharan Africa with satellite-based estimates⁵ of exposure to ambient respirable particulate matter with an aerodynamic diameter less than 2.5 μm ($\text{PM}_{2.5}$) to estimate the impact of air quality on mortality rates among infants in Africa. We find that a 10 $\mu\text{g m}^{-3}$ increase in $\text{PM}_{2.5}$ concentration is associated with a 9% (95% confidence interval, 4–14%) rise in infant mortality across the dataset. This effect has not declined over the last 15 years and does not diminish with higher levels of household wealth. Our estimates suggest that $\text{PM}_{2.5}$ concentrations above minimum exposure levels were responsible for 22% (95% confidence interval, 9–35%) of infant deaths in our 30 study countries and led to 449,000 (95% confidence interval, 194,000–709,000) additional deaths of infants in 2015, an estimate that is more than three times higher than existing estimates that attribute death of infants to poor air quality for these countries^{2,6}. Upward revision of disease-burden estimates in the studied countries in Africa alone would result in a doubling of current estimates of global deaths of infants that are associated with air pollution, and modest reductions in African $\text{PM}_{2.5}$ exposures are

predicted to have health benefits to infants that are larger than most known health interventions.

Epidemiological studies consistently highlight poor air quality as an important contributor to death and disability, with recent estimates showing that exposure to ambient $\text{PM}_{2.5}$ is associated with 3–4 million global deaths annually^{1,2}. Such estimates are influential in a wide variety of research activities^{3,7,8} and policy decisions, including the allocation of health resources, the designation of pollution standards, and the adoption of climate-mitigation policies.

However, the relationship between air quality and mortality in the developing world—where a large proportion of the deaths that can be attributable to poor air quality are thought to occur—remains poorly quantified, limiting our understanding of relative disease burdens and appropriate policy responses. Broad-scale evidence on the health burden of exposure to ambient air pollution comes mainly from developed countries^{1,2,4}, where co-morbidities differ and where both mortality rates and average ambient $\text{PM}_{2.5}$ concentrations are typically much lower (Extended Data Fig. 1). In much of the developing world, limited air pollution data make quantification of dose–response functions challenging^{9,10}, and it is unknown whether large recent declines in infant mortality¹¹ would increase the relative health effects of poor air quality (if other unrelated causes of death are now less important) or decrease them (if children are more resilient).

Here we quantify associations between air quality and the health of infants in Africa by combining recent satellite-based measurements of annual ambient $\text{PM}_{2.5}$ concentrations with household survey data on




Fig. 1 | Spatial patterns of pollution and infant mortality in Africa. **a**, Long-run average $\text{PM}_{2.5}$ concentration for 2001–2015. **b**, Maximum annual $\text{PM}_{2.5}$ concentration for 2001–2015. **c**, Average infant mortality rate

(IMR) in study countries for 2001–2015, derived from Demographic and Health Surveys using previously described methods¹¹. Country outlines were obtained from Global Administrative Areas, version 2.0³⁰.

¹Center on Food Security and the Environment, Stanford University, Stanford, CA, USA. ²School of Global Policy and Strategy, University of California, San Diego, San Diego, CA, USA. ³School of Medicine, Stanford University, Stanford, CA, USA. ⁴Department of Earth System Science, Stanford University, Stanford, CA, USA. ⁵National Bureau of Economic Research, Cambridge, MA, USA. *e-mail: mburke@stanford.edu




Fig. 2 | Mortality of infants in Africa is strongly and linearly increasing with post-birth $\text{PM}_{2.5}$ exposure. **a**, Effect of $\text{PM}_{2.5}$ exposure during the 12 months after birth on mortality rates of infants ($n = 990,696$ births). Response function is centred at mean $\text{PM}_{2.5}$ concentration ($25 \mu\text{g m}^{-3}$) and mean IMR (71 deaths per 1,000 births). Histogram shows the distribution of exposures across sample locations. **b**, Impacts of in utero versus post-birth exposures. **c**, Impacts of post-birth exposure in West

Africa (higher exposure) versus the rest of Africa (lower exposure). See Extended Data Fig. 2b for countries in each region. **d**, Effect of post-birth exposure on child mortality by terciles of household-level asset wealth, measured as the percentage change in infant mortality per $10 \mu\text{g m}^{-3}$ increase in $\text{PM}_{2.5}$ exposure. **e**, Effect of post-birth $\text{PM}_{2.5}$ exposure on IMR over time, measured as the percentage change in IMR per $10 \mu\text{g m}^{-3}$ increase in $\text{PM}_{2.5}$ exposure.

infant mortality (death in the first 12 months of life) as measured in the Demographic and Health Surveys, a set of nationally representative household health surveys. We use data from 65 available Demographic and Health Surveys across 30 sub-Saharan African (SSA) countries carried out between 2001 and 2015, representing 990,696 births over the period (Extended Data Fig. 2). We match the location and timing of each birth to satellite-based estimates of $\text{PM}_{2.5}$ exposure from 9 months before to 12 months after birth⁵ (Fig. 1). These satellite data offer critical advantages in SSA, where only two countries have air-pollution monitoring stations that report to global databases¹², and where chemical transport model-based exposure estimates rely on emission inventories that have high degrees of uncertainty in rural biomass-burning areas^{13,14}.

We model the effects of $\text{PM}_{2.5}$ exposure on infant mortality using fixed-effects regression analyses that flexibly account for time-invariant differences in air pollution and mortality across locations, local seasonality in both air quality and mortality, and trending factors or abrupt shocks common to all locations in our sample (see Methods). Because seasonally adjusted variation in $\text{PM}_{2.5}$ levels over time at a given location is plausibly exogenous, we propose that this approach isolates the role of poor air quality from other confounding variables that affect mortality risk.

Infant mortality in SSA strongly and linearly increases with $\text{PM}_{2.5}$ exposure in our data (Fig. 2a, Extended Data Fig. 3 and Extended Data Table 1). A $10 \mu\text{g m}^{-3}$ increase in $\text{PM}_{2.5}$ exposure in the first 12 months of life is associated with a 9.2% increase in infant mortality ($P < 0.01$). We find no qualitative difference between exposure before and after birth at average exposure levels (Fig. 2b), although prenatal associations appear to decline at higher exposure levels. Consistent with recent US evidence¹⁵, we estimate positive associations between $\text{PM}_{2.5}$ exposure and mortality at exposure levels below the WHO (World Health Organization)-recommended guideline of $10 \mu\text{g m}^{-3}$ annual average exposure¹⁶ (Extended Data Fig. 4).

Our results are robust to models that use only within-household variation in mortality and $\text{PM}_{2.5}$ exposure over time, that allow differential country-level trends in mortality and $\text{PM}_{2.5}$ exposure, models that include a large set of additional controls, including temperature, precipitation and household- and child-specific demographic information, and models that use only cross-sectional variation in $\text{PM}_{2.5}$ exposure and mortality (Extended Data Figs. 3, 4 and

Extended Data Table 1). Similarly, $\text{PM}_{2.5}$ exposure in months 13–24 after birth does not predict mortality in the first year of life (Extended Data Fig. 8i). These findings reduce concerns that results are driven by unobserved factors that are correlated with mortality and $\text{PM}_{2.5}$ exposure (for example, spurious time-trending variables), or by the relocation of higher mortality households into locations with poorer air quality (Methods).




Fig. 3 | Comparing the relative risk curve for all-cause mortality from this study for SSA and the risk curve for respiratory-infection-specific mortality estimated for the Global Burden of Disease (GBD) study.

Data for the GBD project were previously published⁴. The GBD acute lower respiratory infection relative risk curve (red) is an integrated exposure response combining point estimates from ambient air pollution studies, indoor air pollution studies and second-hand smoking studies. The relative risk curve estimated in this study (blue) is derived by empirically relating observed births and ambient $\text{PM}_{2.5}$ concentrations in SSA (see Methods), with the shaded region representing the bootstrapped 5–95th confidence interval. The histograms show the share of population exposed to different ambient $\text{PM}_{2.5}$ concentrations in the regions corresponding to the estimation of each curve. The x axis is restricted to the range of ambient $\text{PM}_{2.5}$ concentrations observed in our SSA sample.




Fig. 4 | Avoided infant deaths from reduced $PM_{2.5}$ exposure.

a, b, Estimated share (**a**) and number of infant deaths (**b**) in 30 SSA countries that would have been avoided in 2015 if observed $PM_{2.5}$ levels were reduced to sample minimum exposure of $2 \mu g m^{-3}$. The colours in **b** correspond to the shares shown on the map in **a**. **c**, Comparison of estimated reductions in infant mortality from achieving 100% coverage

of various health interventions in our study countries from the Lives Saved Tool¹⁹, and estimated reductions in mortality resulting from a $PM_{2.5}$ reduction of $5 \mu g m^{-3}$ calculated using the relative risk functions from GBD⁴ or this study. Country outlines were obtained from Global Administrative Areas, version 2.0³⁰. DRC, Democratic Republic of Congo.

Despite differences in the source and level of $PM_{2.5}$ concentrations across African regions (Fig. 1), we find similar associations between $PM_{2.5}$ and mortality when the dataset is restricted to West Africa (where $PM_{2.5}$ levels are higher and partly arise from dust) versus the rest of Africa (where exposures are lower and sources are mainly anthropogenic; Fig. 2c and Extended Data Fig. 2c). Given that much of the $PM_{2.5}$ exposure in West Africa is not a result of local economic activity, these results also suggest that such (unobserved) activity is not biasing our estimated associations between $PM_{2.5}$ exposure and mortality. Similarly, we find no estimated difference in the $PM_{2.5}$ –mortality relationship between households using ‘clean’ cooking fuels that produce no indoor particulates and households using ‘dirty’ $PM_{2.5}$ -producing cooking fuels, such as biomass, wood, agricultural residues or dung (Extended Data Fig. 5 and Extended Data Table 2), which suggests that unobserved indoor exposures are not biasing our estimated $PM_{2.5}$ –mortality relationship (see Methods).

A common hypothesis in the environmental health literature is that wealthier households can better avoid the negative health effects of hazardous environmental exposures^{17,18}. However, we do not find evidence that wealth is protective in our setting: associations between $PM_{2.5}$ exposure and mortality risk are similar across wealth terciles

in our data (Fig. 2d), are not moderated by other socio-economic or demographic characteristics (Extended Data Fig. 5 and Extended Data Table 1), and have not declined over time (Fig. 2e and Extended Data Fig. 6).

We use model estimates to construct a relative risk curve for SSA (see Methods). Relative to the mortality risk at the lowest observed exposure levels in our sample ($2 \mu g m^{-3}$), we estimate a 31% increase in mortality risk at sample median exposure levels ($22 \mu g m^{-3}$) (Fig. 3). On the basis of this risk curve, we estimate that if $PM_{2.5}$ concentrations in SSA had been reduced to an annual average of $2 \mu g m^{-3}$, 22% (95% confidence interval, 9–35%) of infant deaths would have been averted, with the largest reductions in areas with high average exposure (such as most of West Africa; Fig. 4a and Extended Data Fig. 7). We calculate that exposure to $PM_{2.5}$ levels above $2 \mu g m^{-3}$ was associated with 449,000 (95% confidence interval, 194,000–709,000) additional infant deaths in 2015 in the 30 study countries alone, with over 40% of these occurring in Nigeria (Fig. 4b).

Although reducing $PM_{2.5}$ concentrations to $2 \mu g m^{-3}$ is probably not feasible, substantial reductions in mortality could still be achieved by relatively modest reductions in $PM_{2.5}$ concentrations. We estimate that reducing $PM_{2.5}$ concentrations uniformly by $5 \mu g m^{-3}$ at all locations

in our sample countries—a reduction comparable to that achieved by the US Clean Air Act²⁵ (Methods)—would have reduced infant mortality by 4.6% (95% confidence interval, 1.8–7.4%) and avoided 40,000 (95% confidence interval, 20,000–70,000) infant deaths in 2015. This reduction exceeds the estimated mortality reductions that would be obtained if many key child health interventions—including vaccines, nutritional supplementation and insecticide-treated bed nets—were scaled from current levels to 100% population coverage across our study countries¹⁹ (Fig. 4c and Methods). We caution that this comparison does not account for the feasibility or cost effectiveness of achieving these alternative reductions or interventions, but instead provides a key input for future analysis.

Given that current estimates show that around 13% of the mortality of children under 5 years of age in Africa is attributable to lower respiratory infection (LRI)²⁰, our estimate that 22% of infant deaths are attributable to PM_{2.5} exposure would be implausible if LRI were the only channel through which PM_{2.5} exposure affected infant mortality. However, consistent with many recent studies^{21–23} (Methods), we find substantial evidence for effects mediated by non-respiratory channels (Extended Data Fig. 8), including positive associations between in utero PM_{2.5} exposure and neonatal deaths and negative associations with birth weight, and harmful associations between post-birth exposure and both stunting and diarrhoea; all of which are leading causes or risk factors for infant death that overlap only partially with LRI²⁰. We interpret these findings as strong evidence that PM_{2.5} exposure can affect mortality risk through channels other than LRI. As a placebo test, we find no association between PM_{2.5} exposure and child age, sex or likelihood of a multiple birth. Finally, our main estimate of a 9% increase in mortality per 10 µg m⁻³ increase in PM_{2.5} is the same or smaller than the six locality- or country-level quasi-experimental estimates^{24–29} to which our results can be easily compared (Extended Data Fig. 8j).

Our results indicate that risks from PM_{2.5} exposure could be much higher than current global estimates suggest. At median exposure levels in Africa, our estimated relative risk of mortality is double the risk estimated by the GBD at the same exposure level (Fig. 3), and our estimated number of infant deaths associated with ambient PM_{2.5} exposure in our 30 study countries in 2015 is larger than the current GBD estimate of global infant deaths that are attributable to air pollution^{2,6}. Differences could result from our measured associations with all-cause infant mortality, which is broader than LRI-specific mortality used in current global estimates.

Our results also contrast with the common finding that economic development is protective of health^{17,18}, with our data suggesting consistent effects across wealth levels and over time. One potential explanation for this consistency is that we are studying long-term exposure to a pollutant that is small enough to penetrate buildings, making avoidance difficult even for wealthier households.

The greatest impact of poor air quality in our sample is in West Africa, where high PM_{2.5} concentrations include large fractions of dust carried by winds from the Sahara. Although the comparison between West Africa and the rest of the dataset suggests a common exposure–response function, both the particle size distribution below 2.5 µm and the chemical species comprising the PM_{2.5} are unobserved but known to vary widely across sources and regions. More research is needed to characterize these parameters remotely, and to link measurements to prospective epidemiological studies with more detail on both exposures and health outcomes.

Our results indicate that substantial reductions in mortality can be achieved with even modest decreases in ambient PM_{2.5} concentrations. The strong linear relationship between PM_{2.5} and mortality indicates that, even against a high background exposure level, mitigation efforts could deliver large mortality reductions—on par with or exceeding many leading health interventions. This finding is particularly important given the minimal protective benefit of wealth in our data. However, given the varied sources of particulate matter and its precursors across SSA, multi-sectoral and region-specific approaches to reducing exposure burdens may be necessary, and large benefits may

come from developing and adopting protective approaches in dusty regions.

Online content

Any Methods, including any statements of data availability and Nature Research reporting summaries, along with any additional references and Source Data files, are available in the online version of the paper at <https://doi.org/10.1038/s41586-018-0263-3>.

Received: 14 July 2017; Accepted: 23 May 2018;

Published online 27 June 2018.

1. Lelieveld, J., Evans, J. S., Fnais, M., Giannadaki, D. & Pozzer, A. The contribution of outdoor air pollution sources to premature mortality on a global scale. *Nature* **525**, 367–371 (2015).
2. Cohen, A. J. et al. Estimates and 25-year trends of the global burden of disease attributable to ambient air pollution: an analysis of data from the Global Burden of Diseases Study 2015. *Lancet* **389**, 1907–1918 (2017).
3. Anenberg, S. C. et al. Impacts and mitigation of excess diesel-related NO_x emissions in 11 major vehicle markets. *Nature* **545**, 467–471 (2017).
4. 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**, 397–403 (2014).
5. van Donkelaar, A. et al. Global estimates of fine particulate matter using a combined geophysical-statistical method with information from satellites, models, and monitors. *Environ. Sci. Technol.* **50**, 3762–3772 (2016).
6. Institute for Health Metrics and Evaluation. Global Burden of Disease study 2015 (GBD, 2015) results. <http://ghdx.healthdata.org/gbd-results-tool> (2016). URL.
7. Shindell, D. et al. Simultaneously mitigating near-term climate change and improving human health and food security. *Science* **335**, 183–189 (2012).
8. Zhang, Q. et al. Transboundary health impacts of transported global air pollution and international trade. *Nature* **543**, 705–709 (2017).
9. Ebisu, K., Belanger, K. & Bell, M. L. Association between airborne PM_{2.5} chemical constituents and birth weight—implication of buffer exposure assignment. *Environ. Res. Lett.* **9**, 084007 (2014).
10. West, J. J. et al. What we breathe impacts our health: improving understanding of the link between air pollution and health. *Environ. Sci. Technol.* **50**, 4895–4904 (2016).
11. Burke, M., Heft-Neal, S. & Bendavid, E. Sources of variation in under-5 mortality across sub-Saharan Africa: a spatial analysis. *Lancet Glob. Health* **4**, e936–e945 (2016).
12. Burke, M. et al. Ambient air pollution exposure estimation for the Global Burden of Disease 2013. *Environ. Sci. Technol.* **50**, 79–88 (2016).
13. Zhang, F. et al. Sensitivity of mesoscale modeling of smoke direct radiative effect to the emission inventory: a case study in northern sub-Saharan African region. *Environ. Res. Lett.* **9**, 075002 (2014).
14. Butt, E. W. et al. The impact of residential combustion emissions on atmospheric aerosol, human health, and climate. *Atmos. Chem. Phys.* **16**, 873–905 (2016).
15. Di, Q. et al. Air pollution and mortality in the medicare population. *N. Engl. J. Med.* **376**, 2513–2522 (2017).
16. Occupational and Environmental Health Team. WHO Air quality guidelines for particulate matter, ozone, nitrogen dioxide and sulfur dioxide: global update 2005: summary of risk assessment. *World Health Organization* <http://www.who.int/iris/handle/10665/69477> (WHO, 2006).
17. Patt, A. G. et al. Estimating least-developed countries' vulnerability to climate-related extreme events over the next 50 years. *Proc. Natl Acad. Sci. USA* **107**, 1333–1337 (2010).
18. Smith, K. R. et al. in *Climate Change 2014: Impacts, Adaptation, and Vulnerability* (eds Field, C. B. et al.) 709–794 (IPCC, Cambridge Univ. Press, 2014).
19. Walker, N., Tam, Y. & Friberg, I. K. Overview of the lives saved tool (list). *BMC Public Health* **13**, S1 (2013).
20. Institute for Health Metrics and Evaluation. GBD Compare data visualization. <http://vizhub.healthdata.org/gbd-compare> (2017).
21. Bell, M. L., Ebisu, K. & Belanger, K. Ambient air pollution and low birth weight in Connecticut and Massachusetts. *Environ. Health Perspect.* **115**, 1118–1124 (2007).
22. Pope, D. P. et al. Risk of low birth weight and stillbirth associated with indoor air pollution from solid fuel use in developing countries. *Epidemiol. Rev.* **32**, 70–81 (2010).
23. Pereira, G., Belanger, K., Ebisu, K. & Bell, M. L. Fine particulate matter and risk of preterm birth in Connecticut in 2000–2006: a longitudinal study. *Am. J. Epidemiol.* **179**, 67–74 (2014).
24. Chay, K. Y. & Greenstone, M. The impact of air pollution on infant mortality: evidence from geographic variation in pollution shocks induced by a recession. *Q. J. Econ.* **118**, 1121–1167 (2003).
25. Chay, K. Y. & Greenstone, M. *Air Quality, Infant Mortality, and the Clean Air Act of 1970*. Report No. 10053 (National Bureau of Economic Research, 2003).
26. Arceo, E., Hanna, R. & Oliva, P. Does the effect of pollution on infant mortality differ between developing and developed countries? Evidence from Mexico City. *Econ. J.* **126**, 257–280 (2016).
27. He, G., Fan, M. & Zhou, M. The effect of air pollution on mortality in China: Evidence from the 2008 Beijing Olympic games. *J. Environ. Econ. Manage.* **79**, 18–39 (2016).

28. Knittel, C. R., Miller, D. L. & Sanders, N. J. Caution, drivers! Children present: traffic, pollution, and infant health. *Rev. Econ. Stat.* **98**, 350–366 (2016).
29. Cesur, R., Tekin, E. & Ulker, A. Air pollution and infant mortality: evidence from the expansion of natural gas infrastructure. *Econ. J.* **127**, 330–362 (2017).
30. Global Administrative Areas. GADM database of Global Administrative Areas, version 2.0. <https://gadm.org/> (2012).

Acknowledgements We thank D. Lobell, G. McCord, M. P. Burke and W. Schlenker for useful comments and V. Tanutama for research assistance. We thank the Stanford Environmental Ventures Fund and the National Science Foundation (CNH Award #1715557) for funding.

Reviewer information *Nature* thanks R. Black, J. Lelieveld, L. Waller and the other anonymous reviewer(s) for their contribution to the peer review of this work.

Author contributions S.H.N., J.B., E.B. and M.B. designed the research; S.H.N. analysed the data; S.H.N., J.B., E.B. and M.B. interpreted results and wrote the paper.

Competing interests The authors declare no competing interests.

Additional information

Extended data is available for this paper at <https://doi.org/10.1038/s41586-018-0263-3>.

Supplementary information is available for this paper at <https://doi.org/10.1038/s41586-018-0263-3>.

Reprints and permissions information is available at <http://www.nature.com/reprints>.

Correspondence and requests for materials should be addressed to M.B.

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

METHODS

Data reporting. No statistical methods were used to predetermine sample size. The experiments were not randomized and the investigators were not blinded to allocation during experiments and outcome assessment.

Infant mortality data. Data on infant health outcomes are taken from the Demographic and Health Surveys (DHS), nationally representative surveys that are conducted in many low-income and middle-income countries. DHS have a two-stage design, in which a number of clusters are first selected from a list of enumeration areas created in a recent population census, and then households are randomly selected in each of the clusters, and women aged 15–49 years are selected from those households for in-depth surveys. In most survey waves, enumerators use global positioning system devices to collect geospatial information to identify the central point of each cluster's populated area. We used data from all 65 available surveys that were carried out between 2001 and 2015 to reconstruct a village-level time series³¹. Our sample covers 30 countries and includes 990,696 individual birth outcomes (Extended Data Fig. 2). The outcome of interest for this study is infant mortality, which is represented by a dummy variable equal to one when a child was reported to die within the first 12 months after birth. Children who were alive but less than 12 months old at the time of the survey were not included in our sample. The mean infant mortality rate in our sample is 71 deaths per 1,000 births.

Construction of the household wealth measure. The DHS record information on household ownership of a common set of durable assets. In the public distribution files, DHS release a wealth index obtained using a principal components analysis of these household assets and additional services, such as electricity, water supply and floor material, with the index in each survey normalized to that specific survey (that is, wealth quantiles are defined relative to the survey-specific asset distribution). Therefore, although this index enables identification of relative wealth within surveys, it does not allow for comparisons across countries or over time given the within-survey normalization. In order to create a wealth index that could be compared across surveys, we pooled all households with information on the following assets: water source, sanitation facilities, type of flooring, electricity, the number of rooms per person living in the house, and possession of radio, television, phone (landline or cellphone), motorcycle and car. In our dataset, 85% had information on all of these assets. The wealth index was then created using a principal components analysis procedure similar to the survey-specific DHS approach, but normalizing across the entire 65-survey sample rather than within each survey. Further details, including validation and testing of this approach, are available in a previously published paper³².

PM_{2.5} data. We use satellite-derived data on PM_{2.5} compiled by the Atmospheric Composition Analysis Group at Dalhousie University, consisting of annual bias-corrected average surface PM_{2.5} concentrations at 0.01° × 0.01° spatial resolution with global coverage⁵. Building on earlier efforts to predict PM_{2.5} from satellite observations^{33,34}, these data are derived from a suite of satellite-based atmospheric optical depth measurement instruments, including the two MODIS instruments on the Terra and Aqua Satellites, the Multi-Angle Resolution Spectroradiometer (MISR) on Terra and the Sea-viewing Wide Field-of-View Sensor (SeaWiFS) on the SeaStar satellite. These data are combined with aerosol profile measurements from the Cloud–Aerosol Lidar with Orthogonal Polarization (CALIOP) instrument aboard the Cloud–Aerosol Lidar and Infrared Pathfinder Satellite Observation (CALIPSO), and satellite and weather and seasonality data from the GEOS-Chem Chemical transport model to quantify the relationship between column Aerosol Optical Depth (AOD) and surface PM_{2.5} measured at available ground-based stations.

In order to assign PM_{2.5} exposure to individual birth outcomes, we define two exposure periods for each birth: (i) the in utero period encompassing the 9 months before birth and (ii) the post-birth period encompassing the first 12 months of life inclusive of birth month. Given that pollution exposure data are only available annually, we calculate PM_{2.5} exposure in each of these periods as the weighted averages of the annual data, where the weights represent the share of the year that falls into the time period of interest. For example, a child born in the third month of year t would be assigned an in utero pollution exposure of 2/9(exposure in year t) + 7/9(exposure in year $t-1$) and a post-birth exposure of 10/12(exposure in year t) + 2/12(exposure in year $t+1$). The mean PM_{2.5} exposure level for both pre- and post-birth periods is 25.2 µg m⁻³ in our sample.

Ideally, we would account for heterogeneity in the chemical and physical properties of the particulate mix across space and time, but at present, these properties are not directly observed at scales commensurate with overall PM_{2.5}. Moreover, estimates of PM_{2.5} properties rely heavily on aerosol models and emissions inventories³⁵, which are known to be highly uncertain in biomass-burning regions such as SSA^{13,14}. As an example, the Atmospheric Composition Analysis Group offers a version of the PM_{2.5} dataset used in this analysis with dust and seasalt aerosols removed. However, this 'dust-free' version is not based on observational partitioning methods, but simply scaled based on an emissions inventory. As such, we use only the observationally constrained full PM_{2.5} dataset for our analysis, but note that finer-grained observations of the chemical and physical properties of aerosol

particulate matter could be used in the future to understand whether and to what extent impacts change with PM_{2.5} chemical composition and size distribution.

Empirical approach. We model the relationship between infant deaths y and PM_{2.5} exposure using a least squares linear probability model:

$$y_{icnt} = f(\text{PM}_{icnt}^b, \text{PM}_{icnt}^a) + \lambda \mathbf{X}_{icnt} + \mu_c + \delta_t + \eta_{nm} + \varepsilon_{icnt} \quad (1)$$

where i indexes child, c indexes survey cluster (that is, village), t indexes birth year, and nm indexes country and month. PM_{icnt}^b and PM_{icnt}^a refer, respectively, to PM_{2.5} exposure in the 9 months before and 12 months after birth. $\mathbf{X}_{icnt}$ is a vector of additional controls potentially relevant to the relationship between PM_{2.5} and infant mortality, including household and individual characteristics, such as child sex, birth order, age of the mother, education of the mother, type of cooking fuel used at home and our asset-based wealth index, as well as time-varying climate variables such as temperature and precipitation. We do not include wealth as a control in our main results because the information is not available for the full sample, but show that results are unchanged upon its inclusion (Extended Data Table 1b). μ_c , δ_t and η_{nm} are DHS cluster, birth year and country–month effects, respectively. The cluster effects control for time-invariant cross-village differences (for example, higher or lower average mortality levels), year effects control flexibly for trends or abrupt shocks common to all locations (for example, macroeconomic shocks or declines in mortality over time), and country–month effects control for seasonality in infant mortality and PM_{2.5} exposure. In order to make nationally representative survey data representative of the entire 30-country sample, we follow a previous publication³⁶ and weight observations by the product of country-specific household survey weights (supplied by DHS) and country population weights; however, our results are insensitive to dropping the weights.

GBD estimates suggest that $f(\cdot)$ is nonlinear, with marginal effects of PM_{2.5} exposure declining at higher exposure levels⁴. To explore potential nonlinear responses to PM_{2.5} in our data, we estimate flexible versions of $f(\cdot)$, including higher-order polynomials and restricted cubic splines. However, we find that flexible models for post-birth exposure provide roughly the same shaped response function as simple linear models (Extended Data Fig. 3), and that higher-order polynomial terms for the post-birth period are not statistically significant in the full sample (Extended Data Table 1). We therefore model post-birth PM_{2.5} linearly in our main specification and in the calculation of attributable deaths; including higher-order terms generally steepens the relationship (Extended Data Fig. 3) and yields higher attributable death estimates. The quadratic term for the 9-month in utero period is statistically significant, however, and thus our main specification adopts $f(\cdot)$ quadratic in in utero PM_{2.5} exposure (PM^b) and linear in post-birth PM_{2.5} exposure (PM^a):

$$f(\text{PM}_{icnt}^b, \text{PM}_{icnt}^a) = \beta_1 \text{PM}_{icnt}^b + \beta_2 (\text{PM}_{icnt}^b)^2 + \beta_3 \text{PM}_{icnt}^a \quad (2)$$

To study whether the impacts of PM_{2.5} change over time or by wealth level, we interact linearized post-birth exposure with dummy variables for wealth quantile or for year of survey:

$$y_{icnt} = \sum_d \beta_d (I_d \text{PM}_{icnt}^a) + \mu_c + \delta_t + \eta_{nm} + \varepsilon_{icnt} \quad (3)$$

where I_d is a dummy variable for whether observation i falls into bin d . The β_d coefficients provide the marginal effect of a 1 µg m⁻³ separately for each bin (wealth quantile or time period). For the wealth estimates, we focus on terciles of the wealth index.

The goal of the fixed effects in equations (1) and (3) is to isolate variation in PM_{2.5} exposure from other time-invariant, seasonally varying or time-trending factors that could be correlated with mortality. For instance, by including cluster-fixed effects and thus using only within-village deviations in PM_{2.5} and mortality over time, our approach accounts for time invariant unobservables that could be correlated with both PM_{2.5} exposure and mortality risk at the cluster level (for example, if villages with lower PM_{2.5} exposure also happen to be wealthier). Because we observe more than one birth for most mothers, our data allow an even more stringent test on the potential role of time-invariant household unobservables. In particular, we can include a mother fixed effect in equation (1) (the cluster fixed effects thus drop out), and in this design, the effects of PM_{2.5} exposure on mortality derive from comparing whether a child born to a given mother during a period of high PM_{2.5} exposure is more or less likely to survive relative to a child born to that same mother during a period of lower PM_{2.5} exposure. This within-mother variation eliminates the common concern in pollution exposure studies that households with different levels of pollution exposure could be inherently different in unobservable ways (for instance if wealthier, lower-mortality mothers choose to live in areas of lower pollution exposure). Our results using mother fixed effects are very similar to results using cluster fixed effects (Extended Data Fig. 3), again providing strong evidence that our results are not being driven by time-invariant

unobservables. Whether or not we control directly for a broader set of individual and household characteristics X_{control} also does not change our results (Extended Data Table 1).

Similarly, overall trends in mortality and $\text{PM}_{2.5}$ exposure over time are taken out by the year fixed effects (and in robustness checks, by time trends or country-by-year fixed effects, see Extended Data Fig. 3), helping to reduce concerns that the effects of $\text{PM}_{2.5}$ that we estimate are driven by common time-trending unobservables. However, common time effects would not account for local time-varying factors that could be correlated with both $\text{PM}_{2.5}$ and mortality. Of particular importance is rainfall: rainfall reduces $\text{PM}_{2.5}$ in the atmosphere, and rainfall could be plausibly negatively or positively correlated with mortality—positively if higher rainfall led to favourable conditions for transmission of vector-borne disease (for example, malaria) or negatively if higher rainfall led to greater local food availability and thus lower mortality. To account for this possibility, we control directly for meteorological conditions using high-resolution remote-sensing based gridded data on precipitation and temperature^{37,38}, and find that our results are unaffected (Extended Data Table 1).

A related concern is that local economic activity could be associated with both local $\text{PM}_{2.5}$ levels and mortality. We believe that this is less important in our setting for two reasons. First, for large parts of our sample (particularly in West Africa), much of the variation in $\text{PM}_{2.5}$ is driven by wind-borne dust, which is unrelated to local economic activity. Second, because economic growth is most likely associated with both higher $\text{PM}_{2.5}$ levels (from industrial or agricultural activities) and lower infant mortality³⁹, then this would bias our results towards zero.

A final concern is that although equation (1) identifies impacts using variation over time in $\text{PM}_{2.5}$ exposure that is plausibly orthogonal to other determinants of mortality, individuals over the long run might adapt to differing levels of average pollution exposure in a way that is not picked up in a time series—for example, they might undertake defensive investments or learn how to limit exposures or reduce their consequence. Panel models that use inter-annual $\text{PM}_{2.5}$ variation might then overstate the harm of $\text{PM}_{2.5}$ exposure, because variation around local averages is harder to anticipate and adapt to. Although cross-sectional models that relate location-average mortality to location-average pollution exposure are subject to bias concerns from omitted variables and are considered unreliable for estimating causal effects, they arguably have the benefit of accounting for general forms of longer run adaption. We find that panel and cross-sectional models indicate surprisingly similar responses of mortality to pollution (Extended Data Fig. 4a), suggesting limited adaptation over the longer run.

Calculating relative risk and excess deaths attributable to pollution. Using the full $y(\cdot)$ function estimated in equation (1), we calculate the relative risk (RR) at a given $\text{PM}_{2.5}$ exposure level z as the predicted values from the full model evaluated at $\text{PM}_{2.5} = z$, divided by the predicted values from the full model evaluated at $\text{PM}_{2.5} = 2$:

$$\text{RR}(z) = \frac{y(z)}{y(2)} \quad (4)$$

where $2 \mu\text{g m}^{-3}$ represents the minimum exposure level observed in our data. Our approach to defining the lower bound for risk is thus similar to the approach in the GBD⁴, who define the lower bound in their relative risk curve as the minimum $\text{PM}_{2.5}$ exposure level observed in a constituent cohort study ($z = 5.8 \mu\text{g m}^{-3}$). Although estimates are imprecise at very low exposure levels due to limited sample sizes, both flexible splines and piecewise linear functions suggest that, in our data, mortality is increasing with $\text{PM}_{2.5}$ even at the lowest observed exposure levels in our data, and for this reason we set our ‘reference’ risk level to $z = 2 \mu\text{g m}^{-3}$.

To calculate relative risk across the entire range of observed $\text{PM}_{2.5}$ levels in our data and for all geographical locations, relative risk is calculated for every birth observation at its observed post-birth $\text{PM}_{2.5}$ concentration and then averaged to the cluster level. To calculate the relative risk curve in Fig. 3, we divide the data into $5 \mu\text{g m}^{-3}$ $\text{PM}_{2.5}$ bins, calculate the average relative risk within each bin, and then fit a flexible locally weighted polynomial to these estimates. Confidence intervals are obtained by bootstrapping equation (1) 1,000 times, sampling clusters with replacement, and recalculating equation (4) for each bootstrapped $y(\cdot)$. The 5–95th confidence interval is then the 5th and 95th percentiles of these 1,000 estimates at each point in the $\text{PM}_{2.5}$ distribution. Measurement error in our outdoor $\text{PM}_{2.5}$ measures, which is estimated to be roughly classical⁵, will lead to attenuation bias in our coefficient estimates in equation (1), and thus mean that our relative risk curve is biased towards zero. We discuss an alternate source of non-classical measurement error—the error related to unobserved indoor air pollution exposure—in ‘Indoor versus outdoor $\text{PM}_{2.5}$ ’.

We calculate the share of infant deaths attributable to $\text{PM}_{2.5}$ exposure in each DHS location i as:

$$S_i = 1 - \frac{y(2)}{y(z_i)} = 1 - \frac{1}{\text{RR}_i} \quad (5)$$

The average share of $\text{PM}_{2.5}$ -attributable deaths across the sample is then calculated as the population-weighted average across DHS locations, using high-resolution gridded data on birth counts from WorldPop⁴⁰ as weights. For each country, WorldPop produces a $100 \times 100 \text{ m}^2$ grid of birth counts that, when aggregated, are consistent with UN estimated country totals.

To generate the country-wide surfaces shown in Fig. 4, we apply the relative risk curve in Fig. 3 to all locations in our sample countries, using grid-level observed $\text{PM}_{2.5}$ levels in 2015. Mean exposure levels in 2015 were $30 \mu\text{g m}^{-3}$, or about $5 \mu\text{g m}^{-3}$ higher than overall sample average exposure, and thus the share of attributable deaths in 2015 shown in Fig. 4 is a little above the 22% average that we calculate for the full sample.

Finally, to calculate the total additional infant deaths attributable to $\text{PM}_{2.5}$ in 2015, we calculate for each location i :

$$\text{ED}_i = B_i \text{IMR}_i S_i \quad (6)$$

where B_i is the estimated number of births in location i in 2015 from WorldPop, IMR_i is the estimated average infant mortality rate (IMR) between 2005 and 2015 as calculated by applying previously published methods¹¹ to the more recent infant mortality data used in this study (map shown in Fig. 1c), and S_i is the share of mortality attributable to $\text{PM}_{2.5}$ as calculated above. The total attributable deaths across our sample countries in 2015 is then the sum of ED_i over all locations. Confidence intervals are calculated as above by recalculating ED_i across bootstrapped estimates of S_i .

Comparison to GBD. In a recent study⁴, a global integrated response function was derived that relates ambient $\text{PM}_{2.5}$ exposure to the relative risk of acute lower respiratory infection in infants (reproduced in Extended Data Fig. 1a). (A recent update² to our knowledge did not provide age-specific response functions to which our estimates can be compared.) To develop the global relative risk estimates, the GBD authors relied on the available literature (see the previous study⁴ for datasets and references) at the time, which consisted of: (i) 4 studies that measured the effect of ambient exposures on health outcomes, all from developed countries and with average ambient exposures below our African sample median (Extended Data Fig. 1a); (ii) 23 studies that measured the effect of second-hand smoking on health outcomes, all of which were assigned the same ‘ambient’ exposure level of $50 \mu\text{g m}^{-3}$, because true exposures were unobserved; and (iii) 1 study of household carbon monoxide exposure on child respiratory outcomes in Guatemala, for which $\text{PM}_{2.5}$ exposures had to be inferred for a large proportion of the sample, and for which counterfactual (minimum) ambient exposures were substantially higher than in all ambient studies. As shown in Extended Data Fig. 1b, of these 28 studies, 8 were in developing countries, and only one in Africa, and the median study sample was $n = 1,250$ individuals.

In comparison, our study (i) observes nearly one million individuals, more than the combined sample size of the 28 studies of acute lower respiratory infection described in the previous study⁴; (ii) directly studies the effect of ambient exposure on health outcomes in a developing country setting using quasi-experimental variation in $\text{PM}_{2.5}$ exposure to estimate health effects; (iii) uses a single empirical approach and data source to estimate a relative risk function across a broad range of $\text{PM}_{2.5}$ exposures, meaning that differences in measured responses across exposure levels cannot be attributed to differences in empirical approach or study design in different locations; and (iv) in keeping with growing literature suggesting non-respiratory effects, does not assume that the only pathway linking $\text{PM}_{2.5}$ to overall health outcomes is respiratory infection.

Differences in estimated relative risk functions between our study and Burnett et al. are shown in Fig. 3 and discussed in the main text. Given the methodological and locational differences between our study and the previous study⁴, we emphasize that differences in relative risk between the two studies at specific exposure levels cannot be interpreted as providing evidence on (for example) the relative damages caused by ambient $\text{PM}_{2.5}$ exposure versus exposure to second-hand smoke or indoor air pollution.

To compare our estimates of attributable deaths to those of the GBD, we recalculate equation (5) using the relative risk function published previously⁴, using the same grid-level $\text{PM}_{2.5}$ and population numbers that we used to generate our attribution estimates, but keeping the previously published⁴ counterfactual exposure of $5.8 \mu\text{g m}^{-3}$. Using the previously published relative risk function⁴ and this higher counterfactual exposure, we estimate that 13% of infant deaths in our sample are attributable to $\text{PM}_{2.5}$ exposure. This estimate is contained within the confidence interval for our main estimate of 22% (9–35%).

An important difference between our attribution calculations and previous study⁴ is that the latter uses a counterfactual $\text{PM}_{2.5}$ exposure (that is, theoretical minimum risk exposure level) of $5.8\text{--}8.0 \mu\text{g m}^{-3}$, which were the lowest and fifth percentiles of exposure in their reference studies (although we

note that a recent update by the GBD team² now uses $2.4 \mu\text{g m}^{-3}$ as the counterfactual exposure). As noted in the previous study⁴, these thresholds were chosen based on “(a) the availability of convincing evidence from epidemiologic studies that support a continuous reduction in risk of disease to the chosen distribution, and (b) a distribution that is theoretically possible at the population level.” We follow this logic and set our lower threshold at $2 \mu\text{g m}^{-3}$, as this is the minimum exposure level measured in our sample, and because we find evidence of linear effects of $\text{PM}_{2.5}$ on mortality at levels below $10 \mu\text{g m}^{-3}$ (Extended Data Fig. 4a). However, our choice of a lower counterfactual exposure, while arguably appropriate in our setting, could generate some of the difference that is observed between our attribution estimates and GBD's. To understand the importance of this choice, we re-calculate our population-weighted attributable deaths under different counterfactual exposures, from $2 \mu\text{g m}^{-3}$ up to $10 \mu\text{g m}^{-3}$ (Extended Data Fig. 7). We find that the share of infant deaths attributable to $\text{PM}_{2.5}$ exposure ranges from 22% with our counterfactual of $2 \mu\text{g m}^{-3}$ to 13% with a counterfactual of $10 \mu\text{g m}^{-3}$. Using previously published⁴ minimum counterfactual of $5.8 \mu\text{g m}^{-3}$, our estimate of attributable deaths becomes 18%, compared to the 13% that we would calculate using the relative risk function of the previous study⁴ (as described above).

In order to calculate a comparable estimate of additional deaths attributable to air pollution in 2015 from the GBD studies, data were downloaded from the GBD results tool⁶ for 2015 and infant deaths attributable to air pollution were summed over the 30 countries in our sample. The ‘air pollution’ category in GBD includes the categories ‘ambient particulate matter pollution’, ‘ambient ozone pollution’, and ‘household air pollution from solid fuels’. We compare our estimates of $\text{PM}_{2.5}$ -attributable death to GBD estimates of the total overall ‘air pollution’-attributable death estimates, which will be an upper bound on GBD-attributed deaths from $\text{PM}_{2.5}$ exposures specifically, and thus provide the most conservative possible comparison. This comparison is also preferable on physical grounds: it is both statistically and functionally difficult to distinguish indoor and outdoor air pollution exposures in rural biomass-burning regions, especially since most of the $\text{PM}_{2.5}$ from household cooking and fires makes its way outdoors⁴¹. The inclusion of ‘ambient ozone pollution’ is a small effect because ozone-related mortality⁴² is typically an order of magnitude lower than for $\text{PM}_{2.5}$, and it is also difficult to distinguish from $\text{PM}_{2.5}$ -related impacts because both are often present in local pollution.

The GBD results tool indicates air pollution-attributable neonatal deaths of 126,000 in our 30 countries in 2015 (range = 73,000–198,000), 150,000 in all of SSA (118,000–184,000) and 294,000 globally (234,000–350,000). Our estimate of 449,000 attributable deaths in 2015 is thus $3.6\times$ higher than the GBD estimate for the same countries. Revising the attributable death estimates upward in our 30 countries would result in an additional attributable 323,000 deaths, which would represent a more than doubling of the global estimated attributable deaths to air pollution.

Alternatively, we can apply the approach described in equation (6) to the integrated response curve developed previously⁴. This approach produces an estimate of 336,000 attributable deaths in 2015, closer to—but still substantially smaller than—our estimate of 449,000 attributable deaths. Our finding that lower ranges of $\text{PM}_{2.5}$ exposure are more harmful to infant health than previously thought is one of the factors that contributes to the difference in estimates. For example, only 11% of attributable deaths (38,000) estimated using previously published response curve⁴ occurred in locations with lower than median ($27 \mu\text{g m}^{-3}$) $\text{PM}_{2.5}$ exposures whereas 18% (80,000) of attributable infant deaths estimated using our methods occurred in these relatively lower $\text{PM}_{2.5}$ exposure areas.

Comparison to other health interventions. We compared the estimated effectiveness of a given reduction in $\text{PM}_{2.5}$ exposure from both our model and the previous study⁴, to the estimated effectiveness of other important health interventions based on estimates from the Lives Saved Tool¹⁹ (LiST; <http://www.livessavedtool.org/>, accessed 20 September 2017). LiST is a model designed to estimate the effect of scaling up health and nutritional interventions on child and maternal health. For each intervention of interest, LiST takes as input country-specific demographic information, cause-of-death data, current intervention coverage, and data from randomized controlled trials and quasi-experiments on intervention efficacy. It then combines this information to estimate the effectiveness (in terms of reduced mortality) of scaling the intervention to a desired level of population coverage.

We used LiST to estimate the mortality impacts of scaling the following interventions from baseline 2015 coverage rates to full (100%) coverage: vitamin A supplementation, selected vaccines (rotavirus, pneumococcal, influenza), insecticide-treated bed nets and oral rehydration solution. We did this separately for each of the 30 countries in our sample (except Tanzania, which was not included in the LiST database), using the default demographic datasets provided in LiST. Baseline population-weighted coverage in 2015 for the interventions were: vitamin A supplementation (73%), rotavirus vaccine (32%), pneumococcal vaccine (54%), influenza B vaccine (74%), insecticide-treated bed nets (54%) and oral rehydration

solution (37%). The LiST-estimated baseline infant mortality across these countries in 2015 was 56.7 deaths per 1,000 live births. Country-specific estimates of mortality reductions due to scaling each intervention to 100% population coverage were then averaged (weighting by population) to produce the overall estimates reported in Fig. 4.

We compared the LiST estimates to the estimated effect of a $5 \mu\text{g m}^{-3}$ $\text{PM}_{2.5}$ reduction, using both our model (equation (5)) and the previously published relative risk function⁴ using the approach described above. A $5 \mu\text{g m}^{-3}$ reduction is roughly equivalent to the estimated reduction in $\text{PM}_{2.5}$ induced by the 1970 Clean Air Act (CAA) in nonattainment counties in the United States; although $\text{PM}_{2.5}$ was not routinely measured in the US until the 1990s, the CAA is estimated⁴³ to have reduced total suspended particulates (TSP) by $20\text{--}25 \mu\text{g m}^{-3}$, and evidence from multiple sites in North America^{44,45} suggest roughly 25% of TSP by mass is $\text{PM}_{2.5}$, meaning the CAA led to $\text{PM}_{2.5}$ reductions on the order of $5 \mu\text{g m}^{-3}$.

We emphasize that our comparison of the effectiveness of $\text{PM}_{2.5}$ reductions to that of other health interventions abstracts from the policy, technical and/or financial realities of implementing these reductions or interventions. As with the LiST model, our purpose is rather to provide a basis for further exploration of the comparative feasibility and cost effectiveness of alternate interventions.

Indoor versus outdoor $\text{PM}_{2.5}$. One concern with our results is that although we purport to measure the relationship between outdoor air pollution and infant health, infants could also be exposed to indoor air pollution, and this unobserved exposure could bias our estimates. Here we quantify the likely sign and magnitude of the bias, and show that unobserved indoor air pollution probably leads us to underestimate the effect of outdoor air pollution.

Indoor and outdoor air pollution have traditionally been treated as distinct public health threats, with the separation largely reflecting the difference in pollution sources (as opposed to biological impact mechanisms) and the anticipated differential impacts of technologies or policy responses aimed at those sources. The main ‘indoor’ source of aerosol particulate matter in SSA is cooking using solid, unprocessed fuels such as wood, dung, agricultural residues, charcoal, and coal. Nearly three billion people worldwide still depend primarily on such fuels (rates are extremely high—85%—in our sample), and women and young children tend to be disproportionately exposed to cooking-related emissions given the gendered breakdown of domestic tasks in much of the world. This stands in contrast to outdoor $\text{PM}_{2.5}$ sources such as electric power generation, transportation, and open biomass burning, which are assumed to affect nearby populations more homogeneously.

But while these indoor and outdoor emissions sources may be distinct, the separation of exposures to those emissions is difficult. The basic connections between indoor and outdoor environments are well established: (i) much of the $\text{PM}_{2.5}$ that originates indoors is transported outdoors through chimneys, windows, and doors, meaning that in rural areas in developing regions, cooking-related emissions can actually drive outdoor concentrations⁴¹; and (ii) absent sophisticated filtering, outdoor concentrations represent the lower limit for indoor exposures, because air must be exchanged periodically. Although very few studies in SSA feature simultaneous indoor and outdoor measurements (and most measure PM_{10} and/or carbon monoxide (CO) instead of $\text{PM}_{2.5}$), they support several generalized findings: that a significant amount of cooking happens outdoors or in cooking areas separated from the rest of the house, that outdoor pollutant concentrations track indoor (cooking-related) concentrations⁴⁶, that concentrations in the immediate cooking area/cookstove plume spike much higher, with concentrations rapidly falling off with distance, and that areas elsewhere in the house can be relatively protected and personal exposures can vary widely^{47–49}.

We highlight the literature from SSA, albeit small, as much of the indoor/outdoor literature that includes cooking emissions has focused on highland areas in China, India and Central and South America, where indoor heating is a key service provided by indoor combustion and so ventilation conditions can be very different^{50,51}. These studies also highlight that concentrations within houses and nearby areas vary markedly during cooking hours and across seasons⁵². The indoor/outdoor literature from developed countries focuses on how well buildings (which must nevertheless exchange air with the outdoor environment) keep out pollutants, including $\text{PM}_{2.5}$. These studies⁵³ highlight that indoor:outdoor concentration ratios span 1 when windows are open and there is direct air exchange, as is the case in most of SSA/our DHS sample.

A perfect exposure metric would integrate indoor and outdoor exposures over time spent in the two environments (or weight an average of the two by relative time spent in each location). We do not have indoor concentration data for our study regions or individual exposure data; as noted above, very few simultaneous indoor and outdoor measurements exist anywhere, and especially in SSA. We therefore proxy for integrated $\text{PM}_{2.5}$ exposure by average ambient (outdoor) concentrations, derived from satellites, ground monitors, and chemical transport models, as described in ‘ $\text{PM}_{2.5}$ data’. Therefore although infant mortality rates

undoubtedly are a function of total integrated $PM_{2.5}$ exposure, we estimate the response only to the observable outdoor portion. Here we assess the extent to which this approximation is valid.

We can write total exposure (PM_{tot}) as a weighted average of indoor and outdoor exposures, where a_{out} and a_{in} are the fraction of time exposed to outdoor and indoor concentrations, respectively:

$$PM_{tot} = a_{out}PM_{out} + a_{in}PM_{in} \quad (7)$$

We expect that IMR is a function of total exposure (that is, indoor plus outdoor), but that the health impact of a given amount of either indoor or outdoor exposure is the same. Thus the correct exposure model, assuming linearity for instructive purposes, would be:

$$IMR = IMR_0 + \beta(a_{out}PM_{out} + a_{in}PM_{in}) \quad (8)$$

with β being the ‘true’ response of IMR to $PM_{2.5}$ exposure. Noting that $a_{out} + a_{in} = 1$, we can write:

$$IMR = IMR_0 + \beta(PM_{out} + a_{in}(PM_{in} - PM_{out})) \quad (9)$$

If we assume a general relationship between PM_{in} and PM_{out}

$$PM_{in} = \delta_0 + \delta PM_{out} \quad (10)$$

we can rewrite equation (9) as:

$$IMR = IMR_0 + \beta(PM_{out} + a_{in}(\delta_0 + \delta PM_{out} - PM_{out})) \quad (11)$$

which simplifies to:

$$IMR = (IMR_0 + \beta a_{in} \delta_0) + \beta(1 + a_{in}(\delta - 1))PM_{out} \quad (12)$$

We only observe PM_{out} and thus estimate the coefficient on the model:

$$IMR = IMR_0 + \alpha(PM_{out}) \quad (13)$$

From this regression (which is analogous to our main regression in equation (1)), we recover estimates of

$$\hat{\alpha} = \beta(1 + a_{in}(\delta - 1))$$

The key question is the extent to which $\hat{\alpha}$ diverges from the true response β . The fact that we do not observe PM_{in} , δ_0 , δ , a_{out} , or a_{in} leads to several possibilities for α . Our estimate $\hat{\alpha}$ will clearly be unbiased when $a_{in} \approx 0$, the case if the child is not exposed to indoor concentrations, either because they are not indoors, or (most practically) because cooking happens in a different location (for example, a separate kitchen) and they are effectively protected from those emissions while indoors. Although there is evidence that young children often have less exposure to indoor air pollution than others in the family (particularly adult females)⁴⁷, and some countries represented in our sample have fairly high rates of outdoor cooking, we nevertheless view it as unlikely that most children in our sample are unexposed to indoor pollution.

If $a_{in} > 0$, then the extent of bias depends on δ . Our estimates $\hat{\alpha}$ recover the true effect β when $\delta = 1$, that is, when PM_{in} scales exactly with PM_{out} (up to a constant offset δ_0). When $0 < \delta < 1$, our estimates of α understate the true effect of $PM_{2.5}$ on health; the opposite is true if $\delta > 1$. There are two ways to think about parameter δ . The first is at the household level. In a household model, δ_0 can be thought of as the time-averaged concentration from indoor emission sources, and δ can be thought of as the time-averaged steady-state balance of total $PM_{2.5}$ mass transport from outdoor-to-indoor compared to $PM_{2.5}$ mass transport from indoor-to-outdoor. Because some air must be exchanged between indoor and outdoor environments, $\delta > 0$, and although the volume of air exchanged will be equal, the total mass of $PM_{2.5}$ transported in either direction can differ. A scenario in which $\delta > 1$ means that $PM_{2.5}$ is trapped and builds up indoors; this leads to estimates of α that exceed the true β . This scenario is highly unlikely in an environment such as SSA with unsealed buildings. Instead, the more likely scenario is that $\delta \leq 1$. $\delta = 1$ would be the case in households that do not have indoor $PM_{2.5}$ sources (for example, clean cooking fuels), so indoor concentrations = outdoor concentrations, and air volumes are cycled back and forth. $\delta < 1$ for households with indoor emissions sources that are ventilated to the outdoors: the total mass transported outdoors is larger than the mass transported indoors (a higher concentration of $PM_{2.5}$ in the ventilated air than in outdoor air exchanged for it). The household-level interpretations of δ and δ_0 are summarized in Extended Data Table 2a.

For households using clean cooking fuels, there is no indoor $PM_{2.5}$ source ($\delta_0 = 0$), so the possible scenarios are the top row of Extended Data Table 2a. For houses using dirty cooking fuels ($\delta_0 > 0$, that is, there is an indoor source contributing to a steady-state indoor concentration that is unrelated to outdoor levels), the possible scenarios are the bottom row of the table. As described above, so long

as $\delta > 0$, the bias in $\hat{\alpha}$ does not depend on whether or not the household has an indoor source of $PM_{2.5}$ (for example, it uses only clean cook fuels). Cells 3 and 6 ($\delta > 1$) of Extended Data Table 2a are highly unlikely in steady-state as they would imply that $PM_{2.5}$ mass from outdoors is being transported indoors and concentrating there. Cell 1 is physically impossible without an indoor source of $PM_{2.5}$, and cell 5 would be an idiosyncratic case in which ventilated indoor $PM_{2.5}$ concentrations were exactly matched by outdoor concentrations (this might be approximately the case for cooking emissions ventilated immediately via a chimney just outside the house, where they are pulled back in again). So from a household model, we would expect $\hat{\alpha}$ to be unbiased for households with clean cooking fuels (no indoor $PM_{2.5}$ generation), and underestimated for households with indoor $PM_{2.5}$ sources.

We can evaluate this prediction in our data, given that DHS data do provide information on the use of cook fuels of households for a subset of households. We define ‘clean’ cook fuel households as those who cook with natural gas, biogas, liquefied petroleum gas (LPG) or electricity, and ‘dirty’ fuel households as those cooking with anything else. Consistent with the prediction above, point estimates suggest larger effects of $PM_{2.5}$ exposure on clean fuel households compared to dirty fuel households, although we cannot reject that the estimates are the same given the wide confidence intervals (Extended Data Fig. 5).

The second way to think about δ is at the aggregate level, or the population relationship between indoor and outdoor $PM_{2.5}$ concentrations. Intuitively, we would expect a positive correlation between indoor and outdoor $PM_{2.5}$ concentrations in aggregate, because outdoor $PM_{2.5}$ penetrates porous buildings and households ventilate indoor emissions. Bias in $\hat{\alpha}$ would only occur if that relationship changed across levels of outdoor $PM_{2.5}$ exposures. This would imply that, for example, at higher outdoor $PM_{2.5}$ concentrations, indoor $PM_{2.5}$ is less well-ventilated, or that at higher outdoor concentrations, homes are trapping and concentrating indoor $PM_{2.5}$.

To explore these aggregate-level relationships in our data, we aggregate DHS and ambient pollution data to the cluster level, restrict the sample to within the survey year (given that survey questions ask about current fuel use) and test the relationship between the percentage of households in a cluster using ‘clean’ fuels and ambient $PM_{2.5}$ levels. Results are shown in Extended Data Table 2b. We find that, on average, clusters using entirely clean cooking fuels have lower ambient $PM_{2.5}$ levels, both across the full sample and (more importantly) when restricted to only clusters with at least some clean cook fuel use. To verify that this isn’t simply result of unobserved variables—for example, locations with access to clean cooking fuels might also have reduced open biomass burning, or better (cleaner) electric power generation, or lower transportation-related $PM_{2.5}$ emissions—we conduct separate regressions on urban and rural clusters. We find no difference in the impact of cluster-level clean cooking fuel penetration on ambient levels in rural versus urban DHS clusters (where we would expect non-cooking emissions profiles to differ). This provides additional evidence that our ambient average metric proxies well for overall exposure, and that the relationship is not driven by non-cooking emissions, because cooking-related emissions that originate indoors are ultimately reflected in outdoor average concentrations.

These three pieces of evidence—that the basic physics of air flow suggest that $0 < \delta \leq 1$, that ‘clean’ fuel households have higher point estimates of $PM_{2.5}$ effects, and that outdoor concentrations appear to reflect indoor exposures—suggest that, if anything, unobserved indoor $PM_{2.5}$ exposures likely bias our main estimates down. We thus interpret our main estimates as conservative.

Nevertheless, infants could be extremely vulnerable to quick but marked spikes in indoor particulates during, for example, stove lighting. The existing literature for SSA highlights the high temporal variability in pollutant emissions from cooking, and the tremendous spatial heterogeneity in concentrations over short spatial scales⁴⁷. Simultaneous measurement of direct exposure for individuals, in addition to indoor and outdoor concentrations, is a key area for future research that would help to clarify the role of higher frequency spatiotemporal variation in concentrations/exposures, and allow for comparison of integrated (versus average) exposure.

Impact channels and effect size plausibility. One of our main results is that 22% of infant deaths in our sample can be attributed to $PM_{2.5}$ exposure above $2 \mu g m^{-3}$. This estimate is larger than current GBD estimates of the total child mortality burden of LRI in SSA; the online GBD tool²⁰ estimates that 12.8% of infant deaths in SSA in 2015 were due to LRI. The difference between estimates makes our results seem implausibly large if (i) the GBD estimates are correct and (ii) LRI are the only cause of death linking $PM_{2.5}$ exposure and infant mortality.

As a first point, we note that if the GBD counterfactual $PM_{2.5}$ concentration of $5.8 \mu g m^{-3}$ is used to define our ‘clear air’ baseline, then our estimate of attributable deaths would be reduced to 18%. So different clear-air counterfactuals could result in some of the difference between our estimate and GBDs.

Second, when we use the exposure–response function that was the basis of the 2015 GBD estimates⁴, we calculate for our SSA sample that 13% of infant deaths are attributable to $PM_{2.5}$ exposure (see above). This estimate is within our estimated 95% confidence interval and suggests that even in the GBD, LRI alone cannot

account for the mortality attributable to $PM_{2.5}$, unless we are willing to attribute all deaths from LRI to $PM_{2.5}$, an unlikely scenario given the multiple agents that cause LRI in young children. In other words, even the GBD data appear to indicate that $PM_{2.5}$ affects mortality through causes beyond LRI.

Third, we can directly compare our main results to quasi-experimental studies from wealthy and middle-income countries that also measured the impact of air quality on infant mortality. We are aware of six quasi-experimental studies^{24–29} that measured the impact of longer-run $PM_{2.5}$ exposure on infant mortality for which we were able to calculate effect sizes in units comparable to ours. As many of these studies reported effect sizes for TSP or PM_{10} , comparing effect sizes requires translating units to $PM_{2.5}$. We used the following conversions: $PM_{2.5} = 0.7PM_{10}$ and $PM_{10} = 0.5TSP$. We note that the results of one study²⁷ are for 0–4 year olds, and baseline mortality rates were not reported in the paper for this study. For another study²⁸, we used the results from table 4, dividing the effect of traffic on mortality by the effect of traffic on $PM_{2.5}$.

Results are shown in Extended Data Fig. 8j. Our results are closest to the previously published results for Mexico City²⁶, with both studies estimating an approximately 9% increase in infant mortality per $10 \mu g m^{-3}$ change in $PM_{2.5}$. Our results are a little smaller than previously published results published in the US^{24,25} and substantially smaller than recent estimates from urban China, Turkey, and California^{27–29}.

Fourth, there is also growing evidence from both developing and developed countries that $PM_{2.5}$ exposure increases infant mortality risk through causes other than LRI. In particular, multiple studies (including meta-analyses) find that in utero exposure to $PM_{2.5}$ increases the incidence of pre-term birth and low birth weight^{21–23,54–57}, the leading risk factors for neonatal mortality (defined as death in the first 28 days of life). Although the epidemiological evidence leading from $PM_{2.5}$ exposure to adverse fetal outcomes is now substantial, the basic mechanisms are not fully understood and are thought to include pro-inflammatory effects, endocrine effects, neurophysiological effects, metabolic effects, increasing oxidative stress in both the mother and fetus and disruption of oxygen flow^{22,23,58–61}. Studies from both developing and developed countries also directly show impacts of in utero $PM_{2.5}$ exposure on neonatal mortality^{24,25,62}. Although LRIs cause some deaths among neonates, the majority of neonatal mortality is distinct from LRI²⁰. Neonatal mortality thus represents an important channel linking $PM_{2.5}$ exposure and infant mortality that is largely distinct from LRI, and we test for this mechanism in our data below.

Other work has linked $PM_{2.5}$ exposure to low child height-for-age (that is, stunting)⁶³, an outcome that is primarily reflective of long-term malnutrition⁶⁴ rather than LRI specifically, and which is the leading risk factor for child mortality in Africa. Evidence of a $PM_{2.5}$ –stunting relationship in our data would again be consistent with $PM_{2.5}$ having mortality-relevant health impacts beyond LRI. A few studies also link poor air quality to incidence of diarrhoea^{65,66}. Diarrhoea is a leading cause of infant death in Africa and is mostly considered to be distinct from LRI, even if some conditions, such as poor hygienic conditions, predispose children to both illnesses⁶⁷. Although the mechanisms linking $PM_{2.5}$ exposure and diarrhoeal illness are poorly characterized—for instance it is thought that biological $PM_{2.5}$ could be associated with increased diarrhoea⁶⁶—we nevertheless consider this yet another possible pathway leading from $PM_{2.5}$ exposure to infant mortality that is distinct from LRI, and evidence of this pathway in our data would provide further support for non-LRI effects.

Although the cause of death is unobserved in our data, we use the same empirical strategy as our main analysis to test the extent to which $PM_{2.5}$ levels are related to adverse childhood outcomes that could affect mortality separately from LRI, including neonatal mortality, birth weight, birth size, stunting and diarrhoeal illness. In order to control for differences across mothers (for example, in their ability to estimate birth size), we include mother fixed effects in all specifications so the estimated effects come from comparing outcomes between children born to the same mother at times of different $PM_{2.5}$ concentrations. The point of these analyses is not to isolate the specific causes of death related to $PM_{2.5}$ exposure—this level of ascertainment is not available in DHS, given that DHS does not collect autopsy data—but to provide evidence that $PM_{2.5}$ exposure harms infant health by increasing the risks of several adverse conditions beyond LRI. As shown in Extended Data Fig. 8, we find substantial evidence of non-LRI pathways, as described below. *In utero $PM_{2.5}$ exposure is associated with lower birth weights.* We use all observations in our data for which we have birth weight (directly recorded in grams from birth cards and transcribed by DHS enumerators, $n = 215,975$) or birth size recalled by interviewed mothers (five-point scale from ‘very small’ to ‘very large’, $n = 454,444$). Consistent with existing literature on the relationship between PM exposure and birth weight, we find suggestive evidence that higher in utero $PM_{2.5}$ exposure is associated with lower birth weight and birth size, although estimates are imprecise (particularly for the birth card measurements) given the reduced sample. *In utero $PM_{2.5}$ exposure is associated with higher neonatal mortality.* Consistent with the epidemiologic evidence that links $PM_{2.5}$ exposure to low birth weight, and

our findings that in utero $PM_{2.5}$ exposure is associated with reduced birth weight and birth size, we find a strong positive link between in utero $PM_{2.5}$ exposure and neonatal mortality (Extended Data Fig. 8c). At the mean exposure in our sample ($25 \mu g m^{-3}$), we find a similar $PM_{2.5}$ –mortality relationship as in our post-birth results. However, because baseline neonatal mortality rates are less than half infant mortality rates, this effect translates into a larger percentage change. We find a $10 \mu g m^{-3}$ $PM_{2.5}$ increase associated with an approximately 22% (95% confidence interval, 3–41%) increase in neonatal mortality at mean exposure; the relationship flattens out at higher exposure levels. Overall, the evidence that links in utero $PM_{2.5}$ exposure to birth weight outcomes and neonatal mortality strongly suggests a meaningful non-LRI pathway from $PM_{2.5}$ exposure to infant mortality.

Post-birth exposure to $PM_{2.5}$ reduces child height-for-age among surviving children. Anthropometric measurements are routinely performed on children under five years of age in the majority of DHS surveys. This allows estimation of malnutrition among living children present during the household interviews. Malnutrition is the largest risk factor for mortality in children under five years of age among children in SSA, responsible for over one million deaths of children under-5 from neonatal disorders, diarrhoeal illness, LRI and other common communicable diseases²⁰. Low height-for-age (stunting), in particular, is a reflection of long-term malnutrition. We find that post-birth $PM_{2.5}$ exposure in the first year is strongly associated with reduced height-for-age at the time of survey (Extended Data Fig. 8d). To the extent that high $PM_{2.5}$ represents unhealthy environments for child growth, it could increase the risk of stunting. Stunting, in turn, represents another pathway linking $PM_{2.5}$ exposure with infant mortality that is unrelated to LRI. We note that the sample for this analysis only includes children who survived to the date of the interview, and thus undercounts children who died before the survey and who may be smaller in size; this may lead us to understate the ‘true’ effect of $PM_{2.5}$ on stunting in this analysis.

Post-birth $PM_{2.5}$ exposure is associated with increase diarrhoeal incidence among surviving children. Interviewed mothers are asked whether or not each living child born in the past 3–5 years had diarrhoea in the two weeks before the survey. We again find a positive relationship between post-birth $PM_{2.5}$ exposure in the first year and the probability of experiencing diarrhoea within the two weeks preceding the survey (Extended Data Fig. 8e). We interpret this as additional evidence of a potential pathway beyond LRI, and one that should be investigated further.

Post-birth $PM_{2.5}$ exposure has no effect on child sex, the likelihood of a multiple birth or bed net usage. These measures are the three most easily observable child-level variables in our data (Extended Data Fig. 8f–h). These are placebo tests to ensure that our main effects are not spurious and that $PM_{2.5}$ exposure is uncorrelated with variables that should not be affected by air quality.

$PM_{2.5}$ exposure 13–24 months after birth does not predict mortality in the first 12 months after birth. As another placebo test, we confirm that exposure after the first birthday of a child does not affect the probability of dying before the first birthday (Extended Data Fig. 8i). This again is evidence that our main effects are not spurious.

The balance of evidence thus strongly suggests that $PM_{2.5}$ exposure is associated with infant survival through channels other than respiratory infection, and lends plausibility to our main results. For instance, the GBD estimates that LRI (12%), diarrhoeal illnesses (11%) and neonatal disorders (27%) make up a combined 50% of all deaths of children under five in SSA²⁰. Given evidence that $PM_{2.5}$ is associated with all of these channels, and the fact that our estimates are even smaller than comparable estimates from the US, our estimate that approximately 20% of infant mortality could be attributed to $PM_{2.5}$ exposure is plausible. Nevertheless, we view prospective epidemiological studies that measure both exposures and intermediary outcomes as critical in building a more complete understanding of causal pathways going forward; such studies would shed critical light on the large overall estimates that we provide here.


Reporting summary. Further information on experimental design is available in the Nature Research Reporting Summary linked to this paper.

Code availability. All codes that support the findings of this study are available at <https://purl.stanford.edu/qt056zr6479>.

Data availability. All data and code that support the findings of this study are available at <https://purl.stanford.edu/qt056zr6479>.



31. ICF. The DHS program, Data. <http://dhsprogram.com/data> (1998).
32. Bendavid, E. Is health aid reaching the poor? Analysis of household data from aid recipient countries. *PLoS ONE* **9**, e84025 (2014).
33. Lee, H. J., Coull, B. A., Bell, M. L. & Koutrakis, P. Use of satellite-based aerosol optical depth and spatial clustering to predict ambient $PM_{2.5}$ concentrations. *Environ. Res.* **118**, 8–15 (2012).
34. Hyder, A. et al. $PM_{2.5}$ exposure and birth outcomes: use of satellite- and monitor-based data. *Epidemiology* **25**, 58–67 (2014).
35. 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**, 18916 (2016).

36. Burke, M., Gong, E. & Jones, K. Income shocks and HIV in Africa. *Econ. J.* **125**, 1157–1189 (2015).
37. Funk, C. et al. The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes. *Sci. Data* **2**, 150066 (2015).
38. LP DAAC. MODIS/aqua land surface temperature and emissivity (LST/E). version 6. *USGS Earth Resources Observation And Science Center*. <https://lpdaac.usgs.gov> (NASA, 2016).
39. Baird, S., Friedman, J. & Schady, N. Aggregate income shocks and infant mortality in the developing world. *Rev. Econ. Stat.* **93**, 847–856 (2011).
40. Tatem, A. J. et al. Millennium development health metrics: where do Africa's children and women of childbearing age live? *Popul. Health Metr.* **11**, 11 (2013).
41. Rehman, I., Ahmed, T., Praveen, P., Kar, A. & Ramanathan, V. Black carbon emissions from biomass and fossil fuels in rural India. *Atmos. Chem. Phys.* **11**, 7289–7299 (2011).
42. Anenberg, S. C., Horowitz, L. W., Tong, D. Q. & West, J. J. An estimate of the global burden of anthropogenic ozone and fine particulate matter on premature human mortality using atmospheric modeling. *Environ. Health Perspect.* **118**, 1189–1195 (2010).
43. Isen, A., Rossin-Slater, M. & Walker, W. R. Every breath you take, every dollar you'll make: the long-term consequences of the clean air act of 1970. *J. Polit. Econ.* **125**, 848–902 (2017).
44. Wilson, W. E. & Suh, H. H. Fine particles and coarse particles: concentration relationships relevant to epidemiologic studies. *J. Air Waste Manag. Assoc.* **47**, 1238–1249 (1997).
45. Brook, J. R., Dann, T. F. & Burnett, R. T. The relationship among TSP, PM₁₀, PM_{2.5}, and inorganic constituents of atmospheric particulate matter at multiple Canadian locations. *J. Air Waste Manag. Assoc.* **47**, 2–19 (1997).
46. Dionisio, K. L. et al. Measuring the exposure of infants and children to indoor air pollution from biomass fuels in the Gambia. *Indoor Air* **18**, 317–327 (2008).
47. Ezzati, M., Saleh, H. & Kammen, D. M. The contributions of emissions and spatial microenvironments to exposure to indoor air pollution from biomass combustion in Kenya. *Environ. Health Perspect.* **108**, 833–839 (2000).
48. Kilabuko, J. H., Matsuki, H. & Nakai, S. Air quality and acute respiratory illness in biomass fuel using homes in Bagamoyo, Tanzania. *Int. J. Environ. Res. Public Health* **4**, 39–44 (2007).
49. Yamamoto, S. S., Louis, V. R., Sié, A. & Sauerborn, R. Biomass smoke in Burkina Faso: what is the relationship between particulate matter, carbon monoxide, and kitchen characteristics? *Environ. Sci. Pollut. Res.* **21**, 2581–2591 (2014).
50. Fischer, S. L. & Koshland, C. P. Daily and peak 1 h indoor air pollution and driving factors in a rural Chinese village. *Environ. Sci. Technol.* **41**, 3121–3126 (2007).
51. Ni, K. et al. Seasonal variation in outdoor, indoor, and personal air pollution exposures of women using wood stoves in the Tibetan Plateau: baseline assessment for an energy intervention study. *Environ. Int.* **94**, 449–457 (2016).
52. Parikh, J., Balakrishnan, K., Laxmi, V. & Biswas, H. Exposure from cooking with biofuels: pollution monitoring and analysis for rural Tamil Nadu, India. *Energy* **26**, 949–962 (2001).
53. Cyrus, J., Pitz, M., Bischof, W., Wichmann, H.-E. & Heinrich, J. Relationship between indoor and outdoor levels of fine particle mass, particle number concentrations and black smoke under different ventilation conditions. *J. Expo. Anal. Environ. Epidemiol.* **14**, 275–283 (2004).
54. Pedersen, M. et al. Ambient air pollution and low birthweight: a European cohort study (ESCAPE). *Lancet Respir. Med.* **1**, 695–704 (2013).
55. Chang, H. H., Warren, J. L., Darrow, L. A., Reich, B. J. & Waller, L. A. Assessment of critical exposure and outcome windows in time-to-event analysis with application to air pollution and preterm birth study. *Biostatistics* **16**, 509–521 (2015).
56. Li, X. et al. Association between ambient fine particulate matter and preterm birth or term low birth weight: an updated systematic review and meta-analysis. *Environ. Pollut.* **227**, 596–605 (2017).
57. Blum, J. L., Chen, L.-C. & Zelikoff, J. T. Exposure to ambient particulate matter during specific gestational periods produces adverse obstetric consequences in mice. *Environ. Health Perspect.* **125**, 077020 (2017).
58. Kampa, M. & Castanas, E. Human health effects of air pollution. *Environ. Pollut.* **151**, 362–367 (2008).
59. Slama, R. et al. Meeting report: atmospheric pollution and human reproduction. *Environ. Health Perspect.* **116**, 791–798 (2008).
60. Jerrett, M. et al. Traffic-related air pollution and obesity formation in children: a longitudinal, multilevel analysis. *Environ. Health* **13**, 49 (2014).
61. Suades-González, E., Gascon, M., Guxens, M. & Sunyer, J. Air pollution and neuropsychological development: a review of the latest evidence. *Endocrinology* **156**, 3473–3482 (2015).
62. Jayachandran, S. Air quality and early-life mortality: evidence from Indonesia's wildfires. *J. Hum. Resour.* **44**, 916–954 (2009).
63. Tielsch, J. M. et al. Exposure to indoor biomass fuel and tobacco smoke and risk of adverse reproductive outcomes, mortality, respiratory morbidity and growth among newborn infants in south India. *Int. J. Epidemiol.* **38**, 1351–1363 (2009).
64. Black, R. E. et al. Maternal and child undernutrition: global and regional exposures and health consequences. *Lancet* **371**, 243–260 (2008).
65. Malilay, J., Mariana, G. R., Ramirez Vanegas, A., Non, E. & Sinks, T. Public health surveillance after a volcanic eruption: lessons from Cerro Negro, Nicaragua, 1992. *Bull. Pan. Am. Health. Organ.* **30**, 218–226 (1996).
66. Menetrez, M. Y. et al. An evaluation of indoor and outdoor biological particulate matter. *Atmos. Environ.* **43**, 5476–5483 (2009).
67. Luby, S. P. et al. Effect of handwashing on child health: a randomised controlled trial. *Lancet* **366**, 225–233 (2005).







Extended Data Fig. 1 | Integrated exposure risk curve estimated by the GBD project. Data were obtained from a previous study⁴. **a**, Relative risk curve representing the risk from acute lower respiratory infections in infants (obtained from figure 2 of Burnett et al.⁴). The curve combines point estimates from ambient air pollution (AAP) studies, indoor air pollution (HAP) studies and second-hand smoking (SHS) studies to derive risk responses across the $PM_{2.5}$ exposure distribution. The histograms show the share of population exposed to different long-run (15-year average) ambient $PM_{2.5}$ concentrations in North American and Europe where most GBD studies took place, in SSA countries in our sample,

and globally. In total, 49% of the overall population in Africa, and 51% globally, live in areas with ambient pollution concentrations exceeding the maximum ambient $PM_{2.5}$ concentration from the GBD study ($25 \mu g m^{-3}$). **b**, Most studies used to estimate the GBD integrated exposure response⁴ were carried out in North America or Europe, with the exception of a household air pollution study in Guatemala and second-hand smoking studies in Vietnam, India and South Africa. Median sample size (depicted by marker size in the plot) across these studies is $n = 1,250$. Country outlines were obtained from Global Administrative Areas, version 2.0³⁰.

a DHS cluster locations**b** Number of births observed by year**c** Study regions


Extended Data Fig. 2 | Overview of birth data from DHS surveys and study regions in Africa. **a**, Location of DHS clusters included in our sample. **b**, The number of births observed in each year in our sample. More births are observed in earlier years because births are recalled in the surveys so each new survey round potentially adds births from all previous years. **c**, Regional categorization of countries, for regional analysis in Fig. 2c. Sample countries assigned to West Africa region are Benin,

Burkina Faso, Ivory Coast, Ghana, Guinea, Liberia, Mali, Nigeria, Senegal, Sierra Leone and Togo. Sample countries assigned to 'rest of Africa' are Angola, Burundi, Cameroon, Comoros, DRC, Ethiopia, Gabon, Kenya, Lesotho, Madagascar, Malawi, Mozambique, Namibia, Rwanda, Swaziland, Uganda, Zambia and Zimbabwe. Country outlines were obtained from Global Administrative Areas, version 2.0³⁰.


a Modeling $PM_{2.5}$ with polynomials**b** Different fixed effects**c** Modeling $PM_{2.5}$ with splines**d** Adding country time trends

Extended Data Fig. 3 | Effect of post-birth $PM_{2.5}$ exposure is robust under different regression models. Estimated responses under higher-order polynomials (a), different specifications of the fixed effects (b), restricted cubic spline functions of $PM_{2.5}$ (c) and additional time


controls (d). In each panel, the blue line and shaded region indicate the estimated baseline response shown in Fig. 2a and the bootstrapped 95% confidence interval. Splines in c have knots at $10 \mu g m^{-3}$ (single knot spline) or evenly spaced knots (three- and four-knot splines).



Extended Data Fig. 4 | Piecewise linear and cross-sectional relationships between post-birth $\text{PM}_{2.5}$ exposure and infant mortality. **a**, Piecewise linear estimates of the effect of $\text{PM}_{2.5}$ exposure below and above the WHO $\text{PM}_{2.5}$ guideline of $10 \mu\text{g}/\text{m}^3$. Shaded regions represent bootstrapped 95% confidence intervals. Slopes above and below the $10 \mu\text{g}/\text{m}^3$ threshold are very similar, although confidence intervals are wider below the threshold due to smaller sample sizes. **b**, Cross-sectional and panel models give similar estimated effects of post-birth $\text{PM}_{2.5}$




exposure on infant mortality. Blue line shows baseline panel model, orange line shows a cross-sectional model that relates cluster-average mortality to cluster-average $\text{PM}_{2.5}$ exposure. Each response function is centred at sample median exposure ($25 \mu\text{g}/\text{m}^3$). Histograms at the bottom show counts of exposure at different $\text{PM}_{2.5}$ levels, for the panel sample (blue) and cross-sectional sample (orange); cross-sectional exposures are slightly narrower given that year-to-year variation has been averaged out.




Extended Data Fig. 5 | Heterogeneous effects of post-birth $PM_{2.5}$ exposure. Effects are estimated by interacting a dummy for each modifying variable with linear $PM_{2.5}$, and are measured as the percentage change in infant mortality per $10 \mu g m^{-3}$ increase in $PM_{2.5}$ exposure,

relative to baseline mortality rates in each subgroup. Circles indicate point estimates, and whiskers the 95% confidence interval on the point estimate. The last column shows the baseline estimate from the full (uninteracted) linear model.




Extended Data Fig. 6 | Linear effect of post-birth $\text{PM}_{2.5}$ exposure by year for different time periods. Panels are the same as Fig. 2e but replicated for different time periods, showing effects in each year independently. Circles indicate point estimates, and whiskers the 95%

confidence interval on the point estimate. For each time period 2001 – year t , the sample was restricted to births between 2001 and year t and to surveys that were conducted after year t . These steps help to approximate a consistent geographical sample across the time periods.



Extended Data Fig. 7 | Effect of different assumed counterfactual $PM_{2.5}$ levels on the estimated share of infant deaths attributable to $PM_{2.5}$. Each point represents the same calculation described in the Methods, under different counterfactual minimum $PM_{2.5}$ exposure levels. Data are from Cohen et al.², Burnett et al.⁴ and the WHO guidelines¹⁶.



Extended Data Fig. 8 | Effect of $PM_{2.5}$ on non-respiratory mortality and mortality risk factors. **a–c**, Effect of in utero $PM_{2.5}$ exposure on low birth weight, low birth size as reported by mothers on a scale from 1 to 5, and neonatal mortality (NMR). **d, e**, Effect of post-birth $PM_{2.5}$ exposure on height-for-age and diarrhoeal incidence for living children. In each case, higher $PM_{2.5}$ concentrations worsen health outcomes. **f–h**, Placebo tests that relate $PM_{2.5}$ exposures to child outcomes that should be unaffected:

child sex, whether child was born in a multiple birth, and child's use of a bed net. **i**, $PM_{2.5}$ exposure in the 13–24 months after birth has no effect on mortality in the first 12 months after birth. Shaded regions represent bootstrapped 95% confidence intervals in each panel. **j**, Estimates of the effect of $PM_{2.5}$ on all-cause infant mortality from published quasi-experimental studies^{24–29}, expressed as the percentage change in the infant mortality rate per $10\mu g m^{-3}$ increase in $PM_{2.5}$.

Extended Data Table 1 | Regression results for main specification and for subsample of households with asset data

a						
	Linear Model - full sample			Quadratic Model- full sample		
	(1)	(2)	(3)	(4)	(5)	(6)
PM _{2.5}	0.000656*** (0.0001938)	0.0006412*** (0.0001984)	0.0006492*** (0.0002016)	0.0010768** (0.0005449)	0.0010592* (0.0005523)	0.0011407** (0.0005403)
PM _{2.5} -squared				-0.000003986 (0.000005000)	-0.000003953 (0.000005037)	-0.000004648 (0.000005032)
Observations	990,696	990,696	990,696	990,696	990,696	990,696
<i>Controls:</i>						
Temp & Rainfall	No	Yes	Yes	No	Yes	Yes
Individual Covariates	No	No	Yes	No	No	Yes
Asset Wealth Index	No	No	No	No	No	No
Regression Weights	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
b						
	Linear Model - wealth sample		Quadratic Model - wealth sample			
	(1)	(2)	(3)	(4)		
PM _{2.5}	0.0006912*** (0.0002086)	0.0006877 *** (0.0002089)	0.0015868 *** (0.0005826)	0.0015762 *** (0.0005824)		
PM _{2.5} -squared			-0.000008 (0.000005)	-0.000008 (0.000005)		
Observations	833,001	833,001	833,001	833,001		
<i>Controls:</i>						
Asset Wealth Index	No	Yes	No	Yes		
Temp & Rainfall	Yes	Yes	Yes	Yes		
Individual Covariates	Yes	Yes	Yes	Yes		
Regression Weights	Yes	Yes	Yes	Yes		
Fixed Effects	Yes	Yes	Yes	Yes		

a. Full sample. **b.** Subsample of households with asset data. Mortality is modelled as either a linear or quadratic function of PM_{2.5} (results from additional specifications modelling mortality as a more flexible nonlinear function of PM_{2.5} are shown in Extended Data Fig. 3). The outcome measure is a binary variable equal to one if the child did not survive until 12 months of age. The mean of the outcome variable = 0.071. Standard errors are in parentheses and are clustered at the DHS Cluster level, and asterisks denote statistical significance (two-sided): * $P < 0.10$, ** $P < 0.05$, *** $P < 0.01$.

Extended Data Table 2 | Understanding potential bias from unobserved indoor air pollution exposure

a

	$0 < \delta < 1$ mass outflow > mass inflow	$\delta = 1$ mass outflow = mass inflow	$\delta > 1$ mass inflow > mass outflow
$\delta_0 = 0$ $\Rightarrow$ No Indoor Source	[1] $PM_{in} < PM_{out}$ $\hat{\alpha}$ biased down	[2] $PM_{in} = PM_{out}$ $\hat{\alpha}$ unbiased	[3] $PM_{in} > PM_{out}$ $\hat{\alpha}$ biased up
$\delta_0 > 0$ $\Rightarrow$ Indoor Source	[4] $PM_{in} \geq PM_{out}$ $\hat{\alpha}$ biased down	[5] $PM_{in} > PM_{out}$ $\hat{\alpha}$ unbiased	[6] $PM_{in} > PM_{out}$ $\hat{\alpha}$ biased up

b

Contemporaneous Ambient PM _{2.5} Concentration						
	All DHS Clusters			Clusters w/ Clean Cooking $\neq 0$		
	(1) Full	(2) Urban	(3) Rural	(4) Full	(5) Urban	(6) Rural
Clean cooking fuel penetration rate	-3.177*** (0.347)	-4.931*** (0.403)	-2.687** (0.951)	-8.563*** (0.846)	-9.869*** (0.981)	-6.885*** (1.964)
Constant	22.94*** (0.0866)	24.51*** (0.170)	22.39*** (0.100)	27.52*** (0.669)	28.84*** (0.807)	25.35*** (1.239)
N	25483	7789	17694	2444	1952	492
R ²	0.00307	0.0180	0.000417	0.0414	0.0498	0.0249
RMSE	13.23	13.16	13.22	12.96	12.86	13.28

a. Potential bias at the household level in the estimated effect of PM_{2.5} exposure on infant health $\hat{\alpha}$ as a function of the relationship between indoor and outdoor pollution exposure. Cells show the expected relative magnitudes of time-averaged differences in PM_{in} versus PM_{out} for all combinations of δ_0 and δ . **b.** Relationship between ambient PM_{2.5} and fraction of households using clean cooking fuels at the DHS cluster level. Standard errors are in parentheses, and asterisk denote statistical significance (two-sided): * $P < 0.10$, ** $P < 0.05$, *** $P < 0.01$. RMSE, root mean squared error.

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see [Authors & Referees](#) and the [Editorial Policy Checklist](#).

Statistical parameters

When statistical analyses are reported, confirm that the following items are present in the relevant location (e.g. figure legend, table legend, main text, or Methods section).

n/a Confirmed

- ☐ ☒ The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- ☐ ☒ An indication of whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- ☐ ☒ The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- ☐ ☒ A description of all covariates tested
- ☐ ☒ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- ☐ ☒ A full description of the statistics including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- ☐ ☒ For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- ☒ ☐ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- ☒ ☐ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- ☐ ☒ Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated
- ☐ ☒ Clearly defined error bars
State explicitly what error bars represent (e.g. SD, SE, CI)

Our web collection on [statistics for biologists](#) may be useful.

Software and code

Policy information about [availability of computer code](#)

Data collection

We utilize publicly available households surveys from the Demographic and Health Surveys, and publicly available satellite data. These data were processed in Stata and R.

Data analysis

Data were analyzed using R (open source software package), and all analysis code will be made available upon publication on the Stanford Digital repository. We have included a permanent URL included in revised manuscript at the end of Methods, which links to our data and code repo and which we will make public after the embargo date.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers upon request. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

We are not allowed to re-distributed the raw household survey data from DHS, but users can access the data on the DHS website. Our data and code repo will provide instructions for accessing the data, and all replication code for reproducing our results.

Field-specific reporting

Please select the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☒ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/authors/policies/ReportingSummary-flat.pdf](https://www.nature.com/authors/policies/ReportingSummary-flat.pdf)

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	Data are quantitative observational, derived from household surveys conducted across 30 African countries.
Research sample	Research sample is children born to African mothers between the ages of 15-49 across 30 African countries, as measured by the Demographic and Health Surveys (DHS).
Sampling strategy	DHS take a stratified random sample, first randomly selecting enumeration areas with probability proportional to population size, and then randomly sampling 10-20 mothers of childbearing age within each enumeration area. The sample in each country is meant to be nationally representative.
Data collection	We rely on household surveys collected by the Demographic and Health surveys, which followed the sampling protocol above. Interviews were conducted in respondent's households, and measured a range of health behaviors and outcomes, including the birth and survival of all children born to the interviewed mothers.
Timing	The surveys we utilize were carried out between the years 2001 and 2015 across 30 African countries.
Data exclusions	We use all geo-referenced African surveys conducted by DHS between 2001 (when our air quality data begin) and 2015.
Non-participation	DHS does not provide information on non-response rates among participants. We use every observation reported in the DHS surveys.
Randomization	As described in the paper, we directly controlled for household and individual covariates (mothers age, mothers education, child sex, birth order, etc), we controlled for all time invariant factors at the village level using village fixed effects, and we controlled for all common time-trending variables using year fixed effects and time trends.

Reporting for specific materials, systems and methods

Materials & experimental systems

n/a	Involved in the study
☒	☐ Unique biological materials
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology
☒	☐ Animals and other organisms
☒	☐ Human research participants

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging