

The Effects of Dust Storms on Economic Development

Timothy Foreman*

June 19, 2020

[\(Link to most recent version\)](#)

Abstract

Dust storms are a common occurrence for populations residing in semi-arid environments and can result in a variety of immediate and long-term impacts. While previous literature documents many of these short-term effects, such as increases in various respiratory issues (e.g. asthma attacks, suffocation) and increases in traffic accidents (resulting from disrupted transportation networks), this is the first study to use exogenous variation in dust exposure due to long-range transport to study the effects of dust storms on health and economic activity. I instrument local dust values using dust observed over the Bodélé Depression of the Sahara Desert, the largest dust source in the world. Economic growth in West Africa is reduced by 3% per standard deviation increase in dust exposure over 2 years. Agricultural yields decline in the year of impact on average by 2%. The effects found here could be a contributing factor to reduced economic development in West Africa, and indicate that dust storms should be considered an important part of geographic endowments alongside other climate indicators.

*Email: timothy.foreman@columbia.edu; School of International and Public Affairs, Columbia University. I would like to thank Wolfram Schlenker, Rodrigo Soares, Ronald L. Miller, Geoffrey Heal, Solomon Hsiang, Douglas Almond, Jeffrey Shrader, Stephanie Lackner, Carlos Pérez García-Pando, Belinda Archibong, Jack Willis, Kiki Pop-Eleches, and Miguel Urquiola as well as seminar participants at the 9th International Workshop on Sand/Dust Storms and Associated Dustfall, Heartland Environmental and Resource Economics Workshop, and the Eastern Economic Association meeting and at Columbia University.

1 Introduction

Air pollution is known to have an array of damaging effects, from health to economic activity. While most air pollution research focuses on industrial emissions, the impacts of air pollution caused by natural dust storms has received much less attention despite the potential for both mitigation of and adaptation to the effects. This study extends the work of recent studies which have shown a correlation between this natural source of air pollution and human health. Previous studies have not considered that the measures of dust that are used are not only measuring dust in the atmosphere, but other pollutants as well. These other pollutants could be a byproduct of economic activity, and therefore resulting in endogeneity when trying to estimate the effect of this pollutant on economic activity. This paper uses the physical process of long-range dust transport to provide credibly exogenous variation in order to estimate the effects of dust storms on economic activity.

To do this, I model the dust that is transported from large natural sources. I allow for time-varying emissions at the source, as well as temporal and spatial variation in winds that determine where the dust affects populations. I focus particularly on West Africa, which neighbors the world's largest dust source, the Bodélé Depression in the Sahara Desert, which is estimated to account for half of natural dust emissions across the entire world. West Africa is of particular interest, as it is the region with the highest dust concentrations and is home to some of the world's most vulnerable populations, making the impacts of these storms particularly severe. Compared to a statistical approach commonly used in the economics literature (see e.g. Deryugina et al. (2016); Schlenker and Walker (2016) for similar approaches), modeling the dust transported using an advection-diffusion model allows for significantly more precise estimates.

Recently, policymakers have noted that dust storms pose a critical threat to economic and sustainable development, with the UN General Assembly President Miroslav Lajčák declaring “One sand or dust storm can cost hundreds of millions of dollars. The losses are really felt in the agriculture, transportation and infrastructure sectors. These are resources that could have been channeled towards sustainable development...” in July 2018. These storms can be very intense, with some storms causing “serious economic loss and [being]

as hazardous as a disaster caused by an earthquake.” (Goudie and Middleton, 2006) Dust storms commonly disrupt traffic, rail transport, and have caused plane crashes. Their effects are most commonly felt in West Africa, but they also occur in the Middle East, India, China, and Australia. In this paper, I focus on the macroeconomic and health impacts of dust storms, focusing on the mechanisms where data are available and there is prior evidence of effects, particularly in agriculture.

The contribution of this paper is two-fold. The first is introducing a physical model of the atmosphere in order to isolate the exogenous component of air pollution due to physical processes unaffected by human activity. The second is using this model to generate instruments that allow for the identification of the macroeconomic impacts of dust exposure. Dust exposure is found to significantly reduce economic growth over the year of exposure and the following year. I provide evidence that endogeneity is likely to be a problem for estimation, as economic activity is positively correlated with estimates of particulate matter, despite the dominance of dust as a fraction of the particulate matter. I then show that the estimate reverses sign when using the long-range transported dust to instrument for local pollution readings.

Since ground-level observations of dust storms are especially sparse in this region, I use a newly-developed reanalysis product, the Modern-Era Retrospective analysis for Research and Applications, Version 2 (MERRA-2). This dataset provides daily gridded values for Aerosol Optical Depth (AOD), which measures dust and other aerosols in the atmosphere, and also contains other climate variables including temperature, precipitation, and winds. I combine the MERRA-2 data with a range of outcome data, from micro-level survey data to macroeconomic measures. Given the ubiquitous nature of the dust reanalysis, I am able to match the outcome data at its level of observation, whether that is at the country or a village level.

I find that a one standard deviation increase in dust exposure in a country over a year decreases economic growth by 3%. The prior work on dust storm impacts on economic outcomes is quite limited. Birjandi-Feriz and Yousefi (2018) examine the effects of dust storms on manufacturing firms in Iran, finding that an increased day of dust storm activity reduces firm productivity by 0.081%. Zivin and Neidell (2012) find that air pollution can

negatively affect worker productivity, but they do not study dust storms in particular. My results are comparable in size to the effects seen from tropical cyclones, with Hsiang and Jina (2014) finding that a one standard deviation increase in cyclone exposure decreases long-term economic growth over by nearly 5%.

The nature of dust storms could mean that their impacts are systematically different from other major disasters. While they tend to be persistent throughout the dry season in West Africa, large, massively destructive storms are relatively rare. I find that agricultural yields decline by 2% for a one standard deviation increase in dust in the year of exposure. Given that agriculture is the main source of income in the region, it is not surprising that the effects are close to the size of the growth effects. Taken together, the total effects of dust storms could severely curtail economic growth in West Africa.

These results are consistent with recent work studying the impact of the dust bowl in the United States. Hornbeck (2012) finds that places more affected by the dust bowl had lower agricultural production for a long period after it ended. Arthi (2018) finds that children with large in-utero exposure to the dust bowl had higher poverty rates, but that educational attainment increased and that government support offset some of the damage. In the West African context, where governmental institutions are relatively weak, it is difficult for such remedial interventions to take place.

2 Data

2.1 Dust data

There are many ways to measure dust in the atmosphere. Weather stations, satellites, and climate models all provide estimates of various particulates. While weather stations also typically measure visibility and may even report weather codes for “blowing dust,” there are few weather stations across West Africa, and reporting is unreliable. Satellite-based observations or reanalyses, which provide more reliable spatio-temporal coverage, can be highly valuable in this context. Here, I utilize a new reanalysis product, the Modern-Era Retrospective analysis for Research and Applications, Version 2 (MERRA-2), which “is the first

long-term global reanalysis to assimilate space-based observations of aerosols and represent their interactions with other physical processes in the climate system” (C. A. Randles et al., 2016). This product integrates a climate model with aerosol models, but is also constrained by satellite observations. In particular the Aerosol Optical Depth (AOD) is calibrated to observations from various satellite instruments over the period 1980-present. The MERRA-2 data has been extensively verified through comparison to other data sources, and shows very high agreement with various other measurements (Buchard et al., 2016, 2017). It also outputs weather variables including temperature, precipitation, and winds.

While AOD can vary due to multiple different particulates, including some industrial pollutants as well as smoke, over the world’s “dust belt” across the Middle East and northern Africa, it is largely dominated by dust. Even if it measured only dust, however, it is still not the ideal measurement of dust that will affect human societies, as it is affected by total atmospheric column dust, not just dust at or near the surface.

MERRA-2 AOD is provided from 1980-present every day for a $0.5^\circ \times 0.625^\circ$ grid. Annual average optical depth over a location varies from 0.1 in regions of very little aerosol up to 2 in highly dusty regions. The highest values for much of the year occur over and downwind of the Sahara Desert, as seen in figure 8. Large values are also seen over the other large deserts of the world, including the Arabian Desert, the Gobi Desert, much of Australia, and to a lesser extent the Southwestern United States, the Patagonian Desert, and the Kalahari Desert. I focus on West Africa due to the extreme dust concentrations, as well as the potential vulnerability of these regions due to low income levels.

The region of intense average dust values within the Sahara Desert is highlighted in Figure 1, which also illustrates the difference between the wet and dry season in the monsoon climate of West Africa. During the wet season, the winds blow from the south and east, bringing moist air from the equator over the continent, leading to precipitation. During the dry season, the winds blow from the north and east, bringing dry air from over the Sahara into West Africa, and along with it significant quantities of dust. I exploit the dry season relationship between major dust sources in the Sahara Desert and dust over the populated regions of West Africa.

2.2 Outcome data

Data on GDP is available from three major sources: the Penn World Tables, World Bank Development Indicators, and the Maddison Project Database. Figure 9 compares the three GDP datasets and the data on nightlights for twelve countries in West Africa, with GDP reported in constant US dollars. While some countries show general agreement across sources, others show significant discrepancies. It has been documented that economic production is not well measured in this part of the world (Jerven, 2013; Johnson et al., 2013), and that can easily be seen from the data. I show results from the various data sets and show that the main results are consistent across all of them.

I use the nightlight data to supplement the official data on GDP. Data on nightlights comes from Henderson et al. (2012). The nightlight measure is area-weighted average digital number from the Defense Meteorological Satellite Program. This is my preferred measure, given the limitations of the official data in this region over the whole time period. I provide the main results with all of the official GDP datasets in order to ensure that the results are not being driven by the physical blocking of light by atmospheric dust.

I also consider the effects on agricultural yields at the country level. Agricultural data for multiple crops is available at the country-year level from the FAO. In particular, I consider the six most commonly grown crops: cassava, cowpeas, groundnuts, maize, millet, and sorghum. These crops account for a large share of GDP of countries in this region, but yields are significantly lower here compared to the rest of the world. While there are many factors at play to explain the low yields, dust storms are known to have negative yield impacts (Stefanski and Sivakumar, 2009).

To explore further what the economic implications of dust exposure may be at a more micro level, I use data from the Living Standards Measurement Surveys (LSMS). The LSMS data are collected by the World Bank in many developing countries. The primary data I use are the National Survey on Household Living Conditions and Agriculture from Niger, which is a panel survey of households conducted in 2011 and 2014. This data contains detailed information about household members and their activities, including agricultural production, time spent in household enterprises, and some health outcomes. There are more than 2,000

households that are interviewed in both rounds, representing nearly 23,000 individuals.

Survey data outcomes are geographically specified at a “cluster,” corresponding to an area or population the size of a small town. For clusters, a latitude-longitude point is provided, which is jittered by less than 1km or 5km for the vast majority of the data. For the survey data, the given point is matched to the MERRA-2 grid cell that it is a part of in order to be matched with dust exposure data.

3 Empirical Strategy

3.1 Physical Model

A central contribution of this paper is its use of physical modeling to disentangle the endogenous and exogenous components of air pollution. While it makes use of a simple physical model, to my knowledge, it has not been applied in the way it is applied here. To this end, I lay out the basics of the model here.

The purpose of the model is to determine how much dust is in a particular place due to natural emission that has been blown from source regions in the Sahara. This dust is exogenous to human activity (apart from any slow-moving influence that climate change may have), as it is coming from unpopulated regions and blown by natural wind patterns. The physical model takes into account both the emission level at the source as well as the wind strength and direction from source to where it affects human activity.

The wind-driven model involves using the dust concentrations over source regions and solving an advection-diffusion equation to estimate dust transport. To simplify this procedure, I assume that there is no vertical mixing or deposition, so that dust is transported only horizontally, and that dust which reaches ground-level over populated areas comes from ground-level over the source region. With this simplification, the model solves:

$$\frac{\partial f(\mathbf{x}, t)}{\partial t} + U(\mathbf{x}, t) \frac{\partial f(\mathbf{x}, t)}{\partial \mathbf{x}} = D \frac{\partial^2 f(\mathbf{x}, t)}{\partial \mathbf{x}^2} \quad (1)$$

Here, $f(\mathbf{x}, t)$ denotes the dust distribution, \mathbf{x} is the position vector, $U(\mathbf{x}, t)$ is the wind field, and D is a constant that determines the rate of diffusion. To solve this equation

numerically, the region of interest is broken into a grid the size of the MERRA data. The equation is then solved using finite differences. $U(\mathbf{x}, t)$ is determined by the MERRA data at the hourly level. The model is run in half-hour time steps to ensure stability. The average over the wet and dry seasons is shown in Figure 11. The model successfully reproduces the observed seasonal pattern in this region.

The source regions are defined based on a map of sources in Prospero et al. (2002). The sources in West Africa are shown in Figure 10. This incorporates the Bodélé Depression as well as a region near the border of Mali and Algeria and small parts of the desert in eastern Libya. These regions have close to no populations living in or near them, so dust emission from these areas are reasonably expected to be unaffected by human activity.

3.2 Econometric Model Setup

The first outcome I consider is country-level economic growth. To estimate the impacts of dust storms on growth, I exploit random year-to-year variation in dust exposure in a fixed-effects regression. I largely follow the approach that Hsiang and Jina (2014) take to look at the effects of hurricanes, with country-level regressions of the following form:

$$g_{cy} = \sum_{j=0}^L [\beta_j D_{cy-j} + \rho_j C_{cy-j}] + \alpha_{0c} + \alpha_{1c}y + \alpha_{2c}y^2 + \epsilon_{cy} \quad (2)$$

Here, D_{cy-j} is a measure of dust exposure. Since the dust exposure measure is likely to be endogenous to economic growth, I will instrument this variable using long-range transported dust. I control for climate C_{cy-j} , which includes temperature and precipitation over the year. Trends in growth are controlled for using a quadratic in time, as Burke et al. (2015) use in studying the effects of temperature on growth. The coefficient β_0 captures the level effect of dust exposure on economic production, while

$$\Omega_L = \sum_{j=0}^L \beta_j$$

captures the effect on GDP over L years. The ϵ_{cy} are allowed to be arbitrarily correlated across 5 years for the same country and across countries with centroids that are within

1000km of each other for the same year.

Since industrial pollutants and smoke from biomass burning, as well as locally-emitted dust, could also influence AOD and will be correlated with GDP, it is important to address the potential for endogeneity. Fortunately, given the large variation in natural year-to-year dust emission in unpopulated source regions, it is possible to estimate dust transported from large natural sources and use this dust as an instrument for local dust. To instrument the dust over populated areas, I use dust observed over the Bodélé Depression in the preceding days. The Bodélé Depression is the largest dust source in the world, accounting for nearly half of the world’s natural dust emissions. In a first specification, I run the model as an instrumental variable model, where the instruments are generated by a statistical model that links dust in the source to dust over populated regions. In a second specification, I use the output of a transport model using observed dust over the source regions as the boundary condition, then using the modeled dust as an instrument. The identifying assumption in both cases is that dust over the Bodélé depression is only correlated with country-level GDP through its effects on the dust concentration over the country. It is also necessary that the dust over the Bodélé depression affects dust over the West African countries of study.

The need to instrument local AOD presents a variety of empirical challenges. Dust transport is a process taking place over days to weeks, whereas GDP is typically reported only yearly, and yields are only realized once or twice per year. In theory, running the first stage at a disaggregate level then aggregating to run the instrumental variable regression is a case of generated instruments, and the asymptotics will be the same as the IV regression at the aggregate level. However, when there is a small number of aggregated observations, uncertainty in the first stage specification, or aggregation cannot be done arbitrarily, there are advantages to running the first stage at the disaggregated level. Simulation results for different approaches are shown in the appendix. The results are shown for a two-step first stage procedure, with the first step run at the pixel-day level as detailed below.

3.3 Model Implementation

In the first specification, where dust is instrumented using a statistical relationship between source location dust and dust over the outcome location, dust in a given MERRA pixel on

a given day is predicted using the following equation:

$$D_{pt} = \sum_{j=0}^{10} \gamma_{p,sj} D_{B,t-j} + \theta_p C_{pt} + \nu_{pt} \quad (3)$$

where D_{pt} is the dust in pixel p on day t and $D_{B,t-j}$ is the mean dust over the Bodélé Depression j days prior. The C_{pt} controls for temperature and precipitation in pixel p on day t . Importantly, $\gamma_{p,sj}$ depends on the pixel and the season, dry or rainy. The result of this regression is shown for a particular location on the coast of Togo in figure 2. As expected, the relationship is much more apparent during the dry season than the wet. A 1 unit increase in dust over the Bodélé Depression leads to a 0.06 unit increase over the Togo coast 3 days later during the dry season. This indicates that about 6% of the dust over a point in the source region is present over the coast of Togo 3 days later. The R-squared of the regression is 0.54, indicating that about half of the dust in this region is attributable to the Bodélé Depression on average. While this is high and means that this transport is important to the dust loading over West Africa, it also means that nearly half of the AOD is from other, potentially human-influenced sources.

The dust values, D_{pt} , and predicted dust values, \hat{D}_{pt} are then aggregated to the country level using population weights and then averaged over the year. The population weights are from the Gridded Population of the World Version 4 for the year 2000. (Center for International Earth Science Information Network - CIESIN - Columbia University, 2016) This process yields predicted values, \hat{D}_{cy} , which is used to instrument for the aggregated dust measure D_{cy} .

The second procedure is to use both the observed dust measurement over the source region and use wind directions to estimate how much reaches populated locations. The main advantage of this approach is that the wind data allows me to exploit variation that is both spatial and temporal. The first approach can be thought of as similar to a shift-share instrument, which exploits temporal variation interacted with a location dummy.

4 Empirical Results

4.1 OLS Results

I start by estimating Equation 2 by OLS. The results for 3 leads and 3 lags are shown in Figure 3. The only significant estimate is the positive coefficient on the 1-year lead. If I were to (mistakenly) interpret this as a causal effect, it would imply that an increase in dust exposure causes last year’s growth to increase. This is obviously the incorrect interpretation. A much more sensible conclusion is that if there is higher economic growth in a particular year, the next year experiences higher air pollution, likely due to the increased emission of industrial pollutants or added dust due to more agricultural cultivation.

4.2 First Stage Country-Level

The first stage regression for both versions of the model takes the following form:

$$D_{cy} = \sum_{k=0}^L [\eta_k \widehat{D}_{cy-k} + \psi C_{cy-k}] + \alpha_{1c}y + \alpha_{2c}y^2 + \alpha_c + \nu_{cy} \quad (4)$$

where D_{cy} is the observed dry-season dust over country c in year y and \widehat{D}_{cy-k} is contemporaneous and lagged values of the aggregated predicted values of dry-season dust from Equation 3. C_{cy-k} contains population-weighted country-level temperature and precipitation. The first stage results for the η_k for both the statistical and wind-driven models are shown in Tables 8 and 9. This is the result of a regression of country-year dust on the aggregated pixel-day level contemporary predicted dust as well as up to 5 lags. The instruments become weak after the inclusion of any lags for the statistical model, and the coefficient on the contemporary value begins to fluctuate. The coefficient on the contemporary value is fairly consistent across the inclusion of varying numbers of lags for the wind driven model. The value of the F-Statistic is decreasing across the inclusion of more lags, but remains above 10 for the inclusion of less than 5 lags.

4.3 Macroeconomic Results

I first estimate equation 2 using the wind-driven model described above. The results are shown in Table 2 for $L = 0$. Column (1) shows the results for growth estimated using a combination of official data and nightlights as described in Henderson et al. (2012) and columns (2)-(4) show results for the three major GDP datasets. The estimates in the year of impact are all negative, and largely statistically significant. The estimates for a 1-year lag of dust exposure are also all negative and highly statistically significant. Summing these effects implies that economic output in the year after a one-standard deviation increase in dust exposure during the dry season is between 2.5-6.5 percent lower than it would be otherwise.

These estimates are large and comparable to the effects of other natural disasters. Hsiang and Jina (2014) find that a 1 standard deviation cyclone decreases economic growth by 3.6% after 20 years. The comparability of these magnitudes indicates that while dust storms do not frequently cause severe damages in single events, they can still have significant impacts over the course of a season of dust exposure. They can have effects of a similar size as other major macroeconomic shocks including financial crises or civil wars (see Hsiang and Jina (2014) for comparisons). The dust storm effects occur in the year of exposure or the year after, with no sign of recovery afterward, in contrast to cyclones which exhibit negative impacts many years after the exposure.

Figure 12 shows the summed effect, the Ω_j over 5 years. While the standard errors become large farther out, the point estimates are between 4 and 10 percentage points of lost GDP. These are large and important effects for a normal year-to-year variation in dust exposure. The persistence of the effects is also notable. Apart from studies of the dust bowl, the economic effects of dust storms have largely been ignored, and these results show that they are likely to be important in regions that experience these storms frequently. Just as depressed growth in areas of the Pacific can be partly explained by tropical cyclone activity, depressed level of economic growth across West Africa can be in part explained by the presence of dust storms. A one standard deviation, as calculated as the average within-country standard deviation, is relatively small when comparing dust levels across countries. This is equivalent from moving from the average population-weighted yearly exposure of

Mali to Chad (see Figure 7), but there is a range of over 14 standard deviations between the country with the lowest exposure in the sample, Gabon, to the highest exposure, Niger. While the model is estimated using within-country differences, extrapolating linearly across countries would imply significant reductions in economic growth in the most dust-exposed countries, compared to less exposed places.

When relating dust to nightlight measures, it is important to be careful that the dust in the atmosphere is not physically blocking light from the Earth's surface from reaching the satellites that are measuring the nightlights, as this will not measure changes in economic activity. If this were the case, it could explain a drop in nightlights the same year that dust increases, but then should immediately recover the next year. Since this is not the case and the results including nightlights look similar to the results for only GDP look similar, this is unlikely to be a major concern. One reason this might be expected is that dust storms tend to occur during the afternoon over this part of the world, although some dust will certainly remain in the atmosphere into the night.

One reason that dust storms could have persistent effects on economic activity is their ability to destroy physical capital. While data on depreciation rates at the macro level may not be well-measured, I can use the data from the World Development Indicators to test if there is any effect on the measured rate. As shown in Table 14, the depreciation rate is estimated to increase in response to an increase in dust exposure over 1-3 lags. Electricity consumption is also estimated to decrease, and the number of unlit cells within the country is estimated to increase in response to an increase in dust exposure.

To see the total long-run effects of these dust storms, it is useful to run counterfactuals. I compare countries' actual growth with model-predicted growth and a simulation where each country is set to its minimum value of predicted dust for each year. The results for selected countries are shown in figure 4. Over thirty years, the countries with highly variable dust exposure such as Benin and Niger show GDP per capita approximately 25% under observed levels of dust compared to if they had experienced their lowest observed level each year.

4.3.1 Agriculture

Equation 2 is also estimated for cereal yields in kg/ha as reported by the FAO, only including current year dust. The results are reported in Table 3. The effects are estimated to be a 2-3% decrease in yields for a one standard deviation increase in dust exposure for most of the crops. It should be noted that most agriculture takes place during the wet season, when there are typically fewer dust storms. However, the dry season dust exposure could still reasonably affect agricultural production. Dust storms tend to erode the soil, and may be depositing less productive, sandy soils. They can also cause irrigation canals to fill with sediment, making irrigation less effective. The losses in yields can explain part of the GDP result, but indicate that the damage done outside of agriculture are likely even larger, in order to explain the effects on GDP that are larger than on yields.

The estimated yield losses could be due both to the direct physical impacts of dust, as well as farmers deciding to change their behavior regarding time and effort spent tending or harvesting crops in response to the observed environmental factors over the growing season. These effects are explored further in section 4.5. It is also possible that yields at the country level are measured with significant error or that harvested areas decrease in response to a negative environmental shock, which would mean that yields do not capture the full extent of the harm to agriculture.

4.4 Micro-level results

4.4.1 Household production

To investigate further the effects of dust storms on economic activity, I consider the effects on agriculture, consumption, and work behavior at the household level. It should be noted that households and firms could be taking actions mitigating the effects of these storms. Such mitigation could take varying forms, some of which are observable and some of which are not. All estimated responses are net of any unobserved mitigation.

The LSMS data has detailed information on land used in agriculture by households as well as production by these households. There is a panel of 4,087 households that report area worked for each crop that they grow as well as total production of each crop in both

2011 and 2014. With this data, it is possible to construct a yield estimate for each crop that the household produces. I then run regressions of the following form to estimate the effect of dust storms on agricultural yields for each given crop:

$$Y_{it} = \beta D_{r(i),t} + \rho C_{r(i),t} + \alpha_{r(i)} + \eta_t + \epsilon_{it} \quad (5)$$

where Y_{it} is the log(yield) for household i in year t , $D_{r(i),t}$ is the mean dust exposure in region $r(i)$ in year t , $C_{r(i),t}$ controls for a quadratic of both temperature and precipitation, and $\alpha_{r(i)}$, η_t are region and year fixed effects or household and year fixed effects. Note that the regression is unweighted, meaning that this is the average effect on a household's yield, not necessarily the average yield effect on total production. The results are shown in Table 4. There appear to be relatively small effects of dust storms on agriculture, with the most common crop, millet, experiencing a decrease in yield of 1-2% for a one standard deviation increase in dust exposure. These are large effects, and could be driving the decline in total macroeconomic production.

I again test for nonlinearities in the response. I bin the dust exposure at the day-pixel level into deciles. The results for this model with household fixed effects are shown in Figure 18. The figure shows the percentage change in yields for spending the entire growing season in a given decile. The effects begin to appear in the seventh decile, with the negative effects of dust being clearly larger in the top three deciles. This provides some evidence that it is really the days of intense dust storms that are responsible for the damages. This is not what would be expected if the main channel is through the blocking of light and hindrance of photosynthesis, as this effect would be approximately linear in dust exposure.

Other activities of households may also be impacted. Individuals may become sick as a result of dust storms, which may lead them to miss work due to their illness. I can test for this, as individuals are asked about their recent illnesses for the past four weeks. Using the dust exposure over the past four weeks, I estimate the increase in illness experienced by each individual. I am again able to include individual fixed effects, as well as control for age. As shown in Table 6, there does not appear to be any increase in probability of sickness, nor a change in number of days of work missed in the past four weeks. Overall, it appears the

health channel is not likely to substantially affect economic output.

Consumption patterns may also change in response to dust storms. This could result from changes in income, increases in health care spending, or increases in other defensive spending or replacement of assets. The results of consumption on different categories of expenditure are shown in Figure 5. While there are statistically significant changes, they are not large as a percentage of expenditures. Overall consumption decreases, consistent with a negative income shock that cannot be completely smoothed. However, the fact that there is some evidence of income smoothing provides some evidence that dissaving could be occurring. If this is the case, this could provide a possible mechanism for the observed persistence in the GDP reductions. Repeated shocks to income that must be compensated for by a reduction in savings could lead to lowered investment rates, leading to lower long-run economic growth.

4.5 Labor allocation

To formalize the labor allocation decision of households which are facing a productivity shock due to environmental factors, I set up a simple model of household production and utility.

Households have a utility function over consumption C and leisure L , $U(C, L)$ with $\frac{\partial U}{\partial C} > 0$, $\frac{\partial U}{\partial L} > 0$, $\frac{\partial^2 U}{\partial C^2} < 0$, $\frac{\partial^2 U}{\partial L^2} < 0$. The consumption good is produced according to $Q = Af(H)$ with $f'(H) > 0$, $f''(H) < 0$ where H is the amount of labor and A is a productivity parameter.

The household's maximization problem implies $\frac{\partial U}{\partial L}P_C = \frac{\partial U}{\partial C}w$ where P_C is the price of the consumption good and w is the wage. In equilibrium, the wage will equal the marginal product of labor, so $w = Af'(H)$. The price of the consumption good can be normalized to 1, as can the total time allotment of the household, so that $H + L = 1$. In equilibrium,

$$\frac{\partial U}{\partial L}(C^*, L^*) = \frac{\partial U}{\partial C}(C^*, L^*)Af'(H^*) \quad (6)$$

Suppose there is a shock to productivity such that the new productivity is $A' < A$. Before time use can be reallocated, the right hand side could be larger or smaller, since $A' < A$, but $\frac{\partial U}{\partial C}(C_{new}, L^*) > \frac{\partial U}{\partial C}(C^*, L^*)$ since $\frac{\partial^2 U}{\partial C^2} < 0$. Because of this, the direction of change in leisure is ambiguous.

To test how leisure changes in response to a productivity shock caused by dust, I regress days worked in each sector on growing season dust exposure, temperature, and precipitation, as well as household and year fixed effects. The results for the agricultural and outside sectors are shown in Figure 6. Consistent with the prediction, an increase in dust leads to more days spent on household enterprises. Days spent in agriculture also increase, indicating that the need to make up for lost income drives an overall increase in labor. However, the increase in time spent on enterprises is larger, indicating higher marginal returns to this sector under increased dust exposure. As shown in Table 7, wages in both sectors decrease, with wages outside of agriculture decreasing by about 21% and agricultural wages not significantly changing. While this is a surprising result, there are a number of possible explanations. The composition of workers looking for wages in the agricultural sector could be different from those working on their own farms. It is also possible that more labor is needed in the agricultural sector to try to recover from the damage caused by the dust storms. Since dust storms can also reduce productivity outside of the agricultural sector, it could be a combination of more demand for labor in agriculture and lowered worker productivity, due to electricity outages or transportation network disruption.

4.6 Robustness

4.6.1 Macroeconomic Results

I run several tests to check the robustness of the results. First, I test for the validity of the IV strategy and using the dry season dust. As shown above, the relationship between source dust and dust over West Africa is strongest in the dry season. Since the dust in the source region shows variation across the entire year, I can use wet season dust as a placebo test. This test provides evidence both that the results are not driven by incorrect specification of time trends, as well as confirms that the dry season dust exposure is driving the results, not total year dust.

Table 11 shows the reduced-form relationship between dry season Bodélé dust and economic growth across West Africa. As expected, the results are significant across all GDP datasets and smaller than the IV estimates. Table 12 shows the reduced form relationship

between wet season Bodélé dust and growth. Here, the results are even smaller and not statistically significant, though some of the lagged values are. While there will be some dust that reaches West Africa during the wet season, the amount is significantly less. The wet and dry season dust are also correlated, making this test less than perfect. This provides some evidence that it is the dry-season dust that is important for the effects on growth.

Next, I run the 5-year model “backwards,” using five leads instead of five lags. The leads should be insignificant and close to zero, and that is what I find, as shown in Table 13. Note that while these estimates could be included in a model that also includes the lagged values, the limited time span of the data means that these estimates quickly become imprecisely estimated with the inclusion of additional lags. The confidence intervals in these models are also quite wide, but the contemporary and one-year lag values remain mostly negative. Importantly, they show a pattern distinct from that in Table 10, which shows consistently negative effects across the GDP datasets.

5 Conclusion

Economic production across West Africa is found to be negatively impacted by dust storm activity. Negative effects of dust storms are also found on agricultural production. Household shifts in labor demonstrate some adaptation to environmental impacts, but it does not eliminate them. Other aspects of household behavior including consumption and total labor supply are affected, but the effects are not large. Climate change may lead to increased dust activity in this region, and more broadly across the world, implying that the social cost of carbon should include these effects.

Policies for reducing dust storm impacts, much like for reducing the impacts of climate change more generally, can be considered in two parts: mitigation and adaptation. Both are being implemented across the world, yet little work has been done to assess the cost effectiveness of policies on either side.

On the side of mitigation, efforts have been undertaken to prevent desertification, which could mitigate the occurrence of dust storms. Large scale tree-planting projects were implemented in the United States, China, and across the African Sahel. However, the effectiveness

of these programs has not been clear. Soil conservation programs have had some success in preventing the wind erosion of soil, which implies that they are decreasing dust concentrations downwind. Efforts to promote soil conservation may be one effective means of reducing dust storm activity (Skidmore, 1986).

Policies promoting adaptation could take a number of forms. Government transfers to the most affected groups could offset the damages, and potentially reduce the long-run effects. To improve adaptive behavior at the time of impact, early warning systems could be put in place. Baek et al. (2015) have shown that early warning systems provide some benefits for child health outcomes in Korea. After being warned of the potential dangers and receiving warning of an impending storm, pregnant mothers are able to stay inside, which significantly mitigates the effects on the unborn child. Innovative strategies for reaching mothers in a less developed setting could add significant value, particularly if advanced warning can have similar effects here. Further research could be done to investigate the effectiveness of such early warning systems.

Further work should also be done to examine other potential mechanisms and social consequences of dust storm activity. One interesting question is whether changes in dust storm frequency or intensity drives migration decisions. This could affect both economic growth outcomes as well as influence social unrest and conflict. Other research could explore what populations are most affected by negative environmental shocks. If these groups can be identified, there may be a role for government intervention to provide assistance. Quantification of the impacts on the transportation sectors and on capital depreciation would also shed light on potential channels through which dust storms affect economic activity in the longer run.

References

- V. Arthi. "The Dust Was Long in Settling": Human Capital and the Lasting Impact of the American Dust Bowl. *The Journal of Economic History*, 78(1):196–230, 2018. ISSN 1098-6596. doi: 10.1017/S0022050718000074.
- D. Baek, D. T. Altinga, and N. Mocan. Chinese Yellow Dust and Korean Infant Health. *NBER Working Papers*, October(21613), 2015. ISSN 1098-6596. doi: 10.1017/CBO9781107415324.004.
- M. Birjandi-Feriz and K. Yousefi. When the Dust Settles : Productivity and Economic Losses Following Dust Storms. 2018.
- V. Buchard, A. M. da Silva, C. A. Randles, P. Colarco, R. Ferrare, J. Hair, C. Hostetler, J. Tackett, and D. Winker. Evaluation of the surface PM_{2.5} in Version 1 of the NASA MERRA Aerosol Reanalysis over the United States. *Atmospheric Environment*, 125:100–111, 2016. ISSN 18732844. doi: 10.1016/j.atmosenv.2015.11.004. URL <http://dx.doi.org/10.1016/j.atmosenv.2015.11.004>.
- V. Buchard, C. A. Randles, A. M. da Silva, A. Darmenov, P. R. Colarco, R. Govindaraju, R. Ferrare, J. Hair, A. J. Beyersdorf, L. D. Ziemba, and H. Yu. The MERRA-2 aerosol reanalysis, 1980 onward. Part II: Evaluation and case studies. *Journal of Climate*, 30(17): 6851–6872, 2017. ISSN 08948755. doi: 10.1175/JCLI-D-16-0613.1.
- M. Burke, S. M. Hsiang, and E. Miguel. Global non-linear effect of temperature on economic production. *Nature*, 527(7577):235–239, 2015. ISSN 14764687. doi: 10.1038/nature15725.
- C. A. Randles, A. M. da Silva, V. Buchard, A. Darmenov, P. R. Colarco, V. Aquila, H. Bian, E. P. Nowottnick, X. Pan, A. Smirnov, H. Yu, and R. Govindaraju. The MERRA - 2 Aerosol Assimilation. *Technical Report Series on Global Modeling and Data Assimilation*, 45(December), 2016. URL <https://gmao.gsfc.nasa.gov/pubs/docs/Randles887.pdf>.
- Center for International Earth Science Information Network - CIESIN - Columbia University. Gridded Population of the World, Version 4 (GPWv4): Administrative Unit Center Points with Population Estimates, 2016.

- T. Deryugina, G. Heutel, N. H. Miller, D. Molitor, and J. Reif. The Mortality and Medical Costs of Air Pollution: Evidence from Changes in Wind Direction. *NBER Working Paper Series*, page 55, 2016. doi: 10.3386/w22796. URL <http://www.nber.org/papers/w22796>.
- A. Goudie and N. Middleton. *Desert Dust in the Global System*. Springer, Berlin, 2006. ISBN 9783540323549.
- J. V. Henderson, A. Storeygard, and D. N. Weil. Measuring economic growth from outer space. *American Economic Review*, 102(2):994–1028, 2012. ISSN 00028282. doi: 10.1257/aer.102.2.994.
- R. Hornbeck. The enduring impact of the American Dust Bowl: Short- and long-run adjustments to environmental catastrophe. *American Economic Review*, 102(4):1477–1507, 2012. ISSN 00028282. doi: 10.1257/aer.102.4.1477.
- S. M. Hsiang and A. S. Jina. The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence from 6,700 Cyclones. 2014. URL <http://www.nber.org/papers/w20352.pdf>.
- M. Jerven. Poor Numbers: How We Are Misled by African Development Statistics and What to Do about It. *Cornell University Press*, IX:176, 2013. ISSN 0225-5189. doi: 10.1080/02255189.2014.877382.
- S. Johnson, W. Larson, C. Papageorgiou, and A. Subramanian. Is newer better? Penn World Table Revisions and their impact on growth estimates. *Journal of Monetary Economics*, 60(2):255–274, 2013. ISSN 03043932. doi: 10.1016/j.jmoneco.2012.10.022.
- J. M. Prospero, P. Ginoux, O. Torres, S. E. Nicholson, and T. E. Gill. Environmental characterization of global sources of atmospheric soil dust identified with the Nimbus 7 Total Ozone Mapping Spectrometer (TOMS) absorbing aerosol product. *Reviews of Geophysics*, 40(1):1–31, 2002. ISSN 87551209. doi: 10.1029/2000RG000095.
- W. Schlenker and W. R. Walker. Airports, air pollution, and contemporaneous health. *Review of Economic Studies*, 83(2):768–809, 2016. ISSN 1467937X. doi: 10.1093/restud/rdv043.

- E. L. Skidmore. Wind erosion control. *Climatic Change*, 9(1-2):209–218, 1986. ISSN 01650009. doi: 10.1007/BF00140537.
- R. Stefanski and M. V. K. Sivakumar. Impacts of Sand and Dust Storms on Agriculture and Potential Agricultural Applications of a SDSWS. *IOP Conference Series: Earth and Environmental Science*, 7:012016, 2009. ISSN 1755-1315. doi: 10.1088/1755-1307/7/1/012016. URL <http://stacks.iop.org/1755-1315/7/i=1/a=012016?key=crossref.58fe54d6f248b3c1a161646beff344f7>.
- J. G. Zivin and M. Neidell. The Impact of Pollution on Worker Productivity. *American Economic Review*, 102(7):3652–3673, dec 2012. ISSN 0002-8282. doi: 10.1257/aer.102.7.3652. URL <http://pubs.aeaweb.org/doi/abs/10.1257/aer.102.7.3652>.

Tables and figures

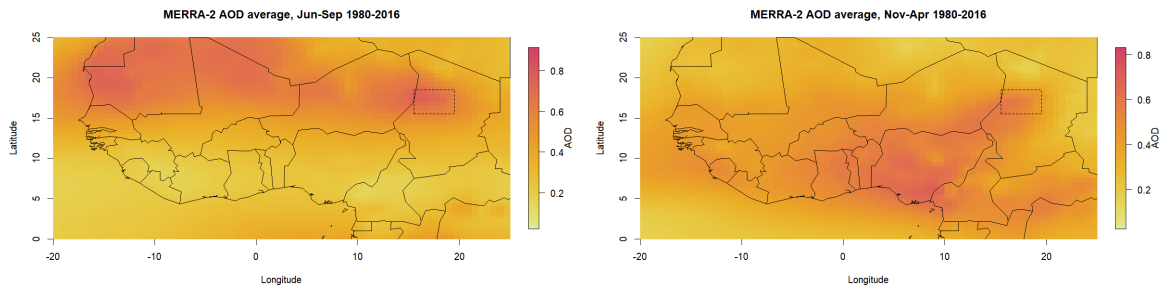


Figure 1: Mean AOD in the wet season (Jun-Sep) and dry season (Nov-Apr) over West Africa. The dashed box indicates the region of the Bodélé depression.

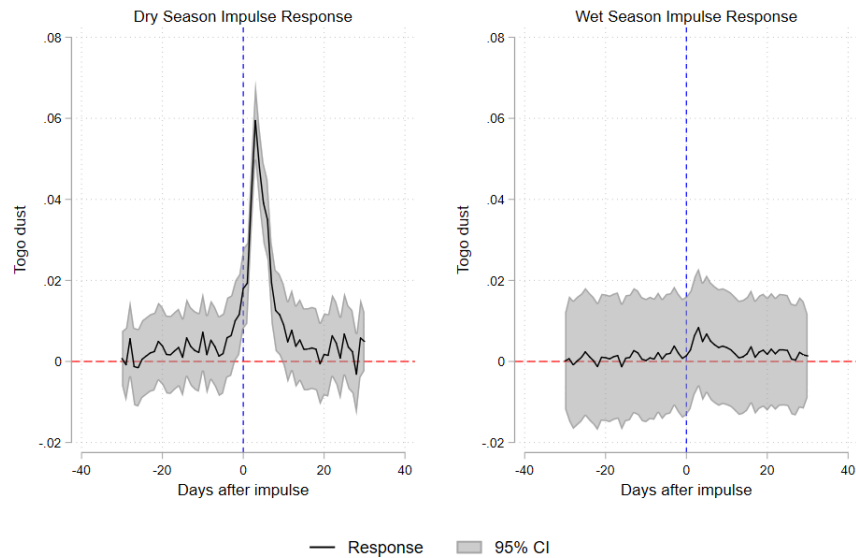


Figure 2: Response of a location in Togo AOD to a 1 unit increase in AOD over the Bodélé depression.

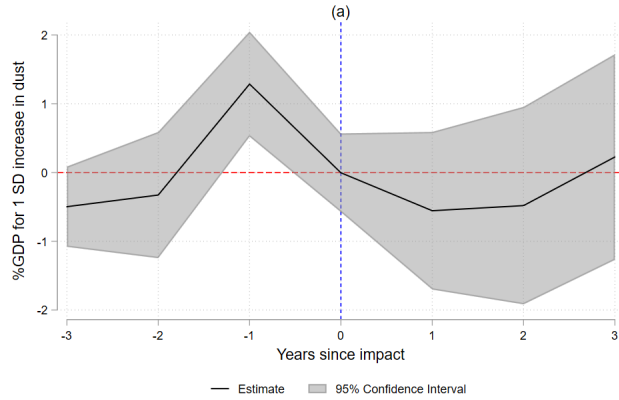


Figure 3: OLS results of Equation 2. If interpreted causally, the result would absurdly imply that an increase in dust exposure causes an increase in GDP growth the previous period. This in fact provides evidence that changes in GDP leads to changes in particulates in the subsequent period.

VARIABLES	(1) Henderson	(2) Penn	(3) Maddison	(4) WDI
AOD	-0.650** (0.318)	-1.428** (0.663)	-0.215 (0.195)	-1.052*** (0.317)
lag AOD	-0.302 (0.227)	-0.630 (0.470)	-0.346 (0.280)	-1.096*** (0.382)
Observations	380	380	380	380
R-squared	0.287	0.280	0.422	0.433
Country FE	X	X	X	X
Country quadratic time trend	X	X	X	X
<i>Contemporaneous + 1 Lag</i>				
Estimate	-0.951	-2.058	-0.560	-2.148
Std. err.	0.399	0.826	0.398	0.617
p-value	0.017	0.013	0.159	0.001
First stage F-stat	30.884	30.884	30.884	30.884

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 1: Estimates including contemporaneous and 1 lag of dust exposure. AOD is instrumented using the aggregated values from the statistical IV model.

VARIABLES	(1) Henderson	(2) Penn	(3) Maddison	(4) WDI
AOD	-0.862** (0.338)	-1.880*** (0.706)	-0.908*** (0.301)	-1.418*** (0.400)
lag AOD	-2.127*** (0.403)	-4.489*** (0.840)	-1.760*** (0.435)	-1.939*** (0.489)
Observations	380	380	380	380
R-squared	0.249	0.243	0.393	0.412
Country FE	X	X	X	X
Country quadratic time trend	X	X	X	X
<i>Contemporaneous + 1 Lag</i>				
Estimate	-2.989	-6.369	-2.668	-3.357
Std. err.	0.551	1.159	0.663	0.819
p-value	0.000	0.000	0.000	0.000
First stage F-stat	27.246	27.246	27.246	27.246

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 2: Estimates including contemporaneous and 1 lag of dust exposure. AOD is instrumented using the aggregated values from the wind-driven model.

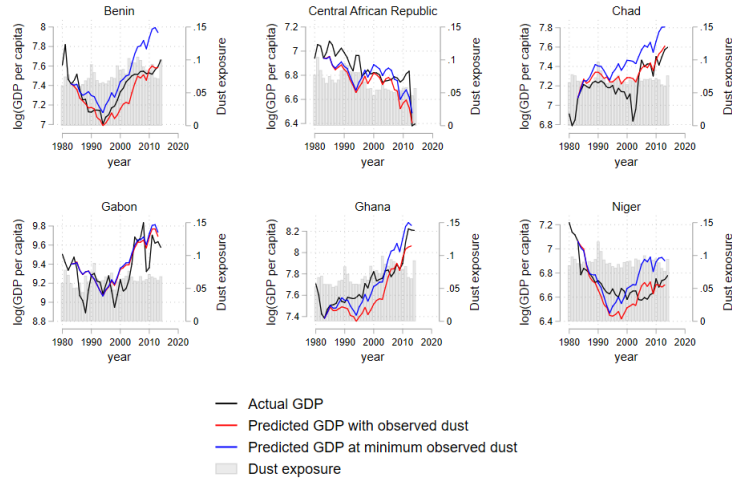


Figure 4: Counterfactual estimates of log(GDP per capita) under the model with observed dust and the minimum observed dust in the sample.

VARIABLES	(1) Cassava	(2) Cowpeas	(3) Groundnuts	(4) Maize	(5) Millet	(6) Sorghum
AOD	-2.794*** (0.952)	2.275 (1.846)	-2.713*** (0.997)	-3.036*** (1.019)	-3.105** (1.564)	-3.196*** (1.115)
Observations	451	218	490	455	455	455
R-squared	0.930	0.775	0.779	0.595	0.764	0.838
country FE	X	X	X	X	X	X
country quad trend	X	X	X	X	X	X

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3: Estimates including contemporaneous dust exposure. Estimates are the percentage change in yield for a 1 standard deviation increase in dry season dust exposure. AOD is instrumented using the aggregated values from the wind-driven model.

VARIABLES	(1) millet	(2) millet	(3) cowpeas	(4) cowpeas	(5) sorghum	(6) sorghum
AOD	-0.562 (0.346)	-0.678** (0.332)	-1.492* (0.805)	-1.846** (0.737)	-0.638 (1.103)	-1.442 (1.108)
Observations	2,472	1,150	987	400	840	348
R-squared	0.309	0.590	0.314	0.594	0.306	0.700
year FE	X	X	X	X	X	X
region FE	X		X		X	
household FE		X		X		X

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 4: Effect of dust on yields in Niger for the three most common crops reported by households, using LSMS data. The estimates are the percentage change in yield for a one standard-deviation increase in dust over the dry season prior to the growing season. The left-side estimates are estimated using cluster fixed effects and the right-side estimates are estimated using household fixed effects.

VARIABLES	(1)	(2)	(3)	(4)
	Probability of reporting sickness in past 4 weeks	Probability of reporting sickness in past 4 weeks	Probability of reporting sickness in past 4 weeks	Probability of reporting sickness in past 4 weeks
AOD	0.00162 (0.00413)	0.00310 (0.00462)	0.00162 (0.00413)	0.00310 (0.00462)
Observations	22,495	22,495	22,495	22,495
R-squared	0.567	0.568	0.567	0.568
person FE	X	X	X	X
year trend	X	X		
month FE	X	X	X	X
weather		X		X
year FE			X	X

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 5: Change in probability of self-reported illness over the past four weeks for a one standard deviation increase in AOD in those four weeks. In the sample, 20% of respondents report being sick in the past four weeks.

VARIABLES	(1)	(2)	(3)	(4)
	Number of days missed past 4 weeks	Number of days missed past 4 weeks	Number of days missed past 4 weeks	Number of days missed past 4 weeks
AOD	-0.000167 (0.00112)	0.000101 (0.00117)	-0.000167 (0.00112)	0.000101 (0.00117)
Observations	945	945	945	945
R-squared	0.574	0.576	0.574	0.576
person FE	X	X	X	X
year trend	X	X		
month FE	X	X	X	X
weather		X		X
year FE			X	X

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 6: Change in number of days missed over the past four weeks for a one standard deviation increase in AOD in those four weeks.

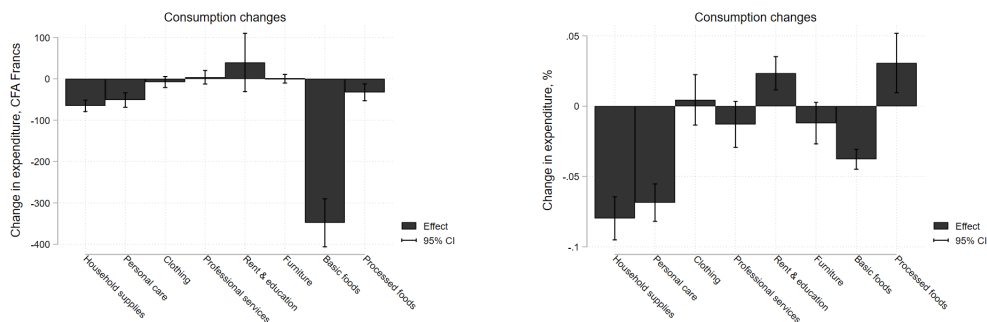


Figure 5: Effect of a one standard deviation increase in AOD on yearly expenditure on various categories, by level of expenditure (left side) and by percentage (right side).

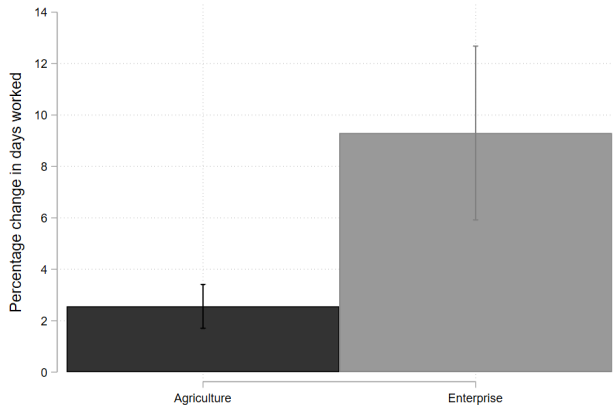


Figure 6: Change in days worked in each sector over a year for a one standard deviation increase in dry season dust exposure.

VARIABLES	(1) log(wage)	(2) log(wage)
AOD	-0.181*** (0.0163)	-0.0485 (0.0353)
Non-Ag x AOD		-0.168*** (0.0400)
Ag		24.09 (17.08)
Observations	1,070	1,070
R-squared	0.850	0.855
person FE	X	X

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 7: Effect of a one standard-deviation increase in dust exposure during the dry season on wages in agricultural and other sectors. The non-agricultural sector workers experience a significant drop in wages, while agricultural workers only see a small, statistically insignificant decline.

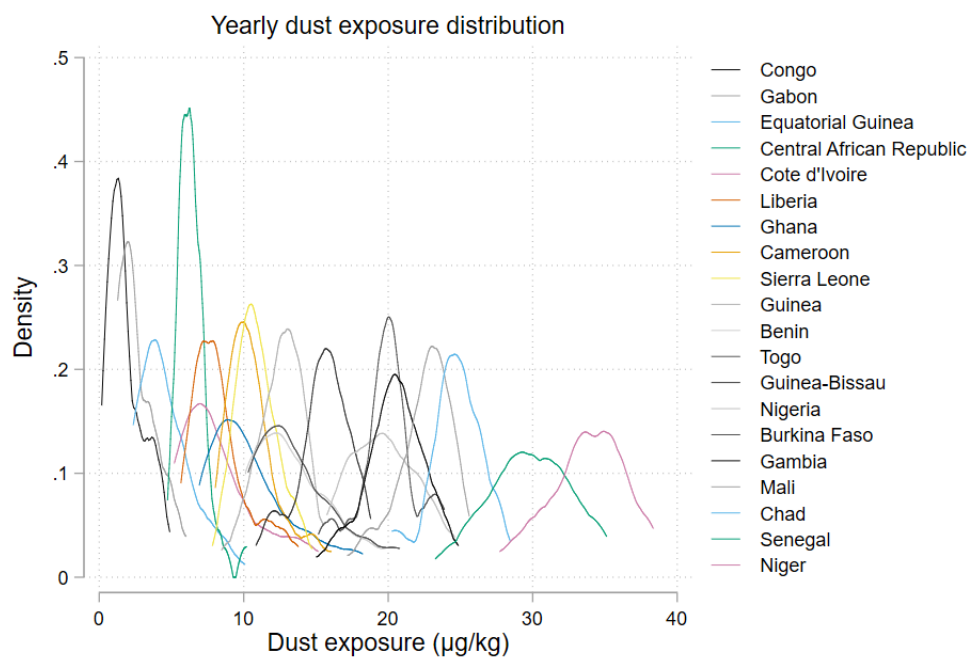


Figure 7: Yearly distribution of dust exposure by country. Legend is ordered by mean dust exposure.

A Additional Tables and Figures

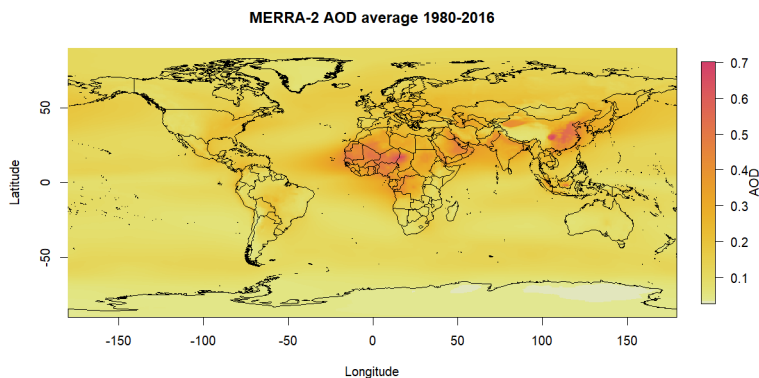


Figure 8: AOD is highest over deserts and eastern Asia.

VARIABLES	(1) L0 AOD	(2) L0 AOD	(3) L1 AOD
L0 pred AOD	0.753*** (0.069)	0.732*** (0.067)	-0.425*** (0.065)
L1 pred AOD		0.081 (0.057)	0.412*** (0.055)
Constant	-0.064 (0.150)	0.297* (0.175)	0.056 (0.168)
Observations	380	380	380
R-squared	0.946	0.949	0.954
country FE	X	X	X
country quad trend	X	X	X

Standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 8: Country-level first-stage regression for the statistical IV approach over the inclusion of up to 5 lags.

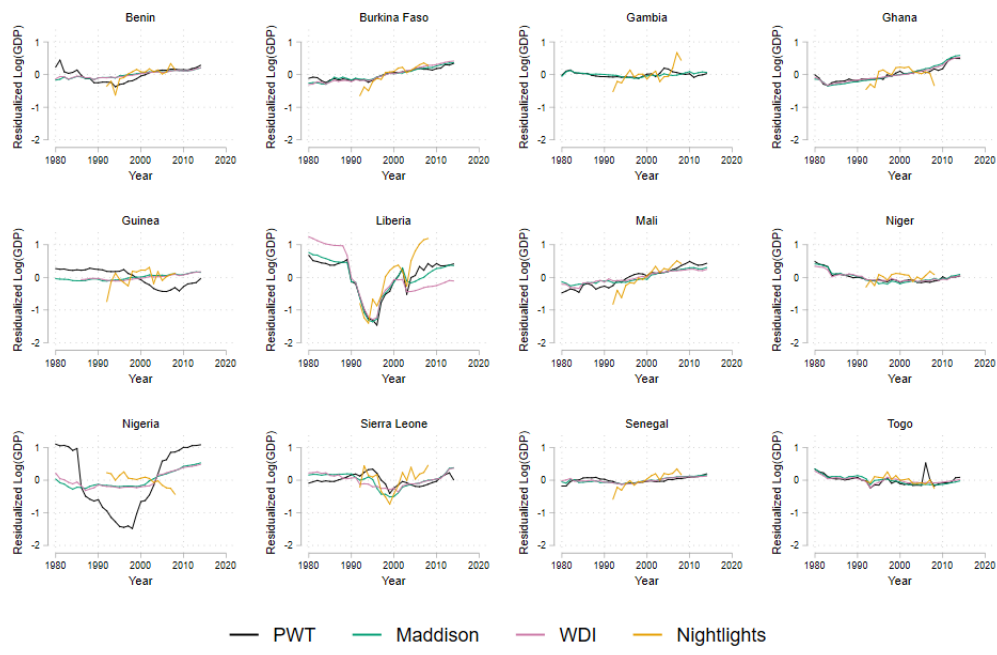


Figure 9: GDP comparison for West African countries.

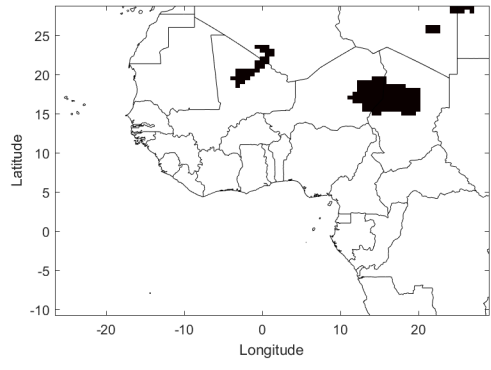


Figure 10: Major dust source regions, as identified by Prospero et al. (2002).

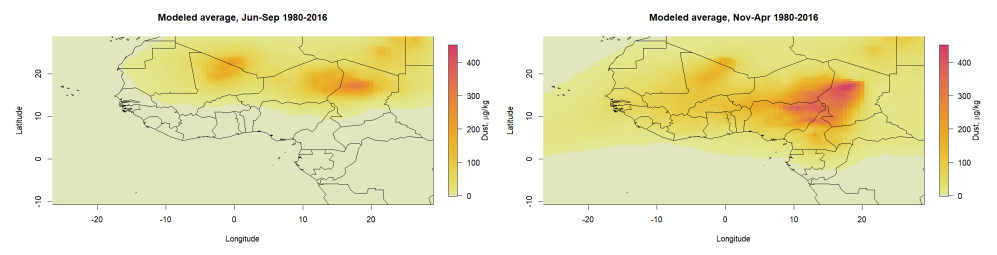


Figure 11: Dust estimates using wind-driven model. Average over the wet season (left) and dry season (right). Units are $\mu\text{g dust} / \text{kg air}$.

VARIABLES	(1) L0 AOD	(2) L0 AOD	(3) L1 AOD
L0 pred AOD	0.153*** (0.017)	0.196*** (0.016)	-0.086*** (0.016)
L1 pred AOD		0.012 (0.017)	0.122*** (0.017)
Constant	0.193 (0.157)	0.934*** (0.173)	-0.472*** (0.173)
Observations	380	380	380
R-squared	0.940	0.951	0.952
country FE	X	X	X
country quad trend	X	X	X

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 9: Country-level first-stage regression for the wind-driven approach over the inclusion of up to 5 lags.

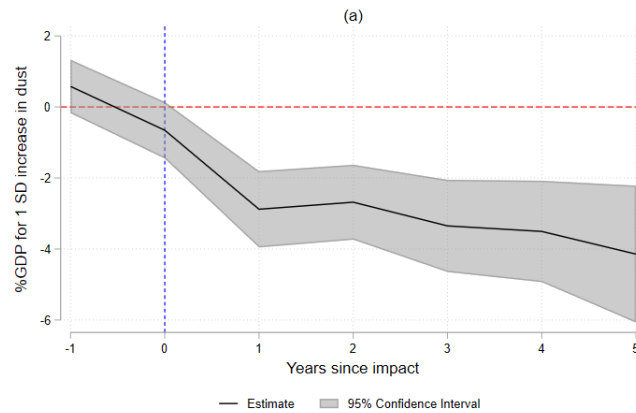


Figure 12: Effect of dust exposure on GDP growth over 5 years of lags, estimated using the wind-driven model. The Ω_j are plotted, showing the reduction in GDP over time using the Henderson et al. (2012).

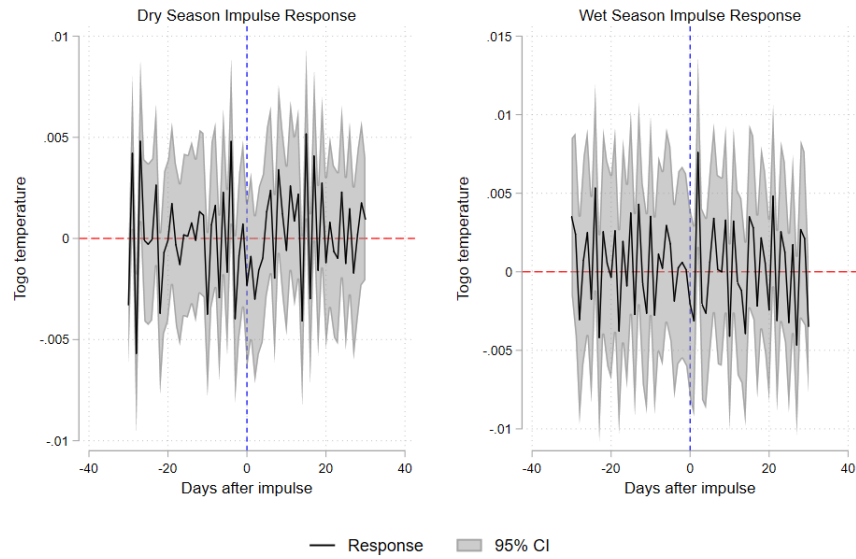


Figure 13: Response of temperature on the Togo coast to a 1 unit increase in AOD over the Bodélé depression.

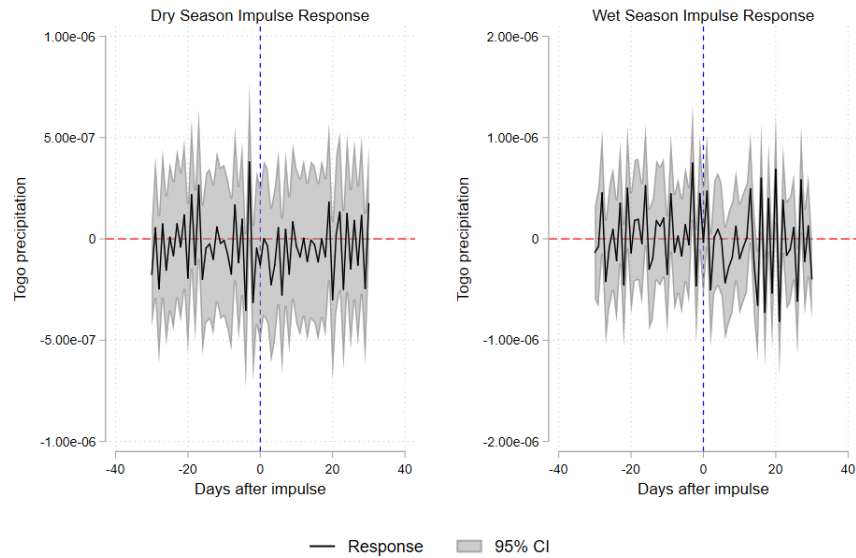


Figure 14: Response of precipitation on the Togo coast to a 1 unit increase in AOD over the Bodélé depression.

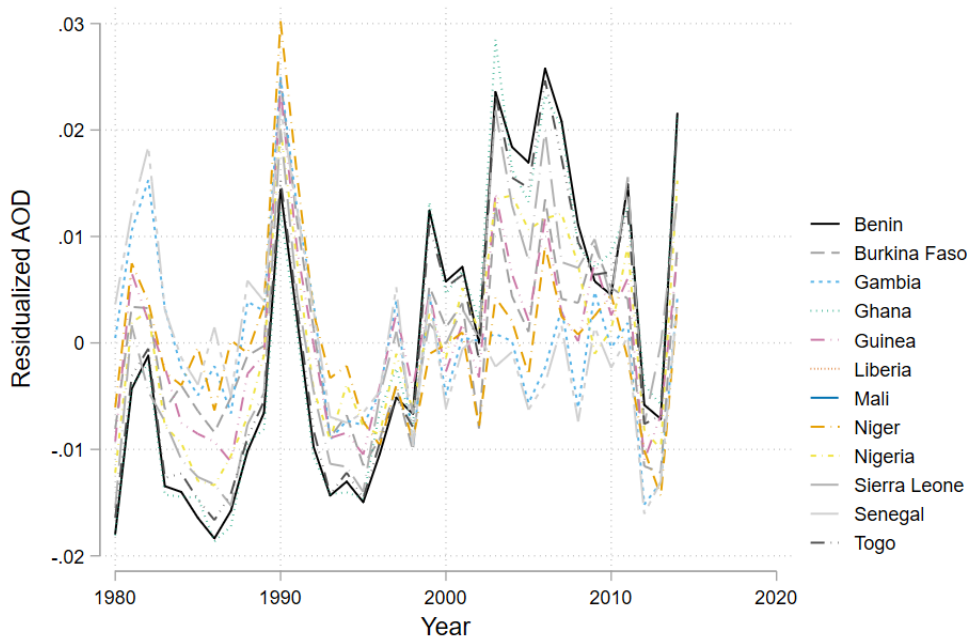


Figure 15: Measured population-weighted AOD over West African countries, net of country fixed effects.

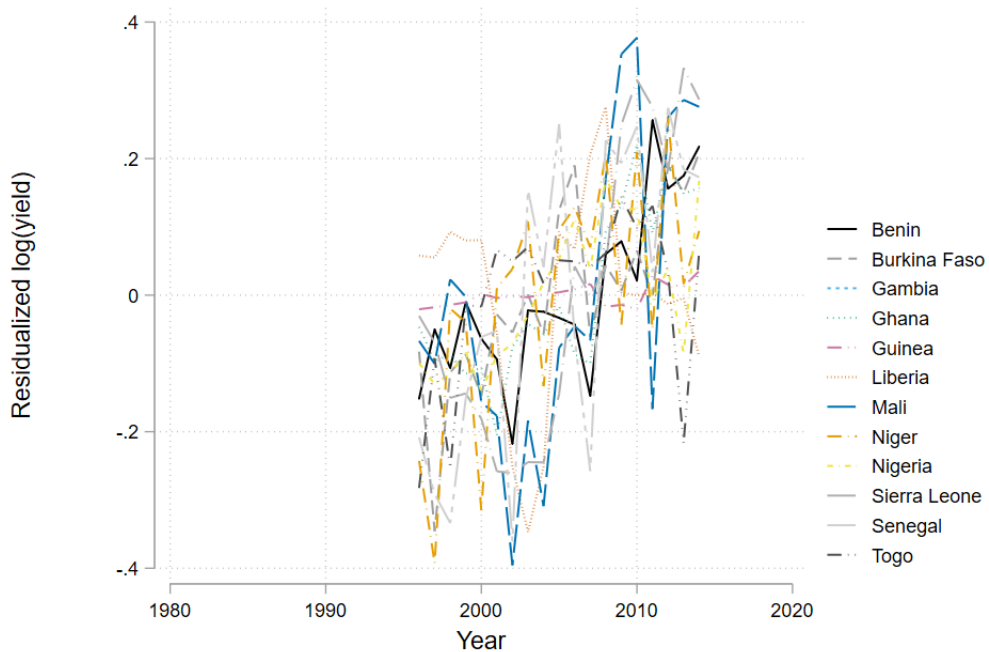


Figure 16: Log of average cereal yields in West African countries, net of country fixed effects.

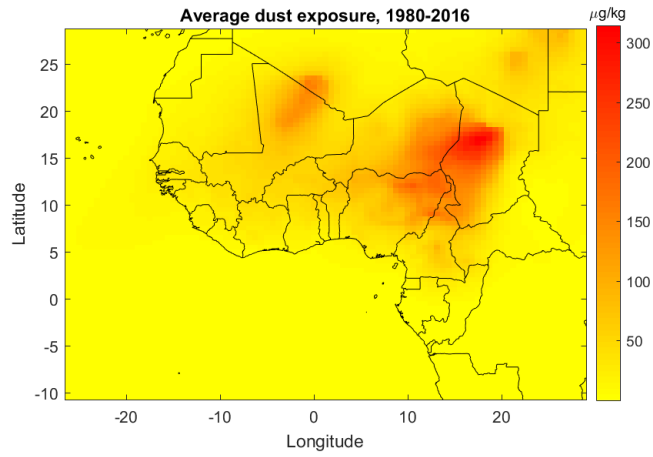


Figure 17: Mean AOD from 1980-2016 as calculated by the dust transport model, constrained by source area dust.

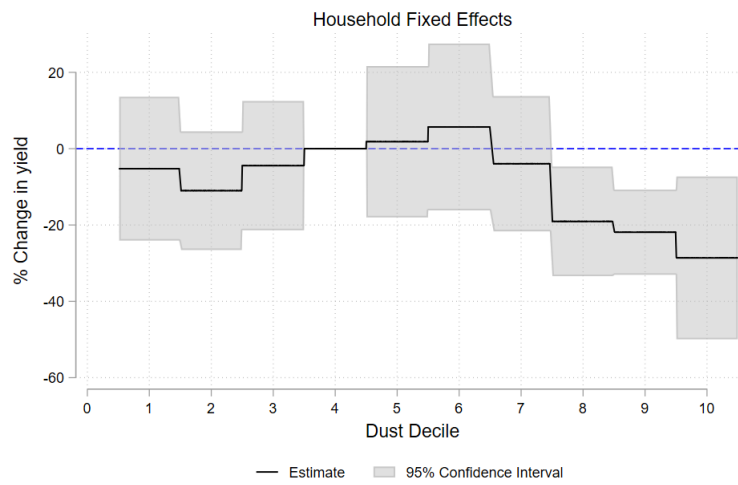


Figure 18: The effect of dust by dust decile on agricultural yields. Deciles are determined by the daily pixel-level exposure over the entire West African region.

VARIABLES	(1) Henderson	(2) Penn	(3) Maddison	(4) WDI
AOD	-0.727* (0.371)	-1.515** (0.767)	0.220 (0.428)	-0.543 (0.446)
lag AOD	-2.409*** (0.476)	-5.002*** (0.977)	-0.116 (0.507)	-0.432 (0.513)
L2 AOD	0.199 (0.371)	0.417 (0.762)	2.496*** (0.625)	2.729*** (0.689)
L3 AOD	-0.396 (0.373)	-0.881 (0.770)	-0.813** (0.400)	-0.160 (0.459)
L4 AOD	0.210 (0.341)	0.362 (0.698)	-1.130** (0.455)	-0.422 (0.393)
L5 AOD	-0.497 (0.395)	-1.006 (0.810)	-1.868*** (0.560)	-1.248** (0.513)
Observations	380	380	380	380
R-squared	0.286	0.283	0.300	0.351
Country FE	X	X	X	X
Country quadratic time trend	X	X	X	X
<i>Sum of all lags</i>				
Estimate	-3.619	-7.625	-1.211	-0.076
Std. err.	0.888	1.804	1.142	1.289
p-value	0.000	0.000	0.289	0.953
First stage F-stat	0.179	0.179	0.179	0.179

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 10: Estimates including 5 lags of dust exposure. AOD is instrumented using the aggregated values from the wind-driven model. L2-L5 are the two to five year lags of dry season dust exposure.

VARIABLES	(1) Henderson	(2) Penn	(3) Maddison	(4) WDI
dry season Bodélé AOD	-0.142 (0.140)	-0.329 (0.293)	0.054 (0.150)	-0.111 (0.151)
lag dry season Bodélé AOD	-0.259* (0.142)	-0.543* (0.292)	-0.135 (0.160)	-0.304 (0.204)
Observations	380	380	380	380
R-squared	0.280	0.197	0.266	0.240
Country FE	X	X	X	X
Country quadratic time trend	X	X	X	X
<i>Contemporaneous + 1 Lag</i>				
Estimate	-0.401	-0.872	-0.081	-0.415
Std. err.	0.192	0.397	0.221	0.243
p-value	0.038	0.029	0.714	0.089

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 11: Reduced-form estimates of growth on dry-season Bodélé dust. Estimates including contemporaneous and 1 lag of dust exposure.

VARIABLES	(1) Henderson	(2) Penn	(3) Maddison	(4) WDI
wet season Bodélé AOD	0.027 (0.112)	0.023 (0.066)	0.020 (0.033)	0.031 (0.029)
lag wet season Bodélé AOD	-0.039 (0.073)	-0.073 (0.071)	-0.104*** (0.035)	-0.133*** (0.038)
Observations	380	680	680	776
R-squared	0.300	0.099	0.156	0.173
Country FE	X	X	X	X
Country quadratic time trend	X	X	X	X
<i>Contemporaneous + 1 Lag</i>				
Estimate	-0.012	-0.050	-0.084	-0.102
Std. err.	0.133	0.089	0.043	0.043
p-value	0.929	0.575	0.053	0.016

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 12: Reduced-form estimates of growth on wet-season Bodélé dust. Estimates including contemporaneous and 1 lag of dust exposure.

VARIABLES	(1) Henderson	(2) Penn	(3) Maddison	(4) WDI
F1 AOD	-3.819 (6.232)	-2.347 (2.239)	1.235 (1.314)	0.984 (1.194)
F2 AOD	-0.308 (2.166)	0.384 (1.810)	1.303 (0.991)	0.285 (0.934)
F3 AOD	3.233 (6.242)	0.549 (1.793)	1.908* (0.995)	1.569* (0.926)
F4 AOD	4.084 (7.989)	-0.581 (1.677)	1.019 (0.905)	0.840 (0.862)
F5 AOD	2.539 (4.682)	-0.810 (1.667)	0.065 (0.907)	0.159 (0.839)
Observations	318	609	609	690
R-squared	-0.323	0.084	0.232	0.316
Country FE	X	X	X	X
Country quadratic time trend	X	X	X	X
<i>Sum of all leads</i>				
Estimate	5.729	-2.806	5.529	3.836
Std. err.	14.135	4.759	2.838	2.587
p-value	0.686	0.556	0.052	0.139

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 13: Estimates yearly leads of dry-season dust exposure. Dust exposure is instrumented using the aggregated values from the wind-driven model.

VARIABLES	(1) Depreciation rate	(2) Electricity consumption per capita (kWh)	(3) Number of unlit cells in DMSP
AOD	-0.048* (0.026)	-19.780** (9.778)	2,017.314 (2,015.073)
lag AOD	-0.039 (0.033)	-35.043 (24.058)	786.064 (2,462.565)
L2 AOD	-0.046 (0.037)	-31.135 (26.075)	-800.755 (1,956.041)
L3 AOD	0.009 (0.033)	-9.311 (15.376)	-2,114.854* (1,166.805)
Observations	640	135	374
R-squared	0.972	0.991	1.000
country FE	X	X	X
country quad trend	X	X	X

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 14: Estimates including contemporaneous dust exposure. Estimates are the change for a 1 standard deviation increase in dry season dust exposure. AOD is instrumented using the aggregated values from the wind-driven model.

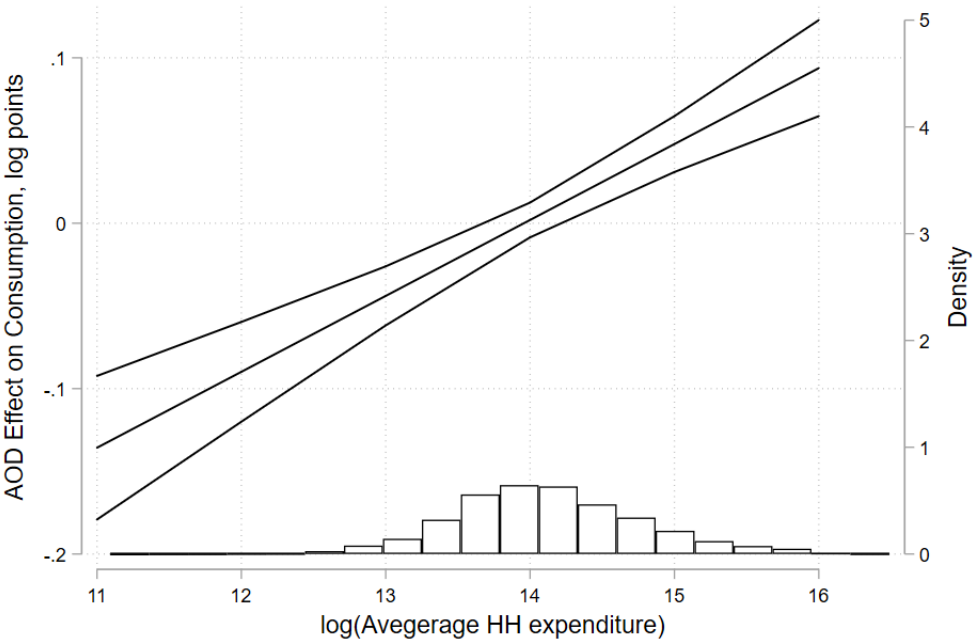


Figure 19: Effect of AOD on household consumption interacted with average household consumption.