

NBER WORKING PAPER SERIES

AIR POLLUTION AND INFANT MORTALITY:
EVIDENCE FROM SAHARAN DUST

Sam Heft-Neal
Jennifer Burney
Eran Bendavid
Kara Voss
Marshall Burke

Working Paper 26107
<http://www.nber.org/papers/w26107>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2019

We thank the National Science Foundation (CNH Award #1715557) for funding. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Sam Heft-Neal, Jennifer Burney, Eran Bendavid, Kara Voss, and Marshall Burke. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Air Pollution and Infant Mortality: Evidence from Saharan Dust
Sam Heft-Neal, Jennifer Burney, Eran Bendavid, Kara Voss, and Marshall Burke
NBER Working Paper No. 26107
July 2019
JEL No. O12,Q53

ABSTRACT

Accurate estimation of air quality impacts on health outcomes is critical for guiding policy choices to mitigate such damages. Estimation poses an empirical challenge, however, because local economic activity can simultaneously generate changes in both air quality and in health impacts that are independent of air quality, confounding pollution-health estimates. To address this challenge, we leverage plausibly exogenous variation in local particulate matter exposure across sub-Saharan Africa due to dust export from the Bodele Depression, a remote Saharan region responsible for a substantial share of global atmospheric dust. Large scale transport of this dust is uncorrelated with local emissions sources and allows us to isolate the causal impact of air quality on infant mortality across Sub-Saharan Africa. Combining detailed information on nearly 1 million births with satellite measures of aerosol particulate matter, we find that a 10mg/m³ increase in local ambient PM_{2.5} concentration driven by distant dust emission causes a 22% increase in infant mortality across our African sample (95% CI: 10-35%), an effect comparable to quasi-experimental pollution-infant mortality estimates from wealthier countries. We also show that rainfall over the Bodele is a significant control on PM_{2.5} export and thus child health, and that future climate-change driven changes in Saharan rainfall could generate very large impacts on African child health through this pathway alone. We calculate that seemingly exotic proposals to pump and apply groundwater to the Bodele to reduce dust emission could be cost competitive with leading interventions aimed at improving child health.

Sam Heft-Neal
Center on Food Security and the Environment
Stanford University
473 Via Ortega
Stanford, CA 94305
sheftneal@stanford.edu

Kara Voss
University of California San Diego
Scripps Institution of Oceanography
8622 Kennel Way
La Jolla, CA 92037
kvoss@ucsd.edu

Jennifer Burney
University of California San Diego
9500 Gilman Dr. #0519
La Jolla, CA 92093
jburney@ucsd.edu

Marshall Burke
Department of Earth System Science
Stanford University
Stanford, CA 94305
and NBER
mburke@stanford.edu

Eran Bendavid
Department of Medicine
Stanford University
ebd@stanford.edu

1 Introduction

Poor air quality is a known determinant of poor health outcomes, with even modest improvements in air quality likely to save millions of premature deaths annually (Shindell et al., 2018; Burnett et al., 2018). Children younger than 5 are particularly vulnerable, and impacts are thought to be largest in developing regions where exposure to high levels of ambient air pollution during childhood is estimated to reduce overall life expectancy by 4-5 years on average (Lelieveld, Haines, and Pozzer, 2018). However, while quantitative assessment of the harms from pollution exposure has greatly improved, only recently have researchers begun using quasi-experimental research designs that plausibly isolate variation in ambient pollution exposure from other correlated factors that also affect health (Cesur, Tekin, and Ulker, 2017; Arceo, Hanna, and Oliva, 2016; Chay and Greenstone, 2003a; Knittel, Miller, and Sanders, 2016; Chay and Greenstone, 2003b; He, Fan, and Zhou, 2016). Moreover, most of these estimates are concentrated in wealthy regions or limited to relatively wealthy cities in developing countries where data are available.

Direct estimation of the local health burdens of air pollution exposures requires addressing two challenges. The first is accurately measuring both exposures and responses, which has become easier with both advances in remote sensing of air quality (Van Donkelaar et al., 2016) and the accumulation of geo-referenced household-survey-based health measurements across broad geographies (Aliaga and Ren, 2006). The second challenge is disentangling pollution exposures from other correlated variables that might also directly affect health. In particular, many economic activities (e.g., transportation, industrial production, biomass clearing in agriculture) produce $PM_{2.5}$ locally but can also directly affect health outcomes by changing households' economic status. These direct effects could positively affect child health, for instance if higher incomes enable greater access or utilization of health services, or negatively affect child health if for example economically-engaged adults are less able to invest time in childrens' well-being (Miller and Urdinola, 2010).

To disentangle the impact of co-varying factors on infant mortality, researchers have utilized study designs that leverage setting-specific sources of variation in $PM_{2.5}$ likely uncorrelated with other factors that affect health outcomes (Cesur, Tekin, and Ulker, 2017; Arceo, Hanna, and Oliva, 2016; Chay and Greenstone, 2003a; Knittel, Miller, and Sanders, 2016; Chay and Greenstone, 2003b; He, Fan, and Zhou, 2016). These approaches use, for example, changes in environmental policy (Chay and Greenstone, 2003a; He, Fan, and Zhou, 2016), naturally occurring local weather phenomena (Arceo, Hanna, and Oliva,

2016), and traffic patterns (Knittel, Miller, and Sanders, 2016) as sources of plausibly exogenous variation in air quality. However, given that these approaches rely on often idiosyncratic local-level events for variation, they are difficult to apply across larger spatial and temporal scales. To build a broader understanding, recent quasi-experimental studies have used local variation in $PM_{2.5}$ derived from natural sources such as dust and argued that this approach allows for causal interpretation of $PM_{2.5}$ impacts on health across broad developing-country geographies (Heft-Neal et al., 2018; Adhvaryu et al., 2019). While using local-level variation in $PM_{2.5}$ from natural sources helps mitigate some concerns of economic activity as a confound, these studies – with the exception of Foreman (2018), discussed below – have no explicit way of distinguishing locally-sourced from distant $PM_{2.5}$. As a result, the possibility remains that existing estimates in these developing-country settings (Burnett et al., 2018; Heft-Neal et al., 2018; Adhvaryu et al., 2019), which are critical for understanding global health burdens and evaluating policy choices, could be biased by unobserved local economic activity.

Here we combine household-survey-based data on the location and timing of nearly 1 million births across Africa with changes in local $PM_{2.5}$ levels driven by remote dust activation events in the Bodélé Depression in Chad. We use satellite-based observations of dust emission over the Bodélé as an instrument for local variation in $PM_{2.5}$ exposure, an approach that allows us to isolate the impact of poor air quality from local unobserved co-varying factors that might also affect child health. In particular, our estimation strategy uses the product of time-varying dust concentrations exported from the Bodélé Depression and spatially-varying average shares of local $PM_{2.5}$ from natural sources as an instrument for local-level $PM_{2.5}$ exposure over the 12 months following each observed birth in our data. We study whether instrumented $PM_{2.5}$ exposure explains the likelihood a child survives to her first birthday.

We find that relatively modest changes in dust export have large effects on infant mortality across much of Africa, with our estimates similar to quasi-experimental estimates from wealthier countries on the impact of $PM_{2.5}$ on infant health. We show that dust export over the Sahara is controlled in part by rainfall patterns, and that these patterns could change substantially under future climate change. Using data from 36 global climate models, we calculate that future climate-change driven changes in Saharan rainfall could generate very large impacts on African child health through this dust pathway alone, albeit with substantial uncertainty due to varying climate model estimates of how rainfall patterns might change. Finally, back-of-the-envelope calculations suggest that seemingly exotic proposals to use solar power to pump groundwater, dampen soil, and reduce dust ex-

port in these remote locations – and thus reduce PM exposure in downwind population centers – could be cost competitive with leading health interventions designed to reduce infant mortality.

Our work is related to two recent papers. Adhvaryu et al. (2019) use re-analysis data on local dust-based $\text{PM}_{2.5}$ exposure to measure the impact of in-utero dust exposure on child health in West Africa. Foreman (2018) uses a separate re-analysis dataset on aerosol optical depth (AOD) to also study impacts of dust exposure on child health in West Africa, and uses AOD over the Bodélé as an instrument for local AOD to isolate the impact of dust from other determinants of health. Our work is similar to Foreman (2018) in that we use dust emission in the Bodélé as an exogenous source of variation in local pollution exposure, an idea long suggested by both scientific and public health articles noting the importance of Saharan dust for a range of downwind outcomes (Taylor, 2002; Koren et al., 2006; Washington et al., 2009; Adhvaryu et al., 2016). We differ from both papers in using recently-developed satellite-based observations of $\text{PM}_{2.5}$, rather than re-analysis data, to study exposures and health impacts. Re-analysis data are output from models that rely heavily on emissions inventories to generate exposure estimates, and such inventories are associated with large uncertainties in developing countries. Recent satellite-derived datasets (Van Donkelaar et al., 2014), on the other hand, allow direct measurement of $\text{PM}_{2.5}$ exposures that do not depend on emissions inventories, and offer $>100\times$ higher spatial resolution than these re-analysis datasets. Finally, we directly study climatic processes that control dust export from the Sahara, allowing us to assess how future climate change might shape child health through this specific channel, and we quantify how direct intervention designed to limit dust export could benefit child health across the region.

2 Data and approach

The Bodélé Depression is the single largest source of dust emissions in the world, including substantial intercontinental transmission (Koren et al., 2006), and is thousands of kilometers away from most of the observed births in our study (Fig. 1). Dust propagation from the Bodélé is associated with variation in climactic and circulation conditions in the Sahara (Washington et al., 2009; Wagner et al., 2016; Voss and Evan, forthcoming), including variation in local rainfall and in strong winds created from funneling between nearby mountain ranges. Activation events propagate dust from the Bodélé across much of West Africa over a matter of days (Fig. 2a- Fig. S1) (Washington and Todd, 2005). Remotely-

sensed dust atmospheric optical depth (DAOD) data (Voss and Evan, forthcoming) indicate that these activation events are a substantial driver of local DAOD levels across much of Africa, with dissipating intensity over space and time (Fig. 2b).

While dust is transported on a daily time scale, satellite retrievals of aerosols, such as dust, are not available under cloud cover and are available within a limited horizontal distance from the satellite track. Gaps in coverage can be partially mitigated by aggregating over longer time spans. Therefore, while we utilize daily data derived from satellite retrievals to examine dust transport patterns, we utilize annual data derived from similar satellite sources in our mortality analysis (Van Donkelaar et al., 2014). Dust transport is seasonal with high levels of emissions occurring during the Harmattan season (\sim November-March). However, while dust export from the Bodélé Depression is highest during the Harmattan season, the Bodélé is one of the only major African dust sources that is active year-round and emissions from it contribute greatly to annual average $\text{PM}_{2.5}$ concentrations across much of Africa, particularly in the north and west where natural sources account for more than 75% of total $\text{PM}_{2.5}$ (Fig. 3a). Annual variation in dust concentrations and rainfall over the Bodélé correlate strongly with temporal variation in the $\text{PM}_{2.5}$ concentrations across our African study locations (Fig. 3b-c). While exported dust is one of the many factors affecting $\text{PM}_{2.5}$ concentrations in our sample, the motivation for using natural $\text{PM}_{2.5}$ from the Bodélé as a source of variation is that emissions from distant natural sources are less likely to be correlated with confounding local economic activity than emissions from local anthropogenic (or natural) sources. Indeed, while we find a positive time-series correlation in our study locations between estimates of $\text{PM}_{2.5}$ concentrations from anthropogenic sources and night-time lights (a proxy for local economic activity), the same correlation does not exist between nightlights and $\text{PM}_{2.5}$ concentrations derived from remote natural sources (Fig. 3d).

To isolate the causal effect of $\text{PM}_{2.5}$ exposure on infant health, we use an instrumental variables strategy, where plausibly exogenous variation in local $\text{PM}_{2.5}$ driven by distant dust emission is used to identify the impact of $\text{PM}_{2.5}$ exposure on health. Our specific strategy is analogous the common shift-share IV strategy, where spatial variation in average exposure levels is combined with common time series variation in a source of exposure. For the spatial variation, we calculate the share of $\text{PM}_{2.5}$ in each of our African locations that comes from non-anthropogenic sources, averaged over our 2001-2015 study period. The mean share of $\text{PM}_{2.5}$ from natural sources in our sample is 42% and varies from 0% in parts of East Africa to 97% in parts of West Africa (Fig. 3a). For the time series variation, we utilize either time variation in $\text{PM}_{2.5}$ emission from the Bodélé Depression, or

cumulative seasonal or annual rainfall over the Bodélé (increasing rainfall reduces dust export). The mean 12-month average $PM_{2.5}$ concentration in the Bodélé Depression is $75 \mu g/m^3$ in our sample but varies between $65-80 \mu g/m^3$ (Fig. 1). Mean cumulative rainfall over the same period is 4.3 cm in our sample and varies between 2.1-8.0 cm. Our instrument is then constructed by multiplying the spatially varying shares of $PM_{2.5}$ from natural sources in all locations i with either the time-varying dust concentrations or cumulative rainfall in the Bodélé Depression across time-periods t .

For $PM_{2.5}$ data, our main analysis uses satellite-derived data on $PM_{2.5}$ compiled by the Atmospheric Composition Analysis Group at Dalhousie University. These data incorporate retrievals from MODIS AOD as well as other satellites instruments and models (MISR, SeaWiFS, GEOS-Chem) to estimate annual bias-corrected average surface $PM_{2.5}$ concentrations at $0.01^\circ \times 0.01^\circ$ spatial resolution with global coverage (Van Donkelaar et al., 2016). As 0.01° is roughly 1km at the equator, the spatial resolution of this product vastly exceeds the re-analysis data used in other studies, which typically have pixel sizes 0.5° (or ~ 50 km pixels). We utilize this data product along with an analogous product produced by the same group that sub-divides estimates of $PM_{2.5}$ into natural (i.e., dust and sea salt) and non-natural sources. Non-natural source $PM_{2.5}$ is estimated by applying simulated compositional information following the process described in Van Donkelaar et al 2014 (Van Donkelaar et al., 2014).

High-resolution remote-sensing based gridded rainfall data come from the Climate Hazards Group InfraRed Precipitation with station data (CHIRPS). (Funk et al., 2015) CHIRPS incorporates satellite imagery with ground station observations to estimate monthly rainfall totals globally at $0.05 \times 0.05^\circ$ spatial resolution. We sum monthly rainfall data across the post-birth period described above separately for where observed births occur and in the Bodélé Depression. Mean annual rainfall totals are 123 cm in our study locations and 4 cm in the Bodélé Depression.

We then link dust and PM exposure to infant health outcomes across Africa. Data on infant health 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, whereby 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 (Aliaga and Ren, 2006). We used data from all 65 available geo-referenced surveys that were carried out between 2001 and 2015 to reconstruct a village-level birth history. Our sample covers 30 countries and includes 990,696 individual birth outcomes (Fig. 1). 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 following 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.

In order to assign $PM_{2.5}$ exposure to individual birth outcomes we match $PM_{2.5}$ concentrations to the post-birth period encompassing the first 12 months of life inclusive of birth month. Given that $PM_{2.5}$ exposure data are only available annually, we calculate $PM_{2.5}$ exposure as the weighted averages of the annual data, where the weights represent the share of the year that falls into the time period. For example, a child born in the third month of year t would be assigned a post-birth exposure of $(10/12)[\text{exposure in year } t] + (2/12)[\text{exposure in year } t + 1]$. The mean $PM_{2.5}$ exposure level for the post-birth periods is $25.2 \mu g/m^3$ in our sample.

Our first stage regression linking distant dust export to local PM exposure is then:

$$PM_{ijcmt} = \lambda_1 DV_{it} + \mu X_{jcmt} + \gamma_j + \delta_t + \nu_{cm} + \varepsilon_{jcmt} \quad (1)$$

where PM_{ijcmt} denotes the post-birth $PM_{2.5}$ exposure for individual i in cluster j , country c , born in month and year mt . $DV_{it} = (\phi_j * D_t)$ is our instrument equal to the share of PM from dust in cluster j (ϕ_j) times dust export from the Bodélé in year t (D_t). We also include a vector of individual, household, and village-level controls X , which include: local rainfall and temperature in time period t , nightlights in time period t , a dummy for whether the household uses clean cook fuel, a dummy for whether mother completed primary school, mother’s age and age-squared at time of child’s birth, child sex, child’s birth order, and a dummy for whether the child was a twin. Fixed effects include DHS cluster (γ_j), birth-year (δ_t), and country-month (ν_{cm}). Our instruments are very strong, with first-stage F-statistics >100 (Figure 4a).

We then regress infant mortality on predicted values of local $PM_{2.5}$ exposure from Equation (1):

$$y_{ijcmt} = \beta \widehat{PM}_{ijcmt} + \mu X_{ijcmt} + \gamma_j + \delta_t + \nu_{cm} + \varepsilon_{ijcmt} \quad (2)$$

where y_{ijcmt} is a binary measure of whether individual i survived to his/her first birthday.

Equation (2) includes the same fixed effects and controls in X as Equation (1). The fixed effects isolate variation in $\text{PM}_{2.5}$ exposure from other time-invariant, seasonally-varying, or time-trending factors that could be correlated with mortality. In particular, inclusion of cluster fixed effects accounts for any time invariant unobservables that could be correlated with both $\text{PM}_{2.5}$ exposure and mortality risk at the cluster level (e.g. any location-specific differences in mortality rates due to average incomes or healthcare access), inclusion of country-month fixed effects accounts for any seasonal differences across locations within a country (e.g. if mortality rates are higher in June than January in Nigeria), and time fixed effects account for any abrupt or trending factors common across the sample (e.g. the overall decline in infant mortality over the last 15 years). Instead of cluster fixed effects, we can instead include mother fixed effects in equations 1 and 2, whereby the effect of $\text{PM}_{2.5}$ exposure on mortality is derived by comparing two siblings born at different times to the same mother. This strategy further reduces the possibility of unobserved confounders, at the cost of eliminating much of the variation in exposure across the sample.

For estimates of $\hat{\beta}$ in Equation 2 to represent causal estimates of the impact of $\text{PM}_{2.5}$ on infant health, it must be the case that, conditional on controls, our instrument is uncorrelated with other factors beyond local $\text{PM}_{2.5}$ exposure that also affect infant health. While this restriction is untestable, we directly include a vector of controls μX_{ijcmt} meant to address potential additional sources of confounding. In particular, one concern is that global or regional climate phenomena could affect both dust emission over the Bodélé as well as local meteorological conditions, and that these latter conditions could themselves shape health outcomes (e.g. local drought lowering food availability). To account for this, we control directly for both temperature and rainfall locally, as well as for the global time series of Pacific sea surface temperature anomalies, the main index used to measure ENSO. A remaining failure of the exclusion restriction would have to involve large-scale climate phenomena that are not picked up either by these global indices or by local meteorological conditions, and we are unaware of any such phenomena that fit this description.

Finally, we emphasize that $\hat{\beta}$ estimates the local average treatment effect (Angrist and Imbens, 1995), which in our setting is the average effect of $\text{PM}_{2.5}$ on child health for those individuals for whom changes in dust emission over the Bodélé cause changes in local $\text{PM}_{2.5}$ exposure. Our estimates thus will not represent treatment effect estimates for other types or sources of $\text{PM}_{2.5}$ exposure, nor estimates for regions where variation in $\text{PM}_{2.5}$ is not driven by emission from the Bodélé.

3 Results

Annual variation in dust concentrations in the Bodélé Depression is strongly and positively associated with $PM_{2.5}$ variation in our study locations (Fig. 4). We find analogous relationships when we use time-varying measures of rainfall accumulation over the Bodélé Depression rather than direct measures of dust concentration as our source of temporal variation; higher rainfall is known to reduce dust emission in the region (Brooks and Legrand, 2000). These relationships are robust to the inclusion of rainfall and temperature in the study locations and to the inclusion of global ENSO variation. Results are similar when analysis is restricted to exposure comparisons between siblings within the same household (Fig. S2). Consistent with the short time-scale of dust transmission, we do not find meaningful associations between $PM_{2.5}$ concentrations in our study locations and dust concentrations or rainfall in the Bodélé Depression in the previous year. These robust estimated relationships are consistent with existing understanding of the role the Bodélé Depression plays in propagating dust throughout the region. (Washington and Todd, 2005; Washington et al., 2009; Wagner et al., 2016; Wang et al., 2017).

To test whether this strong relationship between remote dust emission and local $PM_{2.5}$ concentrations is driven by common time trends, common year-specific shocks across dusty locations or by average spatial differences between dusty and less-dusty locations, we conduct a placebo exercise where we randomly reorder either the time series of dust emissions over the Bodélé or the spatial shares of baseline levels of dust exposure and re-estimate the relationship between these placebo instruments and local-level $PM_{2.5}$ concentrations. In all cases, estimates on these placebo samples ($n=1000$ for each type of reshuffling) are close to zero (Fig. 4b), suggesting our estimated relationships between local $PM_{2.5}$ concentrations and remote dust emission are not spurious.[†]

We then use this predicted variation in local $PM_{2.5}$ driven by remote dust emission to estimate the impact of $PM_{2.5}$ on local health outcomes. We estimate that a $10\mu g/m^3$ increase in local ambient $PM_{2.5}$ concentration causes a 22% rise in infant mortality across our full sample (95% CI: 10-35%) and a 17% rise in our West Africa sample (95% CI: 2-33%; Fig.

[†]The reason that we conduct this placebo exercise for the first stage (Equation 1) and not for the full IV estimation (Equations 1 and 2) is that manually adding uncertainty to our instruments creates an artificially weak first stage and inflates the IV estimates. For a single instrument the estimated IV coefficient is equal to the reduced form coefficient (i.e., the coefficient associated with regressing infant mortality on our instrument directly) divided by the first stage coefficient. Randomly reordering the elements of our instrument and causing the first stage coefficient to be close to zero mechanically causes the IV estimates to be large. We therefore compare the magnitude of first stage coefficients, rather than IV coefficients, for the placebo samples.

5). Our results are similar whether we use either dust emission or rainfall over the Bodélé as our instrument, and are again robust to controlling for both local temperature and rainfall variation as well as global ENSO variation (Fig. S3). When we restrict analysis to between-sibling comparisons, estimates remain large and positive but are noisier likely due to the substantially reduced variance in our explanatory variable.

For all specifications, instrumental variables estimates are larger than analogous least squares estimates that simply regress local health outcomes on local $\text{PM}_{2.5}$ exposures, conditioning on the same fixed effects. This downward bias in OLS estimates is consistent with unobserved local economic activity generating pollution that harms health, but also leading to livelihood improvements that benefit health.

We emphasize again that our IV estimates represent the effect of $\text{PM}_{2.5}$ on infant health for those individuals for whom changes in dust emission over the Bodélé cause changes in local $\text{PM}_{2.5}$ exposure. This helps explain why we cannot reject the hypothesis that estimates for West Africa and our full Sub-Saharan sample are the same (Fig. 5); even when we estimate $\hat{\beta}$ on our full sample, results are driven by those children (many of whom are in West Africa) whose $\text{PM}_{2.5}$ exposure is affected by Bodélé dust emissions. We include the full sample results because children unaffected by Bodélé dust serve as additional controls and make our sample comparable to other studies.

Our estimates of the effect of $\text{PM}_{2.5}$ on infant mortality are greater than previously published empirical estimates for Africa (Heft-Neal et al., 2018). There are several potential explanations for this difference. First, if local economic activity worsens pollution but improves health outcomes, estimation strategies that do not explicitly account for this covariation will underestimate the impact of pollution on health. Our instrumental variables strategy overcomes this challenge by directly isolating variation in pollution exposure from variation in local economic activity. Second, classical measurement error in local measures of $\text{PM}_{2.5}$ exposure will bias least squares estimates of the effect of $\text{PM}_{2.5}$ on health toward zero, a problem that is corrected by instrumental variables estimation if the instrument is uncorrelated with the local measurement error (Angrist and Krueger, 2001). In our setting, and particularly for our rainfall instrument, this condition is plausibly met, providing another reason why instrumental variables estimates are larger than least squares estimates. Consistent with these explanations, our instrumental variable estimates for Africa are comparable in magnitude to other quasi-experimental estimates from elsewhere in the world (Fig. 5b), nearly all of which come from developed or middle-income countries where exposures are often more precisely measured. Our results are also corroborated by

other studies that use different measures of particulate exposure to examine the effect of dust exposure on health and economic outcomes in West Africa (Adhvaryu et al., 2019). Taken together, these estimates clearly indicate that air pollution is a critical determining factor for child health around the world, and that improvements in air quality can be expected to cause large improvements in child health.

In our African setting, our results highlight the particular importance of air pollution that is non-anthropogenic in origin, and suggest two broad pathways for reducing the health burden of this exposure. A first approach would be to better understand how local interventions, ranging from personalized protective equipment to early warning systems, might reduce individuals’ dust exposure in down-wind inhabited areas. A second approach would be to better understand the existing and potential future controls on emissions at their Bodélé source. These could include both known climate controls as well as prospective direct efforts to mitigate emissions at source locations.

Regarding the climate channel, our results, and much past work (Brooks and Legrand, 2000; Wagner et al., 2016; Wang et al., 2017), suggest that climate variation is a key control on dust export from the region. Because anthropogenic climate change is projected to have large potential influence on these patterns (Biasutti, 2013; Skinner and Diffenbaugh, 2014), altering the trajectory of future climate change could have large effects on child health across the continent through its impact on the emission and transport of Saharan dust. To our knowledge, the potential importance of this channel has not been recognized or explored.

To illustrate this channel’s potential importance, we calculate projected changes in rainfall under climate change over the Bodélé by mid-century (2035-2065) relative to a 1995-2015 baseline, using 36 global climate models from the CMIP5 archive run under the RCP8.5 emissions scenario, a “business as usual” scenario. As has been previously characterized (Washington et al., 2009; Biasutti, 2013; Rodríguez-Fonseca et al., 2015), rainfall projections over the regions are highly uncertain, with the ensemble median projected rainfall change close to zero but with individual models drying or wetting by more than 25% during Harmattan months (Fig 6a). Using our historical data, we then estimate that a 1 mm increase in rainfall in the Bodélé during the Harmattan season reduces $PM_{2.5}$ on average in West Africa by $0.71\mu g/m^3$. This estimate comes from a grid-cell level regression where annual $PM_{2.5}$ in our West Africa study locations is regressed on season rainfall totals in the Bodélé Depression. The regression includes cell fixed effects to control for time-invariant differences over space and country-specific linear and quadratic time trends to control for

country specific trends in either rainfall or $PM_{2.5}$ over time. Cells are weighted by population. To characterize the uncertainty associated with the rainfall- $PM_{2.5}$ relationship, we bootstrap this regression, sampling locations with replacement and re-estimating the coefficient 1,000 times. Combining these bootstraps with rainfall projections from the 36 climate models, we get 36,000 estimated changes in $PM_{2.5}$ shown in Fig. 6b. While the median change in $PM_{2.5}$ is 0, the inner 95% of projected changes range from $12 \mu g/m^3$ reduction to $9 \mu g/m^3$ increase in $PM_{2.5}$. Finally, we use our main model for West Africa (Eq. 2) to estimate the change in infant mortality rate associated with changes in local $PM_{2.5}$ concentrations, again bootstrapping this estimate 1,000 times. Combined with the bootstrap sampling of the 36,000 changes in $PM_{2.5}$ concentrations, we get 36 million estimated changes in infant mortality (Fig. 6c).

We find a large spread in potential changes in infant mortality due to changes in dust exposure (Fig 6b), ranging from a -13% decline in mortality (5th percentile estimate) for models projecting precipitation increase to a 12% increase in mortality by mid-century (95th percentile) for models that dry. To our knowledge, these potential impacts in the left and right tail are substantially larger than any known projected health impact of climate change on the continent. These findings suggest that better understanding the current and evolving future climatic constraints on Saharan dust export remains a key scientific and policy priority.

Because mitigating dust *exposure*, particularly for infants, would be extremely difficult (many houses have open windows or permeable roofs or walls, and infants and young children cannot or will not wear masks), direct efforts to mitigate dust *emission* from Bodélé also warrant future evaluation. In particular, existing estimates of even small amounts of renewable groundwater resources in the region (MacDonald et al., 2012), combined with proven technologies currently deployed in the Sahel that use renewable power to bring this water to the surface (Burney et al., 2010), suggest that local groundwater could potentially be used to dampen the surface and mitigate lofting of dust aerosols from Bodélé, reducing downwind exposures and health burdens.

To explore this possibility, we calculate the costs and benefits of deploying solar-powered central pivot irrigation in the Bodélé, using our estimates how a mm of additional water reduces downwind mortality, combined with conservative estimates of recharge and flow rates, equipment and operational costs, and very conservative assumptions of how averted infant deaths are converted to disability-adjusted life years (DALYs; see Methods). We calculate a cost per DALY averted of $\sim \$60$, which even if low by a factor of 5, is well be-

low cost thresholds designating “priority” interventions in low-income countries. For instance, the World Health Organization recommends that any intervention costing less than per capita GDP is “very cost effective” (their highest designation) (World Health Organization, 2001), and average per capita incomes in West Africa range from ~\$400 in Niger to ~\$2000 in Ghana and Nigeria. Pumping groundwater to reduce dust export also appears cost competitive with many leading health interventions currently in use, including a range of vaccines, water and sanitation interventions, and behavioral interventions (Jamison et al., 2017).

While this back-of-the-envelope calculation clearly ignores other plausibly important constraints to project deployment, such as logistical and political impediments to operating in remote and insecure environments in the Sahara, our intent is to suggest that such projects, if feasible, could yield enormous benefits at a modest cost. Better understanding the technical, economic, and political constraints – as well as potential climate (Li et al., 2018) and biogeochemical (Bristow, Hudson-Edwards, and Chappell, 2010) impacts – of the deployment of such a system is a promising avenue for future work.

References

- Adhvaryu, Achyuta, Prashant Bharadwaj, James Fenske, Anant Nyshadham, Richard Stanley et al. 2016. “Dust and death: evidence from the West African Harmattan.” Tech. rep., Centre for the Study of African Economies, University of Oxford.
- . 2019. “Dust and Death: Evidence from the West African Harmattan.” Tech. rep., National Bureau of Economic Research, NBER Working Paper No. 25937.
- Aliaga, Alfredo and Ruilin Ren. 2006. “Optimal sample sizes for two-stage cluster sampling in Demographic and Health Surveys.” .
- Angrist, Joshua and Guido Imbens. 1995. “Identification and estimation of local average treatment effects.”
- Angrist, Joshua D and Alan B Krueger. 2001. “Instrumental variables and the search for identification: From supply and demand to natural experiments.” *Journal of Economic perspectives* 15 (4):69–85.
- Arceo, Eva, Rema Hanna, and Paulina Oliva. 2016. “Does the effect of pollution on in-

- fant mortality differ between developing and developed countries? Evidence from Mexico City.” *The Economic Journal* 126 (591):257–280.
- Biasutti, Michela. 2013. “Forced Sahel rainfall trends in the CMIP5 archive.” *Journal of Geophysical Research: Atmospheres* 118 (4):1613–1623.
- Bristow, Charlie S, Karen A Hudson-Edwards, and Adrian Chappell. 2010. “Fertilizing the Amazon and equatorial Atlantic with West African dust.” *Geophysical Research Letters* 37 (14).
- Brooks, Nick and Michel Legrand. 2000. “Dust variability over northern Africa and rainfall in the Sahel.” In *Linking climate change to land surface change*. Springer, 1–25.
- Burnett, Richard, Hong Chen, Mieczysław Szyszkowicz, Neal Fann, Bryan Hubbell, C Arden Pope, Joshua S Apte, Michael Brauer, Aaron Cohen, Scott Weichenthal et al. 2018. “Global estimates of mortality associated with long-term exposure to outdoor fine particulate matter.” *Proceedings of the National Academy of Sciences* :201803222.
- Burney, Jennifer, Lennart Woltering, Marshall Burke, Rosamond Naylor, and Dov Pasternak. 2010. “Solar-powered drip irrigation enhances food security in the Sudano–Sahel.” *Proceedings of the National Academy of Sciences* 107 (5):1848–1853.
- Cesur, Resul, Erdal Tekin, and Aydogan Ulker. 2017. “Air pollution and infant mortality: evidence from the expansion of natural gas infrastructure.” *The Economic Journal* 127 (600):330–362.
- Chay, Kenneth Y and Michael Greenstone. 2003a. “Air quality, infant mortality, and the Clean Air Act of 1970.” Tech. rep., National Bureau of Economic Research.
- . 2003b. “The impact of air pollution on infant mortality: evidence from geographic variation in pollution shocks induced by a recession.” *The quarterly journal of economics* 118 (3):1121–1167.
- Foreman, Tim. 2018. “The effect of dust storms on child mortality in West Africa.” *Columbia SIPA working paper* CDEP-CGEG WP No. 47.
- Funk, Chris, Pete Peterson, Martin Landsfeld, Diego Pedreros, James Verdin, Shradhanand Shukla, Gregory Husak, James Rowland, Laura Harrison, Andrew Hoell et al. 2015. “The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes.” *Scientific data* 2:150066.

- He, Guojun, Maoyong Fan, and Maigeng Zhou. 2016. “The effect of air pollution on mortality in China: Evidence from the 2008 Beijing Olympic Games.” *Journal of Environmental Economics and Management* 79:18–39.
- Heft-Neal, Sam, Jennifer Burney, Eran Bendavid, and Marshall Burke. 2018. “Robust relationship between air quality and infant mortality in Africa.” *Nature* 559 (7713):254.
- Jamison, Dean T, Hellen Gelband, Susan Horton, Prabhat Jha, Ramanan Laxminarayan, Charles N Mock, and Rachel Nugent. 2017. *Disease Control Priorities, (Volume 9): Improving Health and Reducing Poverty*. The World Bank.
- Knittel, Christopher R, Douglas L Miller, and Nicholas J Sanders. 2016. “Caution, drivers! Children present: Traffic, pollution, and infant health.” *Review of Economics and Statistics* 98 (2):350–366.
- Koren, Ilan, Yoram J Kaufman, Richard Washington, Martin C Todd, Yinon Rudich, J Vanderlei Martins, and Daniel Rosenfeld. 2006. “The Bodélé depression: a single spot in the Sahara that provides most of the mineral dust to the Amazon forest.” *Environmental Research Letters* 1 (1):014005.
- Lelieveld, Jos, Andy Haines, and Andrea Pozzer. 2018. “Age-dependent health risk from ambient air pollution: a modelling and data analysis of childhood mortality in middle-income and low-income countries.” *The Lancet Planetary health* 2 (7):e292–e300.
- Li, Yan, Eugenia Kalnay, Safa Motesharrei, Jorge Rivas, Fred Kucharski, Daniel Kirk-Davidoff, Eviatar Bach, and Ning Zeng. 2018. “Climate model shows large-scale wind and solar farms in the Sahara increase rain and vegetation.” *Science* 361 (6406):1019–1022.
- MacDonald, Alan M, Helen C Bonsor, B É Ó Dochartaigh, and Richard G Taylor. 2012. “Quantitative maps of groundwater resources in Africa.” *Environmental Research Letters* 7 (2):024009.
- Miller, Grant and B Piedad Urdinola. 2010. “Cyclical mortality, and the value of time: The case of coffee price fluctuations and child survival in Colombia.” *Journal of Political Economy* 118 (1):113–155.
- Rodríguez-Fonseca, Belen, Elsa Mohino, Carlos R Mechoso, Cyril Caminade, Michela Biasutti, Marco Gaetani, Javier Garcia-Serrano, Edward K Vizy, Kerry Cook, Yongkang Xue et al. 2015. “Variability and predictability of West African droughts: A review on the role of sea surface temperature anomalies.” *Journal of Climate* 28 (10):4034–4060.

- Shindell, Drew, Greg Faluvegi, Karl Seltzer, and Cary Shindell. 2018. “Quantified, localized health benefits of accelerated carbon dioxide emissions reductions.” *Nature climate change* 8 (4):291.
- Skinner, Christopher Bryan and Noah S Diffenbaugh. 2014. “Projected changes in African easterly wave intensity and track in response to greenhouse forcing.” *Proceedings of the National Academy of Sciences* :201319597.
- Taylor, David A. 2002. “Dust in the wind.” *Environmental health perspectives* 110 (2):A80–A87.
- Van Donkelaar, Aaron, Randall V Martin, Michael Brauer, and Brian L Boys. 2014. “Use of satellite observations for long-term exposure assessment of global concentrations of fine particulate matter.” *Environmental health perspectives* 123 (2):135–143.
- Van Donkelaar, Aaron, Randall V Martin, Michael Brauer, N Christina Hsu, Ralph A Kahn, Robert C Levy, Alexei Lyapustin, Andrew M Sayer, and David M Winker. 2016. “Global estimates of fine particulate matter using a combined geophysical-statistical method with information from satellites, models, and monitors.” *Environmental science & technology* 50 (7):3762–3772.
- Voss, Kara K and Amato T Evan. forthcoming. “A new satellite-based global climatology of dust aerosol optical depth.” *Journal of Applied Meteorology and Climatology* .
- Wagner, Robert, Kerstin Schepanski, Bernd Heinold, and Ina Tegen. 2016. “Interannual variability in the Saharan dust source activation-Toward understanding the differences between 2007 and 2008.” *Journal of Geophysical Research: Atmospheres* 121 (9):4538–4562.
- Wang, Weijie, Amato T Evan, Christophe Lavaysse, and Cyrille Flamant. 2017. “The role the Saharan Heat Low plays in dust emission and transport during summertime in North Africa.” *Aeolian Research* 28:1–12.
- Washington, Richard, Christel Bouet, Guy Cautenet, Elisabeth Mackenzie, Ian Ashpole, Sebastian Engelstaedter, Gil Lizcano, Gideon M Henderson, Kerstin Schepanski, and Ina Tegen. 2009. “Dust as a tipping element: the Bodélé Depression, Chad.” *Proceedings of the National Academy of Sciences* 106 (49):20564–20571.
- Washington, Richard and Martin C Todd. 2005. “Atmospheric controls on mineral dust emission from the Bodélé Depression, Chad: The role of the low level jet.” *Geophysical Research Letters* 32 (17).

World Health Organization. 2001. *Macroeconomics and health: investing in health for economic development*. WHO.




Figure 1: **Location of observed births and their distance to the Bodélé Depression.** (a) Each point represents one of the 28,461 DHS clusters included in our sample. The number of observed births in a single cluster ranges from 1 to 210. Orange points indicate the West Africa sample. The full sample includes all points. The Bodélé Depression is outlined in black. (b) Observed births are 500-5000km away from the Bodélé Depression.




Figure 2: Dust is rapidly transported from the Bodélé Depression across West Africa and beyond. (a) Daily dust propagation, measured as Dust Aerosol Optical Depth (DAOD), over eight days during a large dust activation event illustrates the magnitude of dust concentrations as well as the time scale and spatial extent of transmission. Dust AOD ranges from 0 to greater than 3.0 in this example. For comparison, AOD values during the recent 2018 northern California wildfires were <1.0 . (b) Detrended time series correlation of DAOD in the Bodélé Depression and DAOD at selected population centers dissipate over space and time. Map shows the location of selected population centers across West Africa as well as Nairobi, Kenya. Right panel shows the time-series correlation of detrended DAOD values between the mapped locations and the Bodélé Depression. In general, as distance to the Bodélé increases, peak correlations in the detrended DAOD time series occur later and are lower in magnitude. Variation in DAOD in Nairobi, a populated location distant and not downwind from the Bodélé, is uncorrelated with dust from the Bodélé.

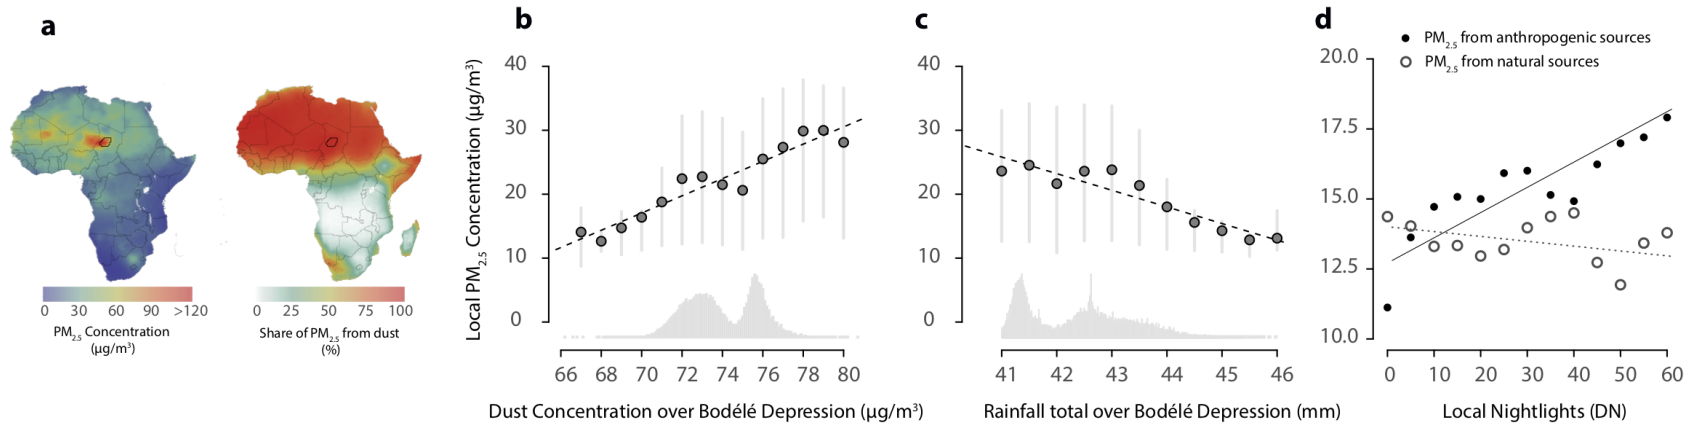


Figure 3: Local air pollution in Africa is driven by both local economic activity and remote natural sources. (a) Average annual $PM_{2.5}$ concentration 2001-2015 (Van Donkelaar et al., 2016), with the Bodélé Depression outlined in black. (b) The share of $PM_{2.5}$ estimated to come from natural sources (Van Donkelaar et al., 2016). (c) Higher dust concentrations over the Bodélé are associated with higher local $PM_{2.5}$ concentrations in our study locations. Annual observations of dust concentrations over the Bodélé were divided into equal sized bins and the average annual $PM_{2.5}$ concentrations (points) and interquartile ranges (gray lines) over all survey locations were calculated within each of these bins. Years with more dust in the Bodélé Depression are associated with higher average $PM_{2.5}$ concentrations across Africa. (d) Analogous to (c) using rainfall totals over the Bodélé in place of dust concentrations, showing that higher Bodélé rainfall totals are associated with lower local $PM_{2.5}$ concentrations in our study locations. (e) Binned scatterplot of nightlight intensity and $PM_{2.5}$ concentrations by source in our study locations. Nightlight intensity was assigned to 1-unit bins and average $PM_{2.5}$ concentrations were calculated within each bin separately for anthropogenic and naturally sourced (i.e., dust and sea salt) $PM_{2.5}$. Higher nightlight intensity is associated with higher local concentrations of $PM_{2.5}$ from anthropogenic sources (solid line), but not with higher $PM_{2.5}$ from natural sources (dotted line).

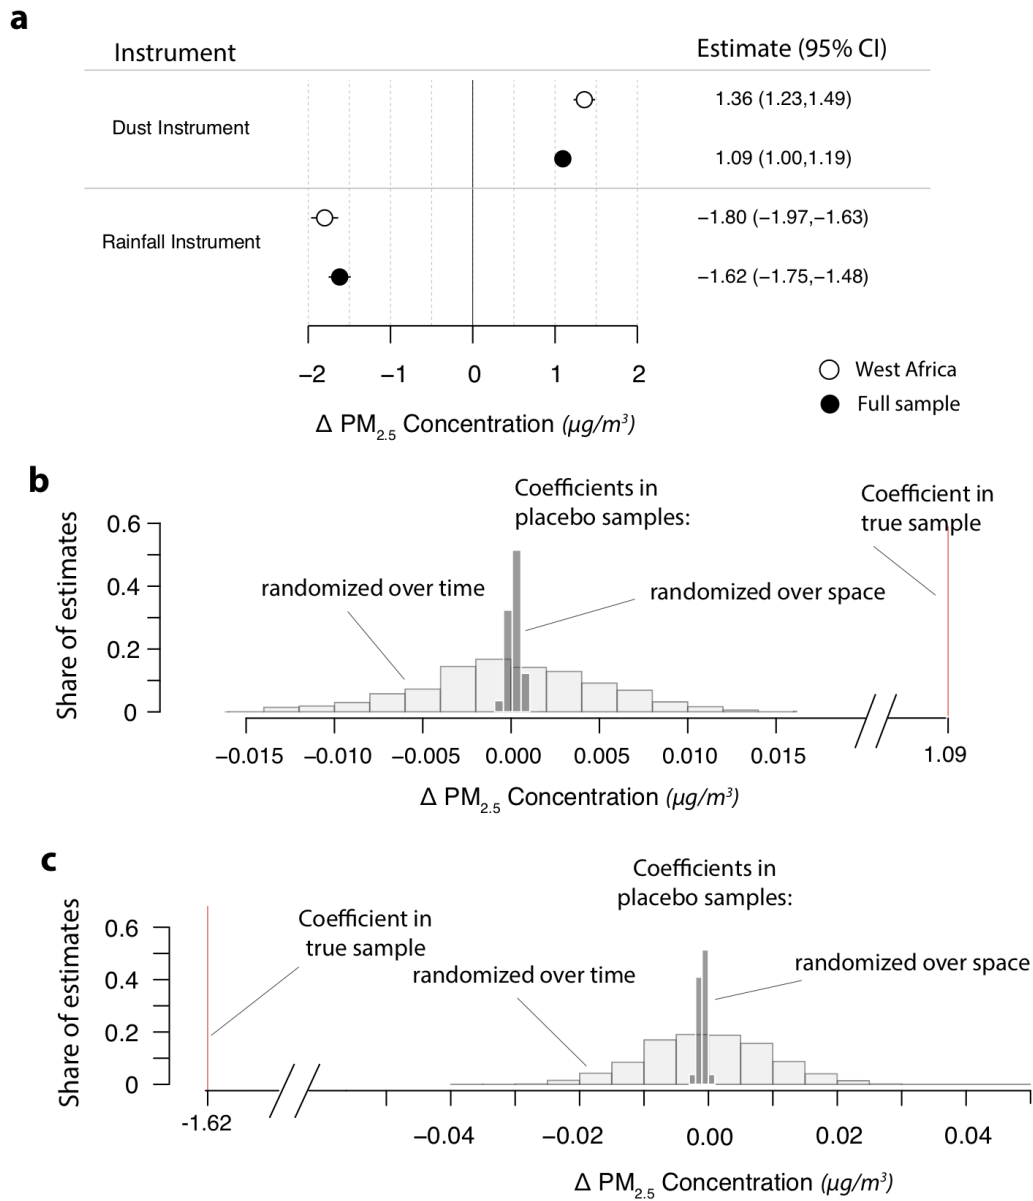


Figure 4: Variation in dust and rainfall over the Bodélé Depression are strong predictors of variation in $PM_{2.5}$ concentrations elsewhere in Africa. (a) First stage estimates corresponding to Equation (1) show the effect of dust and rainfall instruments on $PM_{2.5}$ concentrations across our study locations. Effects are larger in West Africa, closer to and downwind from the Bodélé Depression. (b) In a placebo test, we find no effect of Bodélé dust on local PM when randomly re-order the time series of Bodélé dust (light grey histogram, 1000 resamples), or when we randomly re-sort shares of $PM_{2.5}$ concentrations across space (dark grey histogram, 1000 resamples); estimates for both placebo experiments are centered around zero and several orders of magnitude smaller than the effect estimated with the true ordering. (c) Analogous to (b) but with our rainfall instrument.

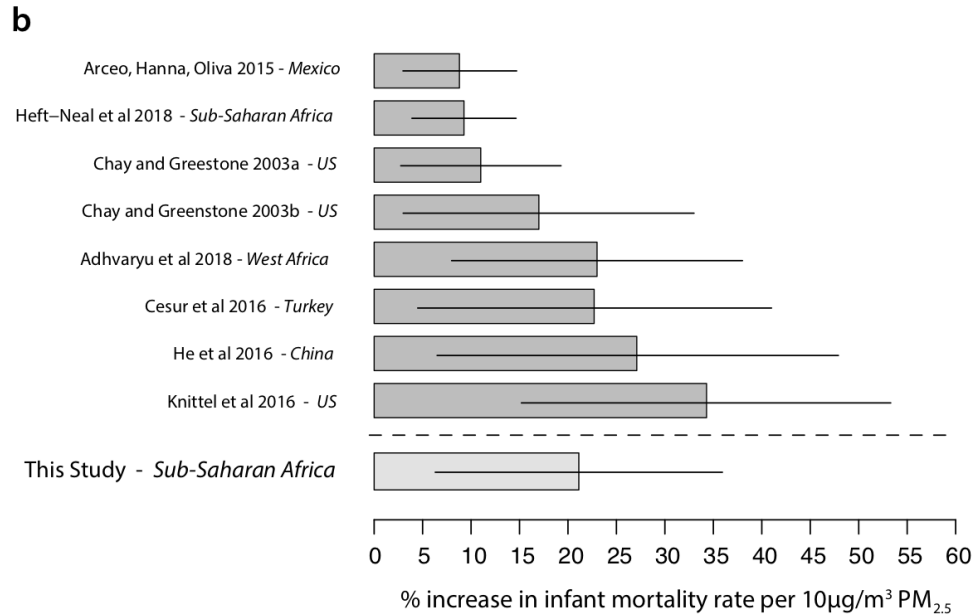



Figure 5: Instrumental variable (IV) estimates suggest large impacts of $\text{PM}_{2.5}$ on infant mortality, with effect sizes similar to quasi-experimental studies from higher income countries. **a** IV estimates indicate a $10\mu\text{g}/\text{m}^3$ increase in $\text{PM}_{2.5}$ concentration causes infant mortality to rise by 15% in West Africa and 22% across the full Sub-Saharan Africa sample; IV estimates are substantially larger than OLS estimates. Effects are similar when using either dust or rainfall over the Bodélé Depression as instruments for local $\text{PM}_{2.5}$. Both IV and OLS estimates include DHS cluster, birth-year, and country-month fixed effects as well as household, mother, and child controls (see Methods). **b** IV estimates are larger than previous all-Africa estimates, and similar to causal estimates in other parts of the world with higher incomes.




Figure 6: **Disagreement among climate models projections of future rainfall changes in the Sahara generates a large spread in projected infant mortality changes in Africa** **a** Projected changes in rainfall levels (2045-2055 relative to 2000-2015) over the Bodélé Depression during the Harmattan season (Nov-Apr) across 36 climate models. Baseline average rainfall during these months is 40mm. **b** Projected changes in PM_{2.5} across West Africa in response to rainfall changes shown in (a), incorporating uncertainty in both the distribution of rainfall projections and in the statistical relationship between rainfall in the Bodélé and local PM_{2.5} across West Africa (See Methods). **c** Changes in infant mortality resulting from the changes in PM_{2.5} modeled in (b). Distribution of outcomes represents uncertainty in PM_{2.5} changes as well as uncertainty in the relationship between PM_{2.5} and infant mortality.

Methods appendix and figures

Interpreting the first stage For several specifications using the dust instrument, the magnitude of the estimated first stage coefficients is greater than 1. At first glance this would seem to imply that a 1-unit increase in dust concentration in the Bodélé causes a more than 1-unit increase in $PM_{2.5}$ concentrations across Africa. However, a coefficient greater than 1 only implies that the rate of dust transmission from the Bodélé to our study locations is greater than the share of $PM_{2.5}$ from dust in those locations. A simplified version of our first stage is $PM_{it} = \gamma x_{it} + \varepsilon_{it}$, where $x_{it} = \delta_t \cdot \theta_i$ is our instrument, δ_t is dust in the Bodélé, and θ_i is share of pm from dust. We are interested in $\frac{dPM_{it}}{d\delta_t} = \frac{dPM_{it}}{dx} \cdot \frac{dx}{d\delta_t}$, or the change in total $PM_{2.5}$ in our study locations per 1 unit increase in dust in the Bodélé Depression.

Let $\frac{dPM_{it}}{d\delta_t} = \beta$ be the transmission rate of PM from the Bodélé to local total $PM_{2.5}$. We know $0 \leq \beta \leq 1$. $\frac{dPM_{it}}{dx} = \gamma$ is the coefficient estimated in our first stage equation above. Our share of PM from dust is $\frac{dx}{d\delta_t} = \theta$, with $0 \leq \theta \leq 1$. Thus: $\beta = \gamma \cdot \theta$ and $\gamma = \frac{\beta}{\theta}$. Our estimates for θ for all of Africa, for West Africa, and for Africa outside of the west are 0.44, 0.74, and 0.17, respectively. Therefore in order to estimate first stage coefficients (γ) greater than 1, the increase in total $PM_{2.5}$ in our study locations in response to a $1 \mu g/m^3$ increase in dust in the Bodélé would need to be greater than 0.44, 0.74 and 0.17 for each of the three groups. In fact, we estimate that, for example, a $1 \mu g/m^3$ increase in Bodélé dust results in an $0.86 \mu g/m^3$ increase in average $PM_{2.5}$ across West Africa and therefore our first stage estimate is greater than 1.

A similar logic applies to our first-stage estimates with the rainfall instrument. The estimated magnitude of the rainfall instrument in our first stage ranges from -1.5 to -2, indicating that a 1-unit increase in our rainfall instrument reduces average $PM_{2.5}$ concentrations in our study locations by -1.5 to -2 $\mu g/m^3$. The same dust shares that are used to construct the dust instruments are used in the rainfall instrument. Therefore, the magnitude of the coefficient on the rainfall instrument in the first stage will be greater than 1 if an additional millimeter of seasonal rainfall over the Bodélé reduces average $PM_{2.5}$ concentrations across our study location in West Africa by more than $0.75 \mu g/m^3$. We estimate that 1 millimeter of additional rainfall in the Bodélé reduces $PM_{2.5}$ in our study locations by an average of $1.2 \mu g/m^3$ and therefore our estimate of the coefficient on the rainfall instrument ranges from -1.5 to -2.

Calculation on mitigation of dust emission Damages from aerosol particulate matter can be mitigated by addressing either exposure or emission. Historically, dust emissions have been considered a natural phenomenon and (limited) efforts have focused on avoidance behaviors and, in farmland, surface roughening and vegetation planting as emission controls. Avoidance behavior would be extremely difficult for infants, especially in the West African context, and agricultural mitigation techniques are not relevant in the non-agricultural Bodélé depression. As such, additional, more exotic-seeming mitigation measures may in fact be appropriate.

One potential mitigation option would be to use limited renewable groundwater resources in and near the Bodélé depression to dampen the surface of the emission region during the Harmattan season. Assuming that the active dust emission region is around 10,000 km² (Washington et al., 2009), and that sustainable recharge is of the order a few mm at average depth of 175 m, with maximum flow rates of 12 L/s (MacDonald et al., 2012), we estimate that a distributed grid of borewells, pumps, and distribution systems, powered by photovoltaic arrays, could substantially and sustainably reduce dust-related infant mortality across West Africa. At the average depth-to-groundwater, PV-powered submersible pumps could provide approximately 8 m³ per day, or enough, over the course of the Harmattan season, 2.2 mm of water over a standard quarter section (800m x 800m, or 160 acres). We calculate, using above estimates, that this would avert 37,000 infant deaths in West Africa annually. Although more than 15,000 such units would be required to cover the Bodélé, off-the-shelf technologies could be installed for ~ \$3000, with lifetimes of >20 years. The most effective and appropriate distribution systems would need to be determined by testing, but as an estimate, distribution costs of \$25,000 per unit (again over 20 years) leads to an annualized cost of ~ \$22M. This translates to ~ \$600 per avoided death in West Africa.

To compare this estimate to existing cost-effectiveness estimates of health interventions, we very conservatively assume that infants whose death before reaching age 1 year has been averted live, on average, 10 healthy life years (and therefore the averted death also averts 10 disability-adjusted life years, or DALYs). This yields a conservative cost-effectiveness estimate of ~ \$60 per DALY averted. This would make such efforts to reduce dust emissions high-value interventions by any benchmark for low-income countries, and would be cost-competitive with many leading health interventions (Jamison et al., 2017). While this back-of-the-envelope calculation clearly ignores other potential constraints to project deployment, such as logistical and political impediments to operating in remote and insecure environments in the Sahara, our intent is to suggest that such projects, if feasible, would

yield enormous benefits at a modest cost.




Figure S1: **Dust from the Bodélé is exported to remote regions within a daily time-scale.** The left columns shows daily DAOD and the right column shows the raw MODIS image for the corresponding date. Daily satellite based measures of dust concentrations illustrate that dust transmission occurs over the period of a few days. However, spatial coverage for daily data is limited by missing observations due to both satellite paths and cloud cover.




Figure S2: Variation in dust concentration and rainfall over the Bodélé Depression are strong predictors of variation in $PM_{2.5}$ concentrations in other parts of Africa. Dust in the Bodélé Depression is positively associated and rainfall negatively associated with $PM_{2.5}$ concentrations in both our West African sample (white dots) and full African sample (black dots; see Fig. S1 for study locations). Contemporaneous effects are large and statistically significant (F-stats of 17.17-366) while lagged effects are not, and results are consistent across a large range of fixed effects and controls. All estimates include the fixed effects and controls listed in the first and second columns.




Figure S3: Instrumental variables estimates of the effect of $PM_{2.5}$ on infant mortality. A $10\mu g/m^3$ increase in $PM_{2.5}$ concentration causes infant mortality to rise by 4-19% in West Africa and 8-24% across the full African sample. Effects are similar when using either dust or rainfall over the Bodélé Depression to instrument for local $PM_{2.5}$. IV estimates are larger than OLS estimates, although IV estimates that use only within-mother variation in child outcomes are imprecisely estimated. All estimates include the fixed effects and controls listed in the first and second columns.