

Isolating the Effects of Pollution Externalities from Natural Resource Extraction on Local Agricultural Output in Africa

Mira Korb

October 6, 2024

Abstract

This paper explores the extent to which pollution externalities and market-based effects from industrial mining affect local agricultural output in Sub-Saharan Africa. I combine mine geolocations, topographical data and satellite-based measures of pollution, yields and weather, to identify areas around mines that are disproportionately exposed to pollution but not market-specific effects. Leveraging the staggered openings of mines across Sub-Saharan Africa, I find that air and water pollution externalities from mines account for almost 50% of the overall reduction in yields caused by industrial mining. Finally, I document that pollution externalities are larger for mines in countries with poor governance and regulatory environments, as well as for mines located in areas with low initial levels of pollution.

1 Introduction

“It would be fortuitous if the choices that maximize economic prosperity also maximized environmental quality. The inconvenient truth is that such a perfect alignment rarely occurs” (Jayachandran, 2022). With burgeoning industrialization, many low and middle income countries face this important potential trade off between economic development and environmental quality. This tension is particularly salient when evaluating how increased industrialization in developing countries affects local livelihoods, specifically through market-based channels and pollution externalities.

Market-based channels can affect local livelihoods through shocks to production inputs or consumption goods, which in turn affects prices or the allocation of inputs across sectors. In contrast, pollution externalities will affect the degree to which inputs can be used productively. To interpret whether the “net” effect of industrialization on local livelihoods is positive or negative, it is valuable to disentangle the role that market-based versus pollution channels play in this overall effect. For example, if a new industry primarily affects local development by reallocating labor out of agriculture towards manufacturing or services sectors, with minimal environmental impacts, this might be interpreted as positive evidence of structural transformation. Conversely, large negative environmental impacts with minimal wage gains might suggest that the new industry is doing more harm than good for the community.

In this paper, I investigate the extent to which extractive industries, specifically the mining and mineral processing sectors, affect local agricultural output in Sub-Saharan Africa through pollution externalities versus market based channels. I show that on average, air and water pollution externalities account for approximately 44% of the total reduction in local yields caused by industrial mining. Mining-induced water pollution is the main driver of this effect, with the effects of water pollution over 10 times as large as the effects of air pollution on yields.

There are several reasons why the mining sector in Sub-Saharan Africa is an important setting to examine the development-environmental quality trade off. To begin, a majority of the countries in Sub-Saharan Africa are heavily reliant on both agriculture and natural resources, such as metals, minerals, oil and gas. Of the 48 countries in Sub-Saharan Africa, 26 are considered “resource-rich” by the International Monetary Fund, meaning at least 20% of exports or fiscal revenue come from natural resources [Cust and Zeufack, 2023]. Furthermore, the number of resource-rich countries in Sub-Saharan Africa has risen over time and is expected to continue to rise due to a combination of new discoveries, new production and rising resource prices [Cust and Zeufack, 2023]. At the same time, over 60% of the population in this region work as small-holder farmers, with 23% of the GDP in Sub-Saharan Africa coming from agriculture (CITE).

Additionally, while mining has increased wealth [von der Goltz and Barnwal, 2019] and reduced poverty [Aragón and Rud, 2013], it has also reduced agricultural production in nearby areas [Aragón and Rud, 2016]. Importantly, the literature is yet to explore the extent to which this decline in agricultural output is driven by market-based channels versus pollution externalities. While this decline may be spurred by economic transitions to more profitable industries or sectors (Kotsadam and Tolonen [2016]; Huang et al. [2023]), it may also be driven by pollution. Literature in the biological sciences documents how air and water pollution from industrial activity can substantially reduce local crop yields (Choudhury et al. [2017]; Maggs et al. [1995]).

To examine the extent to which industrial mining affects local agricultural output through pollution versus market-based channels, I combine a large geo-spatial dataset of mines across Sub-Saharan Africa from S&P Global Market Intelligence, with topographical data as well as satellite-based measures of air, water pollution, yields and weather. I leverage the fact that pollution externality effects are directional but market-based effects are non-directional. That is, pollution externalities will disproportionately affect locations near mining areas that are more likely to be exposed to pollution, for instance by being located downwind or downstream, but market effects will not disproportionately affect these same areas.

First, I use the direction of water flow from rivers near mines to identify areas that are exposed to water pollution from mines. Given mine coordinates, I define upstream and downstream cropland buffers by selecting mines within 1 kilometer of a river and buffering the sides of river segments that are upstream and

downstream from the mine. The cropland in these buffers represents areas that might be affected by river water quality through seepage of the water into surrounding areas or use of the river water for irrigation. I define the downstream side of each mine as “treated” and the upstream side as a “control.” In addition, I measure water quality in the upstream and downstream river segments near mines by applying a novel, remotely sensed indicator, the normalized difference turbidity index (NDTI), which is based on spectral reflectance of water pixels within rivers near mines.

Second, I use highly localized, daily data on wind direction to identify areas near mines that are more likely to be exposed to air pollution. As wind blows across a mine site, pollutants such as particulate matter or sulfur dioxide from activities such as drilling, blasting or smelting are more likely to be transported into the downwind side of a mine. I define a buffer of radius 60-kilometers around the mine centroid and divide this buffer into four sides indicating the cardinal directions that the wind can blow from the mine. The intensity of air pollution exposure experienced by each side is proxied by the number of days the side is downwind from a mine in a given month, while air pollution is measured using remotely sensed aerosol optical density (AOD). Sides receiving more wind from a mine are “treated,” while sides receiving no wind or less wind serve as “controls.”

Within the air and water pollution treatment and control areas, I construct a measure of remotely-sensed yields, the normalized difference vegetation index (NDVI), which is highly correlated with yields from survey and administrative data [Lobell et al., 2022]. To estimate the effect of mining-induced water pollution on yields, I use a difference-in-difference (DID) design that compares NDVI in areas around the river on the downstream side to that of the areas on the upstream side of the mine, before and after a mine opening. Similarly, to estimate the effect of mining-induced air pollution on yields, I use a continuous DID design that compares NDVI on sides of a mine that are downwind more frequently to sides of the same mine that are downwind less frequently, before and after a mine opening.

For both of these DID designs, the treated area (downstream or downwind) is disproportionately exposed to pollution from the mine relative to the control area (upstream or upwind), but treatment assignment is credibly exogenous to other unobservable shocks from mine openings occurring through market-based channels. The high temporal resolution of the satellite data allows me to estimate both event study and distributed lag versions of my DID designs, so I can distinguish between contemporaneous and prolonged effects of exposure to pollution on local yields.

Next, to estimate an “overall” effect of industrial mining on local agricultural output, I use the staggered openings of over 300 large-scale mines in a DID that compares NDVI in areas near mines to areas slightly further away, before and after the mine’s opening. This “overall” effect captures both market-based channels and pollution externalities, as both will disproportionately affect areas near mines relative to areas further

away.

I find that on average, mine openings lead to a statistically significant 2-3% fall in NDVI for areas within 20 kilometers of the mine, relative to areas further away. Isolating only the water pollution channel, I find that downstream areas experience a 3-4% drop in NDVI relative to upstream areas after a mine opening. In contrast, I find negligible and statistically insignificant effects of both contemporaneous and cumulative mining-induced air pollution on yields.

To measure the share of the overall effect of mining on agricultural output that can be explained by pollution externalities, I scale the treatment effects estimated in the air and water pollution analyses by the share of the area in 20km “near” buffer that would be exposed to these types of pollution. I then take the ratio of the scaled pollution treatment effects to the overall effect of mining on NDVI, finding that pollution externalities account for 44% of the overall effect.

Furthermore, I conduct a quantification exercise to translate the estimated NDVI effects to predicted effects on actual plot-level yields in kilograms per hectare. I regress plot-level yields from a survey of over 13,000 smallholder farmers on measures of NDVI around these plots, finding that on average a 0.01 unit decrease in NDVI corresponds to a 55 kilogram per hectare decline in yields. Back-of-the-envelope calculations that combine this correlation between NDVI and actual yields with the estimated overall treatment effects show that on average, mine openings reduce local yields by 36 kg/ha through both pollution externalities and market-based channels.

Lastly, I generate mine-specific DID estimates to investigate heterogeneity in treatment effects across a variety of dimensions including mine characteristics, local governance or regulatory environments, economic factors and pre-period exposure to pollution. I find that both governance and mine type can influence the extent to which mines affect local agriculture via pollution. Both the air and water pollution results reveal that treatment effects are larger around mines in countries that perform worse on governance indicators, such as regulatory quality and government effectiveness. In addition, pollution effects are larger for areas that experienced low levels of pollution prior to the mine opening, suggesting that either crops or farmers in more polluted areas may be better adapted to poor environmental conditions.

My paper contributes to three growing literatures. First, it adds to the literature studying the local impacts of natural resource extraction by examining the understudied question of how the pollution from extractive industries affects local agricultural output. The discussion of whether natural resources are a blessing or a curse has focused heavily on the effects of natural resource extraction on local economic outcomes, such as re-allocation of labor [Kotsadam and Tolonen, 2016], wealth [von der Goltz and Barnwal, 2019] and income [Aragón and Rud, 2013], as well issues of corruption and governance (Caselli and Michaels [2013]; Asher and Novosad [2023]), conflict [Berman et al., 2017] and crime [Axbard et al., 2021]. Our

understanding of the effects of natural resource extraction on outcomes through the channel of pollution externalities in a developing country context is primarily limited to human health outcomes (von der Goltz and Barnwal [2019], Benschaul-Tolonen [2020]) and human capital development (Bonilla Mejía [2020]; Rau et al. [2015]). With the exception of [Aragón and Rud, 2016], to the best of my knowledge there is no other work quantifying the effects of natural resource extraction on agricultural outcomes in developing countries. My contribution is to isolate the effects of pollution from natural resource extraction on agricultural output and examine heterogeneity in these effects. This is valuable for informing policy on natural resource management in developing countries, where local communities may need to make a trade off between local economic boosts from extractive industries and long run environmental degradation in the absence of effective environmental governance. It is additionally policy-relevant to understand what factors may mitigate or aggravate the extent to which mining pollution affects local agricultural output.

Second, my paper contributes to the literature on structural transformation and agriculture. By isolating the effect of pollution externalities from market-based channels, I am able to speak to the extent to which local reductions in yields are driven by pollution rather than structural transformation through the reallocation of labor away from agriculture to other sectors or regions of the country. Huang et al. [2023] and Kotsadam and Tolonen [2016] both find evidence that mining leads to structural shifts out of agriculture into more productive sectors, like low-skilled services. As I find that pollution externalities alone account for the bulk of the effect of mine openings on local yields, this suggests that we must be cautious in interpreting mining-driven structural transformation as a net positive given the costs of environmental degradation, which will persist even after mine closures reduce or even eliminate the local economic opportunities afforded by mining companies.

Finally, my paper contributes to an emerging literature that uses remotely-sensed outcomes to address data gaps in developing countries. Common applications of satellite-imagery to questions of causal inference include NDVI to proxy for crop yields (Sukhtankar, 2016; Emerick, 2018) and AOD to proxy for air pollution (Gendron-Carrier et al. [2022]; Xie and Yuan, 2023; Zhou, 2021). To the best of my knowledge, I am the first to apply a satellite-based measure of water quality used in the remote sensing literature, the normalized difference turbidity index, to estimate causal effects on water pollution. This innovation is valuable as it can allow researchers to answer questions in developing country settings where ground-measured indicators of water quality are spatially and temporally sparse, or completely unavailable.

2 Background

Mining pollution can affect crop yields through shocks to land and labor inputs. Shocks to land may arise through channels such as worsening soil quality, direct toxicity of certain pollutants to plant health and interference with plant photosynthetic processes. Labor shocks primarily arise through pollution affecting human health and resulting labor productivity, which has been documented extensively in the air pollution context (Graff-Zivin and Neidell, 2012, He et al. 2019) but sparsely in the water pollution context (Russ et al., 2022). The following section provides background information on how mining pollution can affect yields through the plant health and soil quality channels, which are less clear than the labor productivity channel.

2.1 Mining, Air Pollution and Crop Yields: Several small-scale environmental impact assessments and qualitative research suggest that mining can generate air pollution (Mwaanga et al. [2019]; Ghose and Majee [2002]). These studies typically focus on a handful of sites within a single country and time period, using ground-based measures of air pollution. A separate literature examines the extent to which air pollution can affect crop yields [Lobell et al., 2022].

The technologies employed by large-scale mining operations can contribute to air pollution through a variety of ways. First, metal smelting and refining produces gaseous emissions, such as carbon dioxide, sulfur dioxide, and nitrogen oxide, as well as particulate matter [Dudka and Adriano, 1997]. Sulfur dioxide emissions are especially common, as most economically important mineral ores occur as metallic sulfides, which are released into the atmosphere during blasting and smelting operations that expose rocks to the air and other chemicals. Furthermore, smelting and refining operations are major contributors to emissions of heavy metals, such as arsenic, cadmium, copper, lead and zinc. Since most mining operations in developing countries tend to process the raw ores on site before export, the pollution effects of these smelting and refining activities are concentrated in areas local to the mine. Second, activities such as blasting, crushing, stockpiling, loading and transportation of mine rocks also release large amounts of dust [Petavratzi et al., 2005]. This dust contributes to fine particulate matter and heavy metal levels in the atmosphere. In fact, levels of air pollution generated by mining operations can be on par with those produced by power plants or other large-scale industrial activities [Dudka and Adriano, 1997]. Finally, there are two main types of mining: underground mining and surface (also known as open pit) mining. While both types can contribute to air pollution, open pit mining is considered more damaging [Sahu et al., 2015].

Firstly, mining-induced air pollution can directly affect plant health. Activities such as drilling, blasting or smelting can release pollutants into the atmosphere that affect plant photosynthesis or other plant growth

processes. Whether air pollution has a positive or negative effect on plant health depends on the type and intensity of pollution. Some pollutants, like ozone, consistently reduce yields by limiting photosynthesis [Feng and Kobayashi, 2009]. Other pollutants, like particulate matter (PM), have mixed effects due to competing absorption and scattering effects. While absorption by PM reduces total solar radiation reaching the surface and thus diminishes a plant’s photosynthetic capabilities, scattering by PM increases the fraction of diffuse light available, which can be more effectively used by the plant [Schiferl et al., 2018]. This suggests that the effect of mining-induced air pollution on local crop health could be highly heterogeneous across mines, agro-ecological zones and climates.

Secondly, air pollution can lead to long-run deterioration of soil quality through atmospheric deposition of pollutants on the soil over time. As dust, heavy metals and other pollutants are released into the atmosphere from mining activities, they are carried by the wind and deposited on the ground. While some literature suggests that atmospheric deposition could erode soil quality over time [Liu et al., 2023], other literature finds that atmospheric deposition from polluting industries can actually serve as key nutrient inputs for modern high yield crops [Sanders and Barreca, 2022].

Lastly, I seek to disentangle contemporaneous and cumulative effects of air pollution on yields. Contemporaneous effects refer to when air pollution affects concurrent yields, that is, air pollution occurring within a given season affects yields in that same season. In contrast, cumulative effects refer to when yields are affected by prolonged exposure to air pollution. While yield effects due to deterioration of soil quality will likely occur due to cumulative exposure, yield effects due to plant health or photosynthetic interference could also occur solely from contemporaneous pollution.

2.2 Mining, Water Pollution and Crop Yields: Mines can pollute surface water and ground-water sources through acid mine drainage (AMD), when water high in sulfates and heavy metals from mining activities is released into the soil or neighboring water bodies [Dudka and Adriano, 1997]. Water contaminated with AMD can then affect crops either by overflow or seepage from contaminated rivers, seepage through groundwater or through irrigation. In some cases, AMD can reduce yields substantially: pot experiments reveal that fields contaminated with AMD have 62% lower grain yields than non-contaminated fields [Choudhury et al., 2017].

Like air pollution, water pollution from mines can have both contemporaneous and cumulative effects on yields. However, unlike air pollution, there is much less temporal variation in which areas are exposed to AMD. While wind direction can vary substantially within a year, blowing air pollutants in different directions around a mine, the direction of water flow is generally consistent over time.

3 Theoretical Framework

To provide a framework for interpreting the extent to which mining affects local agricultural output through pollution externalities versus market-based channels, I assume that farmer i in mining area m produces a single agricultural good with price $p = 1$ in season t , according to the following production function:

$$Y_{imt} = A_{imt}F(L_{imt}, M_{imt}, \epsilon_{imt}) \quad (1)$$

where Y is actual output, A is total factor productivity, M is land, L is labor, ϵ represents unobservable farmer- and season-specific shocks to output, and F is a concave production technology.

The farmer then maximizes profits within “perfect” input markets by choosing land and labor inputs accordingly:

$$\max_{L,M} A_{imt}F(L_{imt}, M_{imt}, \epsilon_{imt}) - wL_{imt} - rM_{imt} \quad (2)$$

where w is the wage rate and r is the cost of land. As in Aragón and Rud [2016],

I assume that A is a function of three factors: farmer-specific heterogeneity (η_i), time-invariant economic and environmental conditions specific to a given mining area (ρ_m), and time-varying factors correlated with the presence of local mining activity (S_{mt}). I further break down S_{mt} into environmental, E_{mt} , and non-environmental factors, N_{mt} , correlated with the presence of mining activity. Importantly, E_{mt} is directional: environmental impacts generated by local mining activity will disproportionately affect areas near mines that are more likely to be exposed to pollution, due to local geography and weather. For example, water pollution will disproportionately affect areas downstream from a mine, where pollutants from the mine will flow. In contrast, N_{mt} is generally non-directional, in that most non-environmental factors correlated with mining activity will tend to affect all neighboring areas around a mine.

Mining can affect agricultural output through both shocks to inputs and total factor productivity. Firstly, mining can increase demand for inputs, which increases input prices w and r . This may reduce the amount of land and labor used in agriculture, as these inputs are allocated towards other sectors that support mining activity. Demand for labor may be driven by the introduction of low-skilled mining jobs or the development of small businesses to support large-scale mining activity (referred to as backwards linkages by Aragón and Rud [2013]). Similarly, demand for land may increase as mining companies buy or appropriate land from existing land owners.

Importantly, shocks to input demand are non-directional, that is, all areas near a mine will likely be affected by these wage or land price increases, not only the areas that are more likely to experience pollu-

tion. Indeed, existing work has shown that mining increases wealth [von der Goltz and Barnwal, 2019] and employment in the non-agricultural sector [Kotsadam and Tolonen, 2016] in areas near mines, relative to areas further away.

Secondly, mining can affect how productively inputs are used, through shocks to S_{mt} . To begin, mining can generate air and water pollution, which affects the quality of land and labor inputs (as described in Section 2). This enters the model through a negative total factor productivity shock via E_{mt} , which will disproportionately affect areas near a mine that are exposed to pollution, rather than all surrounding areas.

Aside from pollution, mining can affect total factor productivity through non-environmental factors such as changing composition of agricultural workers or changing agricultural practices, which are shocks to N_{mt} . Even if the quantity of people employed in agriculture remains the same, the sorting of more productive farmers out of agriculture into other sectors could reduce A . Mining may also change agricultural practices, even if input choice remains the same. One channel through which this effect might arise is through weak property rights. If farmers fear that their land may be appropriated for mining, they may switch to less productive agricultural practices (e.g. not growing cash crops), even if they use the same amount of land and labor on their farms.

My paper seeks to isolate the impact of mining-induced pollution (shocks to E_{mt}) on agricultural output, from the effect of mining on agricultural output through market-based channels. To do this, I leverage the fact that pollution shocks are directional while shocks through market-based channels, such as input price shocks or shocks to non-environmental time-varying factors correlated with mining activity, are non-directional. The literature examining the effects of mining on various economic outcomes predominantly uses a DID design that compares areas near mines to areas slightly further away (Aragón and Rud [2013], Aragón and Rud [2016], Kotsadam and Tolonen [2016], Benshaul-Tolonen [2020], Dietler et al. [2021], von der Goltz and Barnwal [2019]). Importantly, the near versus far DID mixes together the effects of the pollution externalities and market-based channels, both of which would more heavily affect areas closer to mines. This suggests that using a near vs. far DID is insufficient to isolate the pollution externalities of mining on agriculture.

Aragón and Rud [2016] serves as a helpful benchmark for the effects of mining activity on agricultural productivity. The authors estimate an agricultural production function using three rounds of repeated cross-section household data and a near versus far DID design in the context of 12 gold mines in Ghana. They use cumulative gold production as their proxy for pollution exposure, comparing the effect of increased gold production on household agricultural output in mining areas to that of non-mining areas further away. A household falls in a mining-area if it is located within 20 kilometers of a mine. All other households are considered to be living in non-mining areas. The authors find that a 100-tonne increase in cumulative gold

production reduces real agricultural output by about 17% in mining areas, relative to non-mining areas. Of the 12 mines in their sample, 9 were already operating during the first wave of the survey. As a result, their empirical strategy estimates primarily an intensive margin effect of increased mining production on agricultural output.

By controlling for household labor and land use as covariates in their model, Aragón and Rud [2016] are able to remove the influence of market-based input demand shocks when trying to isolate the effect of mining-pollution on agricultural productivity. However, remaining market-based channels that operate by shocking total factor productivity and are correlated with proximity to a mine, such as local demand shocks that change the composition of workers in agriculture, remain a potential endogeneity issue when trying to isolate the pollution effect. This highlights the importance of more carefully defined treatment and control areas.

Furthermore, while the focus on a single commodity and single country over a short time period allows Aragón and Rud [2016] to leverage context-specific information to close down several market-based channels, it limits the ability to study heterogeneous effects of mines on agricultural output. It is policy-relevant to understand whether their finding that mining leads to meaningful reductions in agricultural output is consistent across other countries, commodities, regulatory environments, technologies and time.

4 Data

To investigate the effect of industrial mining on agricultural output in mining communities, I construct a panel of mining areas from 2000-2022. I link spatial buffers around mine geo-locations and mine opening dates, provided by S&P Global Market Intelligence, to remotely sensed measures of pollution, crop yields and weather variables. The unit of observation in my analysis is a mine-side-month, where a “side” refers to part of a buffer around a mine. In this section, I briefly explain the construction of the main variables used in my analysis, with an in-depth discussion detailed in Appendix 10.1

4.1 Mines S&P Global Market Intelligence provides the geo-locations of the centroids of mining operations for large-scale mines across the world. These centroids are used to construct three different types of spatial buffers, which define treatment and control areas for my analysis.

First, to estimate an “overall” effect of mining activity on yields, I define circular buffers around each mine to distinguish between areas near mines and areas slightly further away. Intuitively, areas near mines are more likely to experience both the effects of local labor market shocks and pollution exposure from mining activity, relative to areas further away. I refer to the buffer of radius 20 kilometers around a mine as

the "near" group, which I assume is the "treated" area that would be impacted by both market-based and pollution channels. I discuss the specific reasons for choosing the 20-km cutoff in more detail in Section 7.

Second, I buffer river segments near mines to estimate effects of mining-induced water pollution on yields. I use the HydroRIVERS dataset to identify rivers located in Africa. Each mine within 1 km of a river in this network is "snapped" to the closest river segment, then the upstream and downstream river segments from the mine river segment are identified. The upstream and downstream segments are buffered by 1-kilometer on either side of the line to identify land area around the rivers that would likely be affected by water pollution effects occurring through seepage, irrigation or labor productivity channels. The "side" of a mine in the water pollution analysis refers to the upstream or downstream buffer.

Lastly, to estimate effects of mining-induced air pollution on yields I use slices of circular buffers around each mine that indicate wind direction from the mine. I start by defining a buffer of radius 60 kilometers around each mine. This radius is chosen based on a spatial decay model shown in Appendix Figure 10, which demonstrates that on average, air pollution, as proxied by AOD, is concentrated within 60 km of the mine centroid and is statistically indistinguishable from zero beyond this cut-off. Next, for each mine the 60-km buffer is divided into 4 slices representing the cardinal directions that the wind can blow from the mine centroid. Hence, a "side" of a mine in the air pollution analysis refers to one of these 4 slices of the 60-km buffer around the mine. Note that the circle of radius 5-km around the mine is removed from the 60-km circle, leaving behind a "donut," to limit the influence highly local air pollution effects that would affect all sides of the mine regardless of wind direction.

Figure 1 illustrates the three different types of treatment and control areas constructed for the three types of analysis. The left panel illustrates the near versus far buffers used to estimate an "overall" effect of mining on local yields, where areas within the 20-km buffer are considered "treated." The top right panel illustrates the upstream and downstream buffers used to estimate the water pollution effect, while the bottom right panel illustrates the wind direction buffers used to estimate the air pollution effect.

Additionally, the S&P data provides information on the year that commercial production at the mine first began, the specific minerals extracted and the extraction method. In line with S&P, I define the start of mining activity, also referred to as the opening of a mine, as the year that "the mine/plant has been commissioned or has produced its first metal, concentrates, or bulk commodity at a commercial rate." Other papers have defined mining activity based on years where a mine had non-zero output (von der Goltz and Barnwal [2019], Benschaul-Tolonen [2020]). Although the S&P data provides yearly production data for some mines, this information is missing or incomplete for many mines in Africa, meaning activity would need to be imputed. As the S&P data on the start year of commercial production is more complete, I opt to define mining activity based on this year instead.

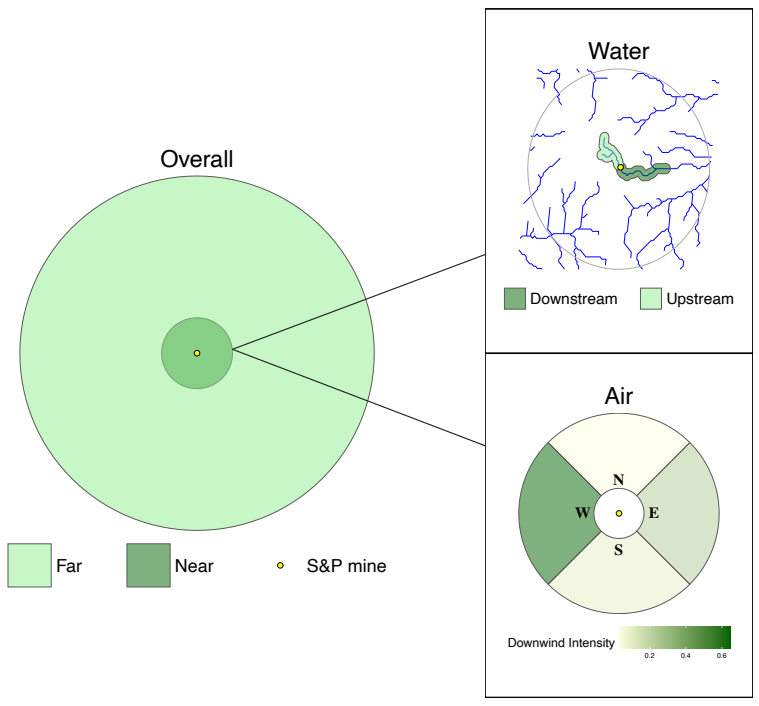


Figure 1: Treatment and Control Buffers

Given that the S&P mining data covers mostly large-scale mines operated by multinationals or national governments, this might suggest that my analysis is not able to speak to the effects of artisanal and small-scale mines (ASM) or those that are illegally operated. As much of ASM and illegal mining operates on the fringes of large-scale mining operations, the mine buffers should encompass areas that are affected by ASM and illegal mining. However, if small-scale mining activities were already contributing to pollution prior to large-scale operations, treatment effects may be underestimated. Since ASM and illegal mining generally uses low capital intensive and high labor intensive technologies, rather than large blasting, drilling and smelting machinery, I do not expect an attenuation of treatment effects in the air pollution analysis. However, underestimation of treatment effects is a potential concern with the water pollution analysis, as ASM or illegal mining often releases heavy metals into the soil and water bodies through the improper and unregulated use of chemicals during the extraction process. For example, artisanal gold mining is associated with mercury pollution as the mercury used in the traditional amalgamation process to separate gold from the ore can leach into the surrounding soil and water [Soe et al., 2022].

4.2 Wind: Wind data for the air pollution analysis is retrieved from the Modern-Era Retrospective analysis for Research and Applications, Version 2 (MERRA2). MERRA2 is a re-analysis product, meaning it combines satellite imagery with algorithms and atmospheric models to create the final data product. I use the daily aggregates of hourly time-averaged U and V wind vector components for heights of 2 meters and 50 meters, at a 0.5 degree x 0.625 degree resolution.

To measure downwind exposure for each side of a mine, I first determine the direction (in degrees from due North) that the wind is blowing from the mine centroid, for each day from 2000-2022.¹, and classify this direction into one of the four 90-degree sides defined by the cardinal directions (North, East, South, West). Then, for each mine-month, I count the number of days that the wind blows from the mine into each of the four sides, as well as determine the average monthly wind speed experienced by each side. Downwind exposure experienced by a given side is defined as the share of days in a month that the side is downwind from the mine.

4.3 Satellite-based air pollution: Daily measures of aerosol optical density (AOD) are obtained from the Moderate Resolution Imaging Spectroradiometers (MODIS) satellites at a 3-kilometer spatial resolution. In brief, AOD is calculated by comparing the light intensity in a particular band against a reference value and attributing the difference to particulates in the air column. The higher the level of AOD, the higher the level of air pollution as more light is being reflected back due to particulates in the atmosphere.

¹ $Direction = 90 - atan2(v, u) \times 180/\pi$

Several studies have shown that AOD is highly predictive of ground-based measures of PM10 and PM2.5. In their preferred specification, Gendron-Carrier et al. [2022] find that one unit of remotely sensed AOD is associated with about $114 \mu\text{g}/\text{m}^3$ of PM10 measured by a ground-based instrument.

I rely on global daily rasters of MODIS AOD generated by Gendron-Carrier et al. [2022]. The authors provide daily tifs from July 4, 2002 until August 31, 2018, for the Aqua satellite and February 24, 2000 until July 31, 2020 for the Terra satellite. Given MODIS data availability, the panel used in the air pollution analysis covers 2003 - 2017. From the daily AOD rasters, I construct mean AOD for each mine-side-month by averaging all pixel-days with non-missing AOD readings that fall within each side of a mine in a given month. I calculate this average monthly AOD only using data from the Aqua satellite, as the Terra satellite suffers from more missing values due to satellite detection errors during my time period of interest. More details on AOD data construction are covered in Appendix 10.1.

4.4 Satellite-based water pollution While satellite-based measures of air pollution have been used extensively to answer causal inference questions related to air pollution ([Gendron-Carrier et al., 2022], Guitierrez and Teshima; Chen, Oliva and Zhang), to the best of my knowledge, no paper is yet to conduct causal inference using similar satellite-based measures of water pollution in developing countries.

Satellite-based proxies for pollution have an key advantage in that they allow researchers to answer questions in contexts where ground-based pollution data is unavailable or limited. Importantly, spatial and temporal gaps in ground-based water pollution data can be even larger than those for air pollution in developing countries (Holger et al., 2022). While low cost air quality sensors have been deployed in many African countries, the development and implementation of low cost water quality sensors is in its infancy. These limitations in ground-based water quality monitoring highlight the value of remotely-sensed measures of water pollution, which offer an affordable way to track quality of African water bodies over time.

While the remote sensing literature has developed several promising satellite-based indicators of water quality (CITE), these measures are yet to make their way into the social sciences literature. To the best of my knowledge, I am the first to apply a satellite-based measure of water quality from the remote sensing literature, the normalized difference turbidity index (NDTI), in the context of causal inference. NDTI uses the spectral reflectance of water pixels to estimate turbidity, which is a measure of water clarity. Similar to AOD, turbidity measures the amount of light scattered by particles in the water column. The higher the level of particles in the water, the more light that will be scattered. High turbidity makes water appear cloudy, muddy or discolored. Turbidity is often used to proxy for total suspended solids (TSS), which are a common indicator of water quality. TSS are particles found in the water column that are larger than 2 microns. Increases in TSS can be driven by natural phenomena, such as stirred bottom sediments from heavy

precipitation or algal blooms, as well as industrial runoff or wastewater discharge. When TSS levels cannot be measured, changes in turbidity can be used to reflect changes in TSS concentration in water (CITE).

To measure turbidity, NDTI leverages how electromagnetic reflectance is higher in green spectrum than the red spectrum for clear water (Lacaux et al., (2007), Gardelle et al. (2010)). Increased reflectance of the red spectrum relative to green corresponds to an increase in turbidity (Islam and Sado, 2006). NDTI is calculated with the red and green spectral bands according to the following formula:

$$NDTI = \frac{\text{Red} - \text{Green}}{\text{Red} + \text{Green}}$$

Generally, NDTI ranges from -0.2 to greater than +0.25, where lower (negative) values indicate clear water and higher (positive) values indicate turbid water (Alka et al. 2014).

To obtain a remotely-sensed measure of river water pollution, I first identify the precise shape of the upstream and downstream river segments near mines using a water mask based on the normalized difference water index (NDWI), which is commonly used to detect water bodies in satellite images and is calculated from the green and near-infrared (NIR) bands.² I generate the NDWI water mask from a mosaic of cloud-free Sentinel-2 pixels from 2020-2023 at a 15-m resolution, to limit the influence of atmospheric interference on water detection. The NDWI water mask is fixed over time, meaning any estimated effect of mining activity on turbidity does not reflect changes in the extent of water bodies over time.

Next, I calculate monthly NDTI for each identified river pixel from 2000-2022, masking out low quality or cloudy pixels, using Landsat 7 data at a 30 meter resolution. Finally, I average monthly NDTI across water pixels within the upstream or downstream river segments, for each mine, to create mine-month-side level measures of river NDTI.

With NDTI, I am able to overcome the poor spatial and temporal resolution of ground-based water pollution monitoring in Sub-Saharan Africa. Global freshwater quality databases, such as the Global River Quality Archive (GRWQ), which harmonizes five water quality datasets to provide the most comprehensive river water quality data from ground-based sampling sites, have extremely sparse coverage of Sub-Saharan Africa. This makes it challenging to use these datasets to investigate the effects of highly localized polluting activities on water quality over time. In Appendix Table 6, I show that NDTI is highly predictive of ground-measured TSS from the GRWQ database. In my preferred specification, a one unit increase in NDTI is associated with a 887 mg/l increase in ground-measured TSS.

Importantly, while NDTI is successful at detecting changes in the overall levels of particles in water,

²

$$NDWI = \frac{\text{Green} - \text{NIR}}{\text{Green} + \text{NIR}}$$

it cannot distinguish between types of particles. While some particles may be pollutants from mine waste effluent, such as heavy metals or acidic sulfides, others may be sediment from rock or soil disturbed during the extraction process. While heavy metals and sulfides have been shown to substantially reduce crop yields (CITE), the effect of excessive sediment in water on local yields is less clear, though some literature suggests potential negative effects. Gurmu et al. find that excess sediment in the water can disrupt formal irrigation systems. However, many smallholder farmers either do not use irrigation or use an informal system (e.g. fetching water from river themselves), suggesting that this channel is unlikely to impact local yields around a river. Alternatively, in extreme cases excess sediment in water may lead to harmful algal blooms, which can interfere with the development of plant structures during irrigation or seepage into neighboring lands (Newton and Melaram). Since NDTI cannot distinguish between polluting particles and sediment, it is important to interpret observed changes in NDTI due to mining as changes in both harmful pollutants and potentially non-harmful sediment.

4.5 Satellite-based yields: Ideally, I would use data on farm-level yields at a high spatial and temporal resolution around mines. While some high quality plot-level agricultural survey data is available, such as the World Bank Living Standards Measurement Survey - Integrated Surveys on Agriculture (LSMS-ISA), these sources are not suited for my research design for several reasons. To start off, these surveys are limited to a handful of countries and do not cover many households near mines, let alone within the precisely defined pollution buffers. Furthermore, surveys are collected infrequently, making it difficult to tie observed effects on yields to the start of mining activity or examine cumulative effects of pollution over time.

To overcome survey data limitations, I proxy for crop yields using the normalized difference vegetation index (NDVI). NDVI is strongly correlated with crop productivity and final yields, especially during the growing season [Panek and Gozdowski, 2021]. I construct a measure of pixel-level daily mean NDVI from 2000 - 2022 using the MODIS MCD43A4 Version 6.1, combined Terra and Aqua product at a 500-m resolution. NDVI is calculated from the MODIS data using readings of light reflected in the near-infrared and red spectrum.³

When aggregating NDVI from the pixel-day to the mine-side-month level, I first mask out pixels that are not classified as cropland according to the Global Food Security Support Analysis Data (GFSAD), which provides data on cropland extent at a 1-km resolution. The coarse resolution of the cropland extent data may constrain the ability to detect cropland of very small farms, resulting in potential measurement error in NDVI over cropland. However, I show in Table 5 that NDVI is still highly predictive of plot-level yields in Sub-Saharan Africa, though the magnitude of the correlation may be attenuated. I construct my final

³NDVI = $\frac{NIR-RED}{NIR+RED}$

measures of cropland NDVI for the treatment and control areas by averaging NDVI across all cropland pixel-days within each mine-side-month, within the relevant buffers.

An emerging literature uses machine learning algorithms to predict crop yields based on vegetation indices, like NDVI, and other covariates (Pelletier et al., Burke and Lobell, Jin et al.). Unfortunately, most of the mining areas in my analysis are not located near sites where ground truth measures of yields, like cropcuts, are available. Furthermore, high quality ground truth data are not available prior to 2015. Given that machine learning models predict more poorly on observations that are spatially distant from training data samples (Proctor et al., 2023) or outside of the time period covered by the training data (CITE), I am concerned that using predicted yields as my main outcome of interest would introduce additional measurement error into the estimation of the effect of mining on yields. As a result, I opt to use NDVI as my main measure of agricultural output rather than predicted yields. This allows me to remove prediction error as a source of measurement error, and instead focus solely on understanding how measurement error in the NDVI data itself could be affecting estimated coefficients. Additional discussion of measurement error in NDVI are covered in Appendix 10.14.

4.6 Weather Controls: Local weather conditions influence actual yields, as well as the ability for satellites to accurately detect these yields (Gendron-Carrier et al. [2022]; Heino et al. [2023]). I calculate mine-side-month averages of cloud cover percentage, temperature, precipitation, evapotranspiration, vapor pressure and wet day frequency using the Climatic Research Unit gridded dataset, available at a 0.5 degree resolution from Harris et al. [2014]. These monthly averages of key weather variables are the main control variables in my regressions.

4.7 Agricultural Seasons: For each mine, I aggregate daily measures of AOD and NDVI to the monthly level. Each mine-month is then linked to one of the following agricultural seasons: planting, early growing, late growing, harvest and non-farm. Aggregating to a less granular temporal resolution allows me to smooth over missing days of data that occur due to cloud cover or other satellite errors. I use the Sacks et al. [2010] global raster dataset on planting and harvesting dates for maize to define the agricultural seasons. This data is available at a 5 arc-minute (approximately 10 km x 10 km) spatial resolution. I focus on maize as it is one of the primary crops grown in Africa and has the most complete crop calendar data. Other important crops like cassava or wheat have missing data across most of Africa and Sacks et al. [2010] caution against using interpolated products. Using seasons defined by maize only could potentially introduce measurement error into regressions where season indicators are used as covariates or to stratify the sample. I discuss this issue in more detail in Appendix 10.14.

5 Mining-induced Water Pollution

To estimate the effect of mining-induced water pollution on yields, I use a difference-in-difference (DID) that compares NDVI between the upstream and downstream sides of a mine, before and after a mine opening:

$$NDVI_{smt} = \alpha_{sm} + \lambda_{mt} + \delta Downstream_{sm} + \eta Post_{mt} + \beta_W Downstream_{sm} \times Post_{mt} + \mathbf{X}'_{smt} \boldsymbol{\Gamma} + \epsilon_{smt} \quad (3)$$

where $NDVI_{smt}$ is mean NDVI on side s of mine m in month t . $Downstream_{sm}$ is a dummy variable equal to 1 if side s is downstream from mine m and equal to 0 if it is upstream. $Post_{mt}$ is a dummy variable equal to 1 after mine m opened, 0 otherwise.

I include mine-side fixed effects (α) to control for time invariant unobservables correlated with agricultural output on each side of a mine, such as soil quality, as well as mine-date fixed effects (λ) control for mine-specific trends in NDVI over time. Given the mine-side and mine-date fixed effects, β_W is identified by within-month differences in the change in NDVI among areas upstream and downstream mine openings. Essentially, β_W is a weighted average of mine-specific DIDs generated from almost 40 mine openings used in this analysis. To address the influence of weather specific shocks to yields that vary across sides and time, I control for linear and quadratic terms of mean cloud cover, vapor pressure, temperature, precipitation, evapotranspiration and wet days.

Estimating the causal effect of mining-induced water pollution on yields requires that the timing and placement of a mine opening is plausibly exogenous to local changes in NDVI. While the placement of mineral deposits is random, the discovery of these deposits may depend on a variety of non-random factors, including local governance quality, business environments, infrastructure access and input prices [Benshaul-Tolonen, 2020]. This issue should not pose a threat to identification in the upstream versus downstream DID, as both sides of the river will be similarly affected by these factors within the narrowly defined buffers.

I estimate Equation 3 over a partially balanced panel of 38 mines within 1 km of a river, where non-missing NDVI is observed on both the upstream and downstream sides of the mine, for at least 6 months of every year, for at least 2 years pre- and 2 years post-mine opening. Additionally, I estimate an event study version that replaces the single $Post$ dummy with a series of event time dummies indicating the number of years since a mine opened. As NDVI during the growing season is more strongly correlated with actual crop yields than NDVI during the non-growing season, I also estimate versions of the DID and event study models separately over months in the growing season and non-growing seasons. Standard errors are clustered at the mine-level.

To establish the first-stage effect of mining activity on water pollution, I estimate Equation 3 with normalized difference turbidity index (NDTI) as the outcome of interest. Monthly NDTI measures are missing more frequently, since NDTI is calculated over narrowly defined river pixels and so are more affected by cloud cover or other atmospheric interference. As a result, I estimate the event study and DID regressions for turbidity over a subset of the mines used in the NDVI analysis.

Figure 2 shows that on average after a mine opens, downstream river turbidity increases significantly. Larger increases in turbidity during the growing seasons could be explained by increased precipitation causing more runoff from mining areas into rivers. Figure 3 reveals a complementary pattern of reductions in NDVI for downstream cropland after a mine opening, with stronger negative effects during the growing seasons.




Figure 2: Average effect of mining-induced water pollution on NDTI - event study

The figure plots coefficient estimates from two separate event study regressions of mean normalized difference turbidity index on a dummy for whether a side is downstream from a mine interacted with a series of event time dummies for years since a mine first opened. The unit of analysis is a mine-side-month, though event time coefficients are binned into yearly increments. The event study regressions are estimated separately over growing season months (early growing and late growing) and non-growing season months (planting, harvest and non-farm). The dotted lines show point-wise 95 percent confidence intervals for the coefficients of the event-time path of mean NDTI. Coefficients can be interpreted as estimated effects relative to the period one year before the mine opened. Both regressions includes linear and quadratic controls for mean temperature, precipitation, vapor pressure, wet days, evapotranspiration and cloud cover, as well as mine-side and mine-year-month fixed effects. The sample includes the 30 mines for which non-missing NDTI is observed on the upstream and downstream sides for at least 2 months in each year, for at least 2 years pre- and 2 years post-mine opening. Event times less than -2 or greater than 2 are binned into two end points.

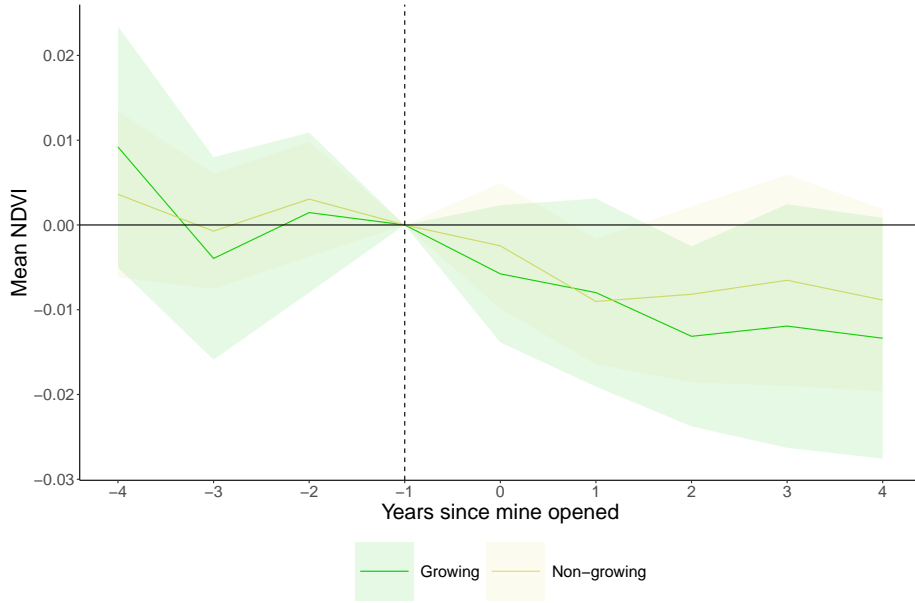


Figure 3: Average effect of mining-induced water pollution on NDVI - event study

The figure plots coefficient estimates from two separate event study regressions of mean NDVI on a dummy for whether a side is downstream from a mine interacted with a series of event time dummies for years since a mine first opened. The unit of analysis is a mine-side-month, though event time coefficients are binned into yearly increments. The event study regressions are estimated separately over growing season months (early growing and late growing) and non-growing season months (planting, harvest and non-farm). The dotted lines show pointwise 95 percent confidence intervals for the coefficients of the event-time path of mean NDVI. Coefficients can be interpreted as estimated effects relative to the period one year before the mine opened. Both regressions includes linear and quadratic controls for mean temperature, precipitation, vapor pressure, wet days, evapotranspiration and cloud cover, as well as mine-side and mine-year-month fixed effects. The sample includes the 38 mines for which non-missing NDVI is observed on the upstream and downstream sides for at least 6 months in each year, for at least 3 years pre- and 3 years post-mine opening. Event times less than -3 or greater than 3 are binned into two end points.

To better quantify the impact of industrial mining, I use the DID regression model outlined in Equation 3. As NDTI is imperfectly correlated with polluting particles in water that may affect yields, I opt to estimate the first-stage and reduced form results separately, rather than as a two-stage least squares model with the downwind and post interaction instrumenting for NDTI. Table 1 presents main results for remotely sensed turbidity and yields, estimated on all months pooled together, as well as separately by growing and non-growing seasons.

Panel A consistently shows that mining increases remotely sensed water turbidity, though these estimates are statistically insignificant, likely due to lack of power. The increase in turbidity is on the magnitude of 12-50% relative to NDTI in the pre-period. Rather than give substantial weight to the magnitude of these effects, I cautiously interpret the reduced form results as evidence of an extensive margin effect of mining on water pollution. Furthermore, the observed increase in NDTI is capturing both increased sediment and increased pollution, meaning that only a fraction of the large effect size is likely attributable to increased water pollution. Using Column 2 of Appendix Table 6 to convert NDTI to total suspended solids suggests that on average, mine openings increase total suspended solids in water downstream by 4.33 - 4.83 mg/l.

Similarly, the reduced form results in Panel B show consistent and statistically significant reductions in remotely sensed yields in downstream areas relative to upstream areas, after a mine opens. These estimates correspond to a 3-4% drop in NDVI, with larger effects observed during months in the growing season. The observed reductions of NDVI in downstream cropland suggest that at least part of the observed increase in remotely-sensed turbidity can be attributed to increased pollution, given that increases in sediment in water are unlikely to affect local yields.

Table 1: Relationship Between Industrial Mine Openings, Remotely-Sensed Water Turbidity and Crop Yields

	(1)	(2)	(3)
	All seasons	Growing	Non-growing
Panel A: First Stage for Turbidity			
Downstream \times Post	0.00488 (0.00649)	0.00426 (0.00761)	0.00544 (0.00636)
Mines	30	30	30
Observations	11,330	5,304	6,026
Mean NDTI (t-1)	.01	-.01	.04
Panel B: Reduced Form for Yields			
Downstream \times Post	-0.01404** (0.00587)	-0.01770** (0.00714)	-0.01093** (0.00531)
Mines	38	38	38
Observations	20,518	9,888	10,630
Mean NDVI (t-1)	.47	.54	.39
Mine-side FE	Yes	Yes	Yes
Mine-year-month FE	Yes	Yes	Yes
Weather	Yes	Yes	Yes

Notes: Each column reports the results of a linear regression. The unit of analysis is a mine-side-month. In Panel A, the dependent variable is the mean normalized difference turbidity index (NDTI) of the river water on the upstream or downstream side of a mine, in a given month. I derive remotely sensed turbidity from Landsat 7. In Panel B, the dependent variable is the mean normalized difference vegetation index (NDVI) of the land within the 1km buffer along the river on either the upstream or downstream side of a mine, in a given month. I derive remotely sensed yields from the MODIS Combined Terra and Aqua product (MCD43A4.061). *Downstream* is equal to 1 if the side is downstream from the mine and 0 if it is upstream from the mine. *Post* is equal to 1 after the mine opened, 0 otherwise. All models include linear and quadratic controls for mean temperature, precipitation, vapor pressure, wet days, evapotranspiration and cloud cover, as well as mine-side and mine-year-month fixed effects. For turbidity, the sample includes the 30 mines for which non-missing NDTI is observed on the upstream and downstream sides for at least 2 months in each year, for at least 2 years pre and 2 years post-mine opening. For yields, the sample includes the 38 mines for which non-missing NDVI is observed on the upstream and downstream sides for at least 6 months in each year, for at least 2 years pre- and 2 years post-mine opening. Column 1 reports results estimated by pooling months over all 5 seasons: planting, early growing, late growing, harvest and non-farm. Column 2 reports results estimated only over months in the early growing and late growing seasons and Column 3 reports results estimated only over months in non-growing seasons: planting, harvest and non-farm. Standard errors in parentheses are clustered by mine.

6 Mining-induced Air Pollution

To identify areas near mines that are disproportionately exposed to air pollution, I rely on the assumption that wind blows pollutants from the mine towards the downwind side of the mine. I apply high-frequency data on wind direction and the staggered openings of over 100 large-scale mines in two methods that estimate the effect of mining-induced air pollution on local yields: (1) a DID and event study that primarily capture the effects of contemporaneous exposure to air pollution and (2) a distributed lag model that allows for cumulative effects of prolonged exposure to air pollution.

6.1 Contemporaneous effect: The model capturing contemporaneous effects of mining-induced air pollution on local yields is specified as follows:

$$NDVI_{smt} = \alpha_{sm} + \lambda_{mt} + \delta Wind_{smt} + \eta Post_{mt} + \beta_A Wind_{smt} \times Post_{mt} + \mathbf{X}'_{smt} \boldsymbol{\Gamma} + \epsilon_{smt} \quad (4)$$

where $NDVI_{smt}$ is mean NDVI on side s of the 60-km buffer around mine m on date t . $Wind_{smt}$ is the share of days in month t that side s is downwind from mine m , while $Post_{mt}$ is a dummy variable equal to 1 after mine m opened. As in the water pollution specification, I include weather controls, as well as mine-side and mine-year-month fixed effects. Standard errors are clustered at the mine-level. Additionally, I estimate an event study version of Equation 4, which replaces the single $Post_{mt}$ dummy with a series of dummy variables for time since the opening of mine m .

Since wind direction is exogenous to local conditions, factors that may be correlated with the timing and placement of mine openings, as well as local yields, are likely constant across all four directional sides of the buffer around the mine. Essentially, β_A is estimated by comparing sides of a mine that receive more wind to sides of a mine that receive less wind, before and after a mine opening, across over 100 mines. In other words, β_A is identified based on within-month variation in exposure to wind from the mine, across four directional sides.

To establish that mining activity does indeed increase air pollution, I use an almost identical model to the one outlined in Equation 4, with mean AOD as the outcome of interest. Unlike the NDVI data, the AOD data is available only for 2003-2017 and has more missing values due to errors in AOD detection from cloud cover. To ensure a consistent sample of mines between the NDVI and AOD regressions, I use a partially balanced panel of 102 mines, where each mine must be observed on all sides for at least 4 months of every year, for at least 5 years pre- and 5 years post-mine opening. To address measurement error in AOD, I additionally control for the number of non-missing pixel days used to construct mean monthly AOD.

6.2 Cumulative effect: Importantly, the DID and event study designs outlined above are limited in that they primarily estimate the effect of contemporaneous air pollution exposure on yields. This is because the mine-side and mine-date fixed effects ensure a comparison of sides of the same mine, within the same time period, which does not account for pollution exposure in prior time periods. Due to seasonality in winds, the side of a mine that is most frequently downwind varies within the year. This means that the level of downwind exposure of a given side in one month may not necessarily be similar to that side’s exposure in prior months.

To overcome this limitation, I use a distributed lag model that allows me to estimate the effect of cumulative air pollution exposure from mines on yields:

$$NDVI_{smt} = \alpha_{sm} + \lambda_{mt} + \sum_{p=0}^P \beta_p Wind_{sm,t-p} + \sum_{p=0}^P \delta_p Wind_{sm,t-p} \times Post_{m,t-p} + \mathbf{X}'_{smt} \mathbf{\Gamma} + \epsilon_{imt} \quad (5)$$

$NDVI_{smt}$ is mean NDVI on side s of mine m at time t , while $Wind_{sm,t-p}$ is downwind exposure on side s of mine m , p months before t . Downwind exposure is defined as the number of days a side is downwind from the mine, normalized so that one unit corresponds to 30 days of downwind exposure. $Post_{m,t-p}$ is an indicator for when production first began at mine m , lagged by p months. By interacting lagged wind exposure with a lagged indicator for mining activity, I ensure that only wind exposure in the months after a mine opening is included in the cumulative effect estimated for the post period. I include a vector of contemporaneous mine-side-month controls for weather and wind speed, as well as mine-side fixed effects and mine-date fixed effects. Standard errors are clustered at the mine-level.

Before a mine opening, the cumulative effect of exposure to air pollution during the current month and the P months prior is given by $\sum_{p=0}^P \beta_p$. The cumulative effect of pollution exposure after a mine opening is $\sum_{p=0}^P \beta_p + \delta_p$. Thus, we can interpret $\sum_{p=0}^P \delta_p$ as the difference in the cumulative effect of exposure to air pollution during the current and P prior months, between the pre- and post-mine opening periods. Standard errors, p-values and confidence intervals for each lag structure are obtained by testing whether $\sum_{p=0}^P \delta_p = 0$, for each P . In my main specification, I include lagged months up to 2 years prior (24 months) so estimate 25 total coefficients (1 contemporaneous and 24 lagged).

To estimate the distributed lag model, I define a different sample of mines from the set used in the AOD and contemporaneous yield effect estimation shown in Table 2. First, I expand the set of mines to include those where AOD data may be missing or incomplete for certain months, as estimating the coefficients on the main effects and interaction terms for each of the lagged variables in Equation 5 is computationally

demanding. Second, I restrict to only the set of mines where NDVI is observed on all four sides, for all 12 months of every year, for at least 5 years pre- and 5 years post-mine opening. This ensures that NDVI for each mine-side-month is observed for all lags included in the main specification, so that the composition of mines does not vary with the lag structure.

6.3 Results: The AOD event study in Figure 4 reveals that on average, sides of a mine that are downwind more frequently experience higher levels of AOD after a mine opening. This increase in AOD begins about 3 years before the mine officially starts commercial production and remains relatively constant over time. Prior to this 3-year period before a mine opens, we see evidence of parallel trends between areas with high downwind exposure and areas with low downwind exposure.




Figure 4: Average effect of mine openings on air pollution - event study

The figure plots coefficient estimates from a regression of mean AOD on the share of days in a month that a side is downwind from a mine, interacted with a series of event time dummies for years since a mine first opened. The unit of analysis is a mine-side-month, though event time coefficients are binned into yearly increments. The dotted lines show pointwise 95 percent confidence intervals for the coefficients of the event-time path of mean AOD. Coefficients can be interpreted as estimated effects relative to the period three years before the mine opened. The regression includes linear and quadratic controls for mean temperature, precipitation, vapor pressure, wet days, evapotranspiration and cloud cover, as well as mine-side fixed effects, mine-year-month fixed effects, mean wind speed and a linear control for the number of non-missing pixel days used to form mean AOD. The sample includes only the 102 mines for which non-missing AOD is observed on all 4 sides of the mine, for at least 4 months in each year, for at least 5 years pre- and 5 years post-mine opening. Event times less than -5 or greater than 5 are binned into two end points.

Why would air pollution occur before a mine officially starts producing? To answer this question, it is helpful to reconceptualize mining activity as a gradual ramp up over time rather than a binary “on/off.” Detailed work histories of the mines in the S&P database provide evidence in support of this gradual increase in activity. S&P roughly classifies mining activity into the following phases: discovery, exploration/feasibility study, construction/pre-production, production and closure. The discovery phase describes the time at which mineral potential was first determined in a mining area. Exploration and feasibility studies encompass activities such as drilling and testing, which are used to estimate the quality and quantity of reserves. Construction and pre-production include activities related to the building of infrastructure needed for large-scale mineral extraction (e.g. roads, electricity grid, buildings). Production refers to industrial-level extraction and processing of raw materials. Closure refers to the permanent shut-down of a mine and the associated reclamation activities that are supposed to return the land to its previous state.

Literature estimating emissions rates from different activities at a mine site suggests that air pollutants are most likely to be released during the exploration, construction/pre-production and production phases [Patra et al., 2016]. In particular, the use of heavy machinery or vehicles during exploration, as well as the clearing of overburden or construction may contribute to local air pollution levels before the official start of production. As a result, it is not unreasonable to observe increased air pollution in downwind areas prior to the official start of mine production.

To better understand the polluting effects of the different phases of a mine’s life cycle, I use satellite data to estimate a structural break in the mean of remotely sensed nighttime light intensity for each mine that opened between 2003 and 2012 using the methods of Andrews [1993] and Andrews and Ploberger [1994].⁴ I hypothesize that the construction of infrastructure in mining areas would lead to a dramatic increase in nightlights and so interpret the structural break in nightlights as aligning with the construction/pre-production phase. Additional details on the structural breaks estimation procedure are covered in the Appendix 10.13.

For the average mine, I find that a structural break in nightlights occurs about 3 years prior to the S&P start date of commercial production. This is similar to Benschaul-Tolonen [2020], who also uses nightlights to demonstrate an “investment phase” of mining activity occurring 2 years before the start of production. In addition, this finding aligns with Figure 4, which shows that increases in air pollution begin about 3 years before mines start commercial production. I cautiously interpret this result as evidence that air pollutants generated during the construction and pre-production phase of a mine are major drivers of local air pollution from mines, which aligns with emission rate studies from mine sites [Patra et al., 2016].

⁴In 2013 there was a switch from the DMSP OLS to the VIIRS instrument as the source of nightlights. Since this switch introduced an artificial break in nightlights in 2013, I only estimate structural breaks for mines that opened prior to 2013.

After establishing that mining starts to increase air pollution in downwind areas about 3 years before a mine officially begins producing, I turn to estimating the effect of contemporaneous air pollution on remotely sensed yields in mining areas. The event study in Figure 5 reveals no economically or statistically significant effects of contemporaneous exposure to air pollution on NDVI in either the growing or non-growing seasons, at any point in a mine’s life cycle. Importantly, these event study estimates are likely only capturing how concurrent exposure to air pollution a certain number of years after a mine opened affects yields, rather than the effect of prolonged exposure to pollution from a mine. This is because the direction that the wind blows from a mine varies substantially across months within the same year, for most mines in the sample.

The event study findings for both AOD and NDVI are supported by the DID estimates shown in Table 2. Given that mines start to release air pollution about 3 years before official production begins, I define $(Post_{mt}-3)$ to be equal to 1 after the “investment” phase of a mine begins, which is 3 years prior to the official S&P start year. Panel A of Table 2 shows that relative to sides a mine that are downwind less frequently, sides of a mine that are downwind more frequently experience about a statistically increase in AOD after the “investment” phase of a mine begins. The magnitude of this effect is equivalent to roughly a 2% increase and is larger when focusing only on “strongly treated cases”: months where the side that is downwind from the mine most frequently out of all four sides experiences at least 20 days of wind exposure from the mine. In contrast, Panel B shows no economically or statistically significant effect of contemporaneous air pollution on NDVI.

Next, I use the distributed lag model in Equation 5 to examine how cumulative exposure to air pollution over an extended period of time affects crop yields. Instead of being determined solely by contemporaneous wind, air pollution exposure is now a vector of monthly downwind exposure from the current month all the way back to 24 months prior. Pooling all seasons together, Figure 6 reports point estimates and 95 percent confidence intervals for the cumulative effect of lagged and concurrent pollution exposure, proxied by days of wind downwind from a mine. Along the horizontal axis, I increase the maximum response delay from $P = 1$ to $P = 24$ months, estimating a separate distributed lag model for each lag increment. On the vertical axis, I report the NDVI effect for a +30 day increase wind sustained over the $P + 1$ preceding and concurrent months. For example, at $P = 12$ lagged months on the horizontal axis, Figure 6 reports the pre-post difference in cumulative effects, $\sum_{p=0}^{12} \delta_p$, of a sustained 30 day increase in wind from a mine on NDVI over the prior 12 months and the current month. For interpretation, I normalize the coefficients using mean NDVI from 3 years before a mine opens. Appendix Table ?? presents the coefficients, standard errors and p-values for each of the lag increments.

On average, I find no effect of cumulative mining-induced air pollution on NDVI in downwind areas, both when pooling all months (Figure 6) as well as when estimating the model separately by growing and

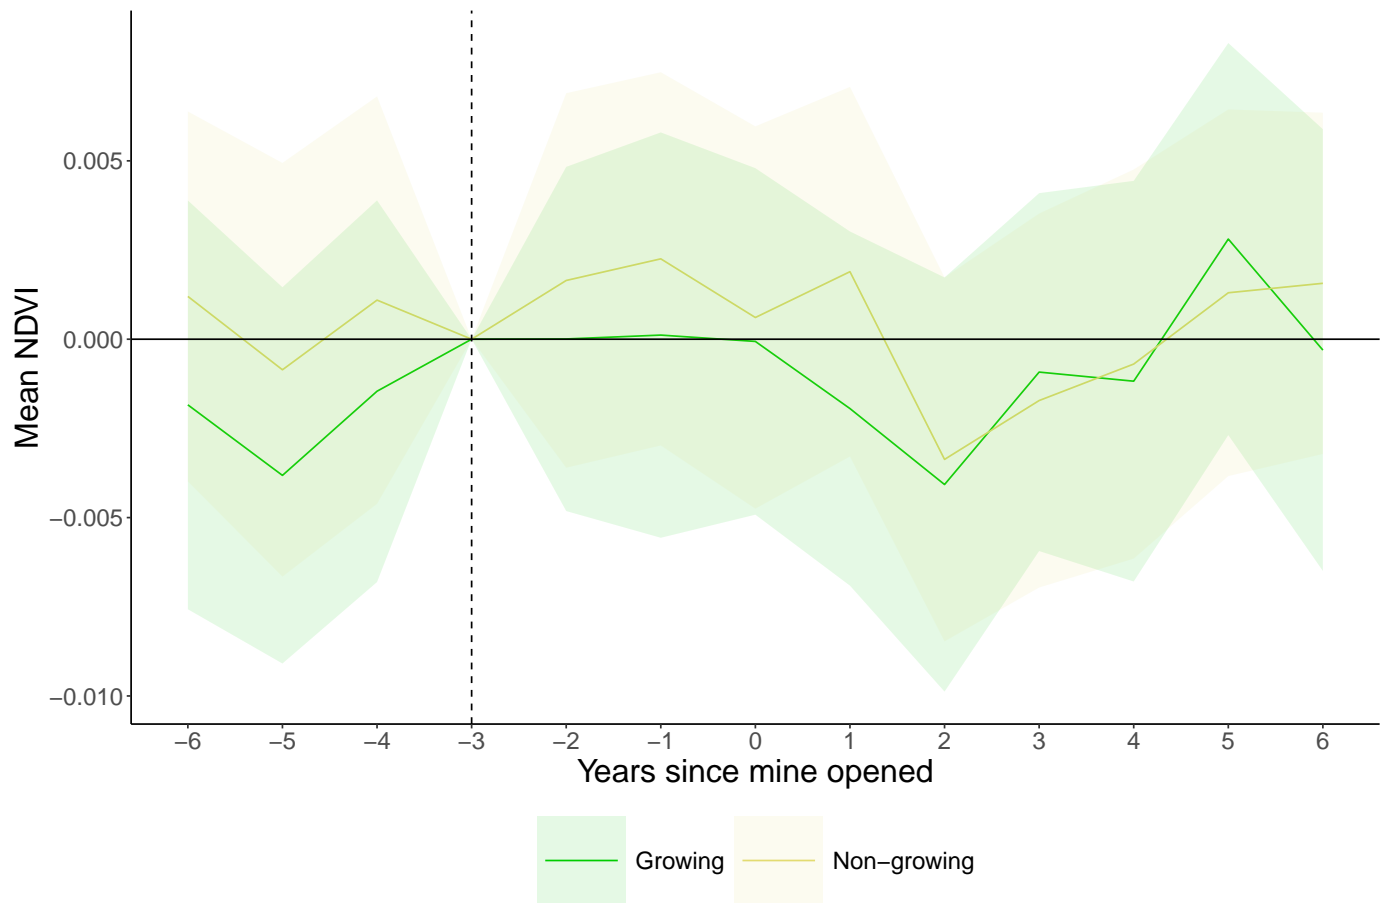


Figure 5: Average effect of contemporaneous mining-induced air pollution on NDVI

The figure plots coefficient estimates from two separate event-study regressions of mean NDVI on the concurrent share of days in a month that a side is downwind from a mine, interacted with a series of event time dummies for years since a mine first opened. The unit of analysis is a mine-side-month, though event time coefficients are binned into yearly increments. The regressions are estimated separately over growing season months (early and late growing) and non-growing season months (planting, harvest and non-farm). The dotted lines show pointwise 95 percent confidence intervals for the coefficients of the event-time path of mean NDVI. Coefficients can be interpreted as estimated effects relative to the period three years before the mine opened. The regression includes linear and quadratic controls for mean temperature, precipitation, vapor pressure, wet days, evapotranspiration and cloud cover, as well as mine-side, mine-year-month fixed effects and mean wind speed. The sample includes only the 102 mines for which non-missing NDVI and AOD is observed on all 4 sides of the mine, for at least 4 months each year, for at least 5 years pre- and 5 years post-mine opening. Event times less than -5 or greater than 5 are binned into two end points.

Table 2: Relationship Between Industrial Mine Openings, Remotely-Sensed Air Pollution and Crop Yields

	(1)	(2)	(3)
	All	20+ days downwind	25+ days downwind
Panel A: First Stage for AOD			
Wind \times (Post - 3)	0.00403*** (0.00147)	0.00443*** (0.00158)	0.00601*** (0.00212)
Mines	102	102	71
Observations	72,596	31,384	16,852
Mean AOD (t-3)	.22	.22	.22
Panel B: Reduced Form for Yields			
Wind \times (Post - 3)	0.00028 (0.00146)	0.00061 (0.00140)	0.00009 (0.00158)
Mines	102	102	71
Observations	72,596	31,384	16,852
Mean NDVI (t-3)	.49	.49	.49
Mine-side FE	Yes	Yes	Yes
Mine-year-month FE	Yes	Yes	Yes
Weather	Yes	Yes	Yes

Notes: Each column reports the results of a linear regression. The unit of analysis is a mine-side-month. In Panel A, the dependent variable is mean aerosol optical depth (AOD) for one of the four sides (N,S, E, W) of a mine within a 60km buffer, in a given month. I derive AOD from the MODIS Aqua Level 2 Daily product (MYD04.3K). In Panel B, the dependent variable is mean normalized difference vegetation index (NDVI) for one of the four sides of a mine within a 60km buffer, in a given month. I derive remotely sensed yields from the MODIS Combined Terra and Aqua product (MCD43A4.061). *Wind* is defined as the share of days in a month that a side is downwind from the mine. *Post - 3* is equal to 1 after the investment phase of a mine, which occurs 3 years before the mine opened, 0 otherwise. All models include linear and quadratic controls for mean temperature, precipitation, vapor pressure, wet days, evapotranspiration and cloud cover, as well as mine-side fixed effects, mine-year-month fixed effects and mean wind speed. The sample includes the 102 mines for which non-missing AOD and NDVI are observed on all four sides of the mine for at least 4 months in each year, for at least 5 years pre and 5 years post-mine opening. Column 1 reports results from a model estimated over all mine months while Column 2 (3) subsets to only mine-months where the side that is downwind most frequently receives at least 20 (25) days of wind from the mine. Standard errors in parentheses are clustered by mine.

non-growing seasons (Appendix Figure 12). Importantly, these average treatment effects mask substantial heterogeneity. For instance, Appendix Table 13 Column 2 shows that downwind areas near open pit mines experience a small but statistically significant reduction in NDVI due to contemporaneous mining-induced air pollution exposure. This aligns with literature suggesting that open pit mines tend to pollute more than underground mines [Sahu et al., 2015], while the latter could be explained by older mines using less advanced, more polluting technologies in their production processes. Other dimensions of heterogeneity are discussed in more detail in Section 8.




Figure 6: Average effect of cumulative exposure to mining-induced air pollution on mean NDVI - pooling seasons

The figure plots the cumulative impact on mean NDVI of +30 days of downwind exposure to air pollution from a mine in each of the concurrent and previous P months. Each value reports the point estimate and 95 percent CI on $\sum_{p=0}^P \delta_p$ from a different distributed lag model, as I increase P along the horizontal axis. $\sum_{p=0}^P \delta_p$ can be interpreted as the average difference in cumulative effects of air pollution on NDVI between the pre and post period, where the post period is defined as the time period after a mine opens. The sample includes the set of 219 mines with non-missing NDVI data observed on all 4 sides, for all 12 months in each year, for at least 5 years pre- and 5 years post-mine opening. Each distributed lag model is estimated using mine-side-months observed after the time period 3 years before a mine first opens, to ensure that lagged wind exposure is observed for at least 2 years before the current month. Each distributed lag model is estimated over months in all seasons.

7 Mixing Pollution Externalities and Market-based Channels

To estimate an “overall” effect of industrial mining on agricultural output through both pollution externalities and market-based channels, I use the staggered openings of 308 mines across Sub-Saharan Africa in a DID that compares NDVI between areas near mines to areas further away, before and after a mine opening. To causally identify the effect of mining on yields through pollution externalities and market-based channels, I rely on the assumption that areas near mines and areas further away will be similarly affected by unobservables that may be correlated with mine openings and local yields, but proximal areas will be more affected by both market and pollution shocks.

First, to motivate the assumption that market channels will disproportionately affect areas near mines, I use findings from existing literature to identify the spatial extent of local markets around mining areas, where we would expect the effects of local input demand shocks to be concentrated. To begin, empirical evidence on commuting distances in urban and rural Africa suggests that areas of 5, 10 or 15 km are likely integrated markets (Amoh-Gymiah and Aidoo, 2013, Kung, Shafer, 2000). Furthermore, the existing literature finds that wealth increases and structural transformation are concentrated within 20-km of a mine. Kotsadam and Tolonen [2016] find that mining leads to large and statistically significant shifts out of agriculture into services for households up to 25-km away, while von der Goltz and Barnwal [2019] find statistically significant wealth increases for households up to 20-km away, with negligible effects further away. Taken together, this literature suggests that areas within 20-km of a mine likely define integrated markets that will be affected by mining activity.

Second, to demonstrate that pollution effects will be concentrated close to mines, I estimate a spatial lag model with AOD as the outcome of interest. Appendix Figure 10 shows that AOD increases are largest within 20-km of a mine, though smaller, statistically significant effects persist up to 60-km away, with no statistically distinguishable effects at greater distances.

Both the existing literature findings are supported by a spatial lag model, shown in Appendix Figure, which highlights that the largest reduction in NDVI occurs within 20 kilometers of a mine, with smaller reductions in NDVI occurring between 20-60 kilometers of a mine and no statistically significant effect at distances greater than 60 kilometers away.⁵

$$NDVI_{gmt} = \alpha_m + \lambda_t + \delta Near_{gm} + \eta Post_{mt} + \beta_O Near_{gm} \times Post_{mt} + \mathbf{X}'_{gmt} \mathbf{\Gamma} + \epsilon_{gmt} \quad (6)$$

where g indexes the distance group (near or far), m indexes the mine and t indexes the date. The near

⁵In the spatial lag model, the omitted category is the 100-150km ring.

group for each mine is defined as the buffer of radius 20 kilometers around the mine centroid, while the far group is defined as the ring between the buffer of 150 km and the buffer of 100 km away from the mine. This radius is chosen based on a spatial lag model: Appendix Figure 11 shows that the largest reduction in NDVI occurs within 20 kilometers of a mine, with smaller reductions in NDVI occurring between 20-60 kilometers of a mine and no statistically significant effect at distances greater than 60 kilometers away.⁶ As a result, I define $Near_{gm}$ equal to 1 within the 20 kilometer buffer and 0 for the 100-150km ring. $Post_{mt}$ is equal to 1 after mine m opened, 0 otherwise. Standard errors are clustered at the mine level.

I include mine-distance group and mine-date fixed effects, as well as linear and quadratic controls for mean temperature, vapor pressure, wet days, precipitation, evapotranspiration and cloud cover. The main parameter of interest is β_0 , which captures the differential impact of an open mine on cropland “near” the mine, relative to cropland further away. Given that the model includes mine-by-distance and mine-by-date fixed effects, β_0 is identified by within-month differences in the change in NDVI in cropland “near” and further away from a mine opening.

In Equation 6, β_0 can be interpreted as the joint effect of pollution externalities and market channels on local yields, as both channels will disproportionately affect areas near mines relative to areas further away. It is theoretically ambiguous whether this joint effect on yields should be larger in magnitude than the effect of just the pollution channels on yields. Based on the theoretical framework outlined in Section 3, the market channels can both generate negative effects on crop yields, for instance by drawing workers out of agriculture into other sectors. However, market channels can also increase local household wealth, which can allow households to invest more in capital inputs to agriculture or switch to crops that may be less sensitive to pollution, helping them to adapt to the effects of mining pollution on yields. If the effect of market channels on local yields are net negative, we would expect that the joint effect of pollution and market based channels estimated through the near-far DID would be larger than the pollution-only effect. The opposite would hold if market effects on yields are net positive or mitigate some of the pollution effect.

Table 3 shows that on average, mine openings lead to a statistically significant decrease in NDVI of about 1.5% in areas near mines relative to areas further away, with slightly larger effects in the growing season. These results provide a helpful baseline against which the estimates generated by isolating the pollution effect can be compared to.

⁶In the spatial lag model, the omitted category is the 100-150km ring.

Table 3: Average effect of mine openings on NDVI - mixing pollution and market channels

	(1)	(2)	(3)
	Pooled	Growing only	Non-growing only
Near \times Post	-0.00638*** (0.000980)	-0.00691*** (0.000990)	-0.00589*** (0.00108)
Number of mines	307	307	307
Obs.	168626	80582	88044
Mine-distance group FE	Yes	Yes	Yes
Mine-year-month FE	Yes	Yes	Yes
Weather	Yes	Yes	Yes
Mean NDVI (t-1)	.476	.543	.415

Each column reports the results of a linear regression. The unit of analysis is a mine-distance group-month. The dependent variable is mean NDVI within a given distance group, in a given month. *Near* is equal to 1 for area within the 20km buffer around the mine and equal to 0 for the area in the ring between the buffers of 100km and 150km. *Post* is equal to 1 after the mine opened, 0 otherwise. All models include linear and quadratic controls for mean temperature, precipitation, vapor pressure, wet days, evapotranspiration and cloud cover, as well as mine and year-month fixed effects. The sample includes the mines for which non-missing NDVI is observed in both distance groups for at least 4 months in each year, for at least 3 years pre- and 3 years post-mine opening. Column 1 reports results estimated by pooling months over all 5 seasons: planting, early growing, late growing, harvest and non-farm. Column 2 reports results estimated only over months in the early growing and late growing seasons and Column 3 reports results estimated only over months in non-growing seasons: planting, harvest and non-farm. Standard errors in parentheses are clustered by mine.

8 Heterogeneous Treatment Effects

8.1 Dimensions explored I investigate heterogeneity in the polluting effects of mines across three key dimensions: (1) governance/regulatory environments, (2) mine characteristics and (3) local economic factors. First, to explore the role of governance on polluting effects, I use the country-level World Bank World Governance Indicators from 2000, the year prior to the start of my NDVI panel. I focus on five main indicators: control of corruption, government effectiveness, regulatory quality, rule of law, and voice and accountability.⁷ Each indicator is based on data from multiple sources, such as survey data, commercial business information providers or non-governmental organizations. The underlying data is aggregated into a score for each of the five indicators, in units of a standard normal distribution from approximately -2.5 to 2.5. Additionally, I examine whether the mine country’s membership to the Extractive Industries Transparency Initiative (EITI), designed to promote transparency and accountability in national resource extraction, can explain variation in the polluting effects of mines.

Second, I examine heterogeneity in effects on NDVI according to the following mine-specific characteristics: extraction method (open pit or non-open pit), commodity type (gold, coal, copper, diamonds), mine

⁷Kaufmann et al. [2010] define each of the indicators as follows: “Control of corruption captures perceptions of the extent to which public power is exercised for private gain, as well as elite or private capture. Government effectiveness captures perceptions of the quality of public services, the quality of policy formulation/implementation and the credibility of government commitment to these policies. Regulatory quality captures perceptions of the government’s ability to formulate and implement sound regulations. Rule of law captures perceptions of the extent to which agents have confidence in and follow the rules of society, such as the quality of contract enforcement, property rights, the police and the courts. Voice and accountability captures perceptions of the extent to which a country’s citizens can exercise rights to freedom of expression, freedom of association and free media, as well as whether they can participate in free and fair elections.”

age, distance to the nearest neighboring mine, distance to the nearest town, AOD in 2003 (start of the AOD data), NDVI in 2001 and precipitation in 2001. Finally, I test for heterogeneity by the following economic factors: mineral rents as share of GDP in 2000, GDP in 2000 and population in 2000. While the mineral rents share is defined at the country level, GDP and population in 2000 are defined specific to the treatment buffers used. Specifically, I calculate mean GDP and population across treatment and control buffers, within each mine.

8.2 “Standard” heterogeneity analysis First, I conduct the “standard” practice used for heterogeneity analysis in the economics literature. To estimate heterogeneous treatment effects of mining activity on local yields, I include a triple interaction of the near dummy, the post dummy and the dimensions of heterogeneity in Equation 6. Similar triple interactions are included in the water and contemporaneous air pollution specifications, with the downstream dummy and the continuous measure of downwind exposure respectively. Lastly, to test for heterogeneous treatment effects of prolonged exposure to air pollution on NDVI, I estimate Equation 5 separately for different groups.

Across the overall, water and air pollution analyses, weak governance and regulatory environments are revealed to be an important driver of the effects of mining activity on yields. Table 4 reveals that reductions in NDVI are 2-3 times larger in areas near mines located in countries with below median performance on governance indicators. The triple interactions are sizable in magnitude and statistically significant across all governance indicators, with the largest effect detected for mines located in countries that are not members of the EITI.

Table 4: Heterogeneous Treatment Effects - Governance and Regulatory Environments

	Base	EITI	Control of Corruption	Govt. Effectiveness	Rule of Law	Reg. Quality	Voice and Accountability
Near × Post	-0.00638*** (0.000980)	-0.00278*** (0.000900)	-0.00405*** (0.00108)	-0.00415*** (0.00110)	-0.00429*** (0.00108)	-0.00426*** (0.00108)	-0.00436*** (0.00111)
Near × Post × Z		-0.00714*** (0.00191)	-0.00499** (0.00198)	-0.00472** (0.00197)	-0.00453** (0.00200)	-0.00472** (0.00202)	-0.00414** (0.00196)
Number of mines	307	307	307	307	307	307	307
Mean NDVI (t-1)	.476	.476	.476	.476	.476	.476	.476

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Similar patterns in heterogeneous treatment effects are uncovered from the air and water regressions. While the triple interactions with the governance and EITI indicators from the water pollution regressions shown in Appendix Table 16 are not statistically significant, likely due to the small sample size, the estimates are overwhelmingly negative. For air pollution, Appendix Figures 15 - 17, which show the heterogeneous treatment effects from the air pollution distributed lag model, reveal that there is a consistent pattern of decreases in NDVI occurring only for mines with below median governance indicators of government

effectiveness, regulatory quality and accountability. Most strikingly, Figure 7 highlights how statistically significant reductions in NDVI occur for mines located in countries that are not members of the EITI, after about 8 months of exposure to air pollution from mines.




Figure 7: Average cumulative effect of mining air pollution on NDVI - heterogeneity by EITI membership

Each panel in the figure plots the cumulative impact on mean NDVI of +30 days of downwind exposure to air pollution from a mine in each of the concurrent and previous P months. The left panel plots the cumulative impact for mines located in countries that are EITI members, while the right panel plots it for mines in non-member countries. Each value reports the point estimate and 95 percent CI on $\sum_{p=0}^P \delta_p$ from a different distributed lag model, as I increase P along the horizontal axis. $\sum_{p=0}^P \delta_p$ can be interpreted as the average difference in cumulative effects of mining air pollution on NDVI between the pre and post period, where the post period is defined as the time period after a mine opens. The sample includes only the set of mines with non-missing NDVI data observed on all 4 sides, for all 12 months in each year, for at least 5 years pre- and 5 years post-mine opening. Each distributed lag model is estimated using mine-side-months observed after the time period 3 years before a mine first opens, to ensure that lagged wind exposure and NDVI is observed for at least 4 years before the current month.

In addition to weak governance, I find that exposure to pollution before a mine opens influences the extent to which mining activity impacts local agriculture. I find across both the air and water pollution analyses that NDVI reductions are higher for cropland in areas initially exposed to lower pollution levels. Appendix Table 13 Column 2 finds that NDVI reductions are about twice as large in cropland near rivers that had below median levels of initial turbidity, with the triple interaction being statistically significant. Similarly, Appendix Figure 28 highlights how statistically significant NDVI reductions occur only for cropland with lower levels of initial exposure to AOD. These effects could potentially be explained by adaptation by plants or farmers. Crops growing in areas that were initially more polluted may have adapted to growing in these sub-optimal conditions and so may be less affected by small increases in air and water pollution due to mining (Hutchinson 1984, Oksanen and Kontunen-Soppela, 2021). Alternatively, farmers that are used to growing crops in more polluted areas may have already developed strategies to mitigate the effects of pollution on their crops, such as investing in fertilizer to improve soil quality that has been eroded by atmospheric deposition or irrigation with contaminated water.

8.3 Machine-learning heterogeneity analysis A unique feature of my empirical setting is that the large, heterogeneous sample of mines allows me to estimate mine-specific DID estimates, then examine both the distribution of these estimates, as well as which characteristics are most predictive of variation in these effects.

Figure 8 plots the distribution of mine-specific overall treatment effects, as well as the water and air pollution effects. It reveals substantial heterogeneity in treatment effects across mines, for all three types of analysis. First, the top panel, which shows the distribution of “overall” effects of mining on local yields, highlights that the majority of treatment effects are negative and concentrated between 0 to -5%. A similar pattern is observed in the water pollution analysis, shown in the bottom left panel, although more mines experience drops in yields that are much larger than the magnitudes observed in the “overall” effect. This provides supporting evidence that some kind of pollution-offsetting effect may be occurring through market-based channels, which is being captured by the near vs. far DID but not the water pollution DID.

In contrast, the bottom right panel of Figure 8 highlights that contemporaneous mining-induced air pollution increases yields in roughly half of the sample but decreases yields in the other half. This aligns with literature documenting a theoretically ambiguous effect of air pollution on local yields, which can depend on local pollution regimes, such as the extent to which the scattering versus absorbing effect of particulate matter dominates, as well as other atmospheric conditions [Schiferl et al., 2018]. One proposed explanation for these findings is the intensity of pollution: Behrer and Wang (2024) show that high density smoke plumes from wildfires decrease yields while low density plumes increase yields. Indeed, Appendix

Figure 35, which plots the normalized AOD DID estimates against the normalized NDVI estimates for each mine, shows that small increases in AOD lead to small increases in yields, while larger increases in AOD lead to yield reductions.




Figure 8: Distribution of Mine-Specific DIDs

The figure plots three histograms showing the distribution of mine-specific DID estimates. The top panel shows the distribution of mine-specific near vs. far DID estimates obtained by estimating a version of Equation 6 for each mine. The bottom left panel shows the distribution of mine-specific downstream vs. upstream DID estimates obtained by estimating a version of Equation 3 for each mine. The bottom right panel shows the distribution of mine-specific continuous wind intensity DID estimates obtained by estimating a version of Equation 4 for each mine. Each mine-specific DID estimate is normalized by mean NDVI in 2000 for that mine, then multiplied by 100 to represent the magnitude of the effect relative to the mean of the outcome at the start of the panel in percent terms.

Next, I apply machine learning (ML) methods to uncover which mine characteristics are most important in explaining variation in the mine-specific treatment effects. The goal of this exercise is to use a data-driven approach to identify dimensions of heterogeneity that drive differences in treatment effects and see whether these findings are complementary to those uncovered in the “standard” heterogeneity analysis. This type of exercise could be valuable and highly applicable in other settings where treated units vary across a large number of characteristics and the researcher wishes to demonstrate that statistically significant findings from “standard” heterogeneity analysis are not spurious.

I train a collection of ML models to predict treatment effects, based on 8 algorithms including linear regression models such as OLS, LASSO, Ridge, and elastic net, as well as non-parametric models such as random forest, a classification and regression tree (CART) and a bagged CART. The unit of observation is a mine, where the outcome of interest is the mine-specific DID estimate for the “overall” effect of mining activity on yields. The main predictors include the five governance indicators, a dummy variable for whether the mine is located in an EITI-member country, GDP in 2000 for the mining area, population in 2000 for the mining area, mineral rents as a share of GDP in 2000, an indicator for whether the mine is open pit, dummy variables indicating commodity type, dummy variables indicating the country that the mine is located in, distance to the nearest mine and distance to the nearest town.⁸

The ML models are trained on the set of 304 mines for which DID estimates are generated and non-missing data for all predictors is available, using 75% of observations for training and the remaining 25% for validation. Model hyperparameters are selected based on a grid search. Out-of-sample accuracy measures values are calculated using 5 fold cross-validation of the training dataset. I therefore obtain 25 hyperparameters combinations for each model (except for OLS), totaling 176 different ML models.

Out of all models trained, the random forest model is identified as the “best” performing model, defined by lowest root mean-squared error (RMSE).⁹ I re-train this random forest model over the entire dataset of DIDs and obtain a measure of variable importance for each predictor.

To estimate variable importance, I construct a measure of “permutation importance,” which leverages the out-of-bag samples for each tree using the following method. First, the RMSE on the out-of-bag sample is calculated. Then, the values of the predictor of interest in the out-of-bag-sample are randomly shuffled, keeping the values of all other predictors the same. Finally, the decrease in RMSE on the shuffled data is measured. Larger reductions in RMSE after permutation imply that a variable is more “important” in contributing to predictive capacity. I opt for permutation importance over Gini impurity, a common

⁸For the country and commodity variables, which are originally categorical, categories accounting for less than 5% of the observations in the sample are binned together into “other.”

⁹The best random forest model is grown on 1000 trees, with 1 predictor randomly sampled at each split and a minimum of 15 data points in each node required for additional splitting to occur.

alternative measure of variable importance for random forest models, as impurity-based importance is less reliable in settings where over-fitting may occur or where many predictors are dummies.

Figure 9 plots the ten most important variables from the best random forest model, based on permutation importance. One striking result is that all five country-level governance indicators are determined to be highly important. This aligns with the results from the “standard” heterogeneity analysis of Table 4, which finds the triple interactions on these five indicators to be highly statistically significant. Similarly, the indicator for whether the mine is located in an EITI-member country, which was also statistically significant in Table 4, falls in the set of the ten most important predictors. These consistent findings between the standard and ML heterogeneity analysis lend credibility to the interpretation that local governance and regulatory environments are important drivers of the extent to which mining activity negatively impacts local agriculture.

However, I note instances where the findings from the “standard” and ML heterogeneity analyses diverge. While the estimates on the triple interactions with mineral rents as a share of 2000 GDP and distance to the nearest neighboring mine are both large and statistically significant in the standard heterogeneity analysis, neither of these variables are considered highly important in the ML analysis.

In summary, this exercise highlights the value of conducting ML heterogeneity analysis to cross-check “standard” heterogeneity analysis in settings where unit-specific treatment effects can be estimated. In Appendix Figure 36, I also show that the main conclusions drawn from the ML analysis discussed above are robust to estimating a classification model, where the outcome is a dummy variable equal to one if the mine-specific DID estimate is negative. While the ranking of the most important variables slightly changes, four of the five country-level governance indicators still rank within the top five most important predictors. Interestingly, distance to a neighboring mine and mineral rents share, which were highly significant in the “standard” heterogeneity analysis but determined less important in the main ML analysis discussed above, are revealed to be highly important in this alternative classification random forest model.




Figure 9: Most Important Variables from Best Random Forest Model

9 Quantifying Results

In this section, I ask two important questions to quantify the effect of mining on yields. First, what share of the average “overall” effect of industrial mining on local agricultural output can be explained by pollution alone? Second, how can the estimated NDVI effects be mapped to effects on actual yields in kilograms per hectare?

9.1 Estimating Share of Effect on Yields Driven By Pollution From the distributed lag results in Table ??, we see that at most, NDVI falls by 0.0007 units in downwind areas due to cumulative and contemporaneous mining-induced air pollution. I scale this estimate by 0.25, to reflect the fact that on average, 25% of the 20 km buffer around mine is downwind most frequently in a given month. This yields a very small scaled treatment effect of 0.000175 for air pollution.

In contrast, we see a much larger effect when looking at exposure to water pollution: Table 1 Column 2 shows that on average, NDVI is 0.0177 lower in downstream areas during the growing seasons. I re-scale this treatment effect by 0.17, which is the 25th percentile of the share of the 20km buffer that is at a lower elevation than the mine, and so could be affected by water flow from the mine, across all mines within 1km of a river. The resulting scaled effect of water pollution on NDVI is 0.003009, which serves as a lower bound for the combined effect of both air and water pollution on local yields.

Taken together with the overall effect of industrial mining on growing season NDVI of -0.00691 through both pollution and market based channels, from Column 2 of Table 3, a basic back-of-the-envelope calculation suggests that approximately 44% of the overall effect of mining on NDVI can be explained by air and water pollution externalities. As this share is quite large, it is important to keep in mind several caveats. First, the heterogeneity analysis suggests that this share likely varies substantially across mines. Second, market-effects could spur household adaptation to pollution occurring in areas near mines, which may attenuate the near-far DID estimate and make the share of this effect attributed to pollution appear larger.

As wealth increases from local mining activity are non-directional, they should not bias the air and water pollution DID estimates. However, wealth increases will disproportionately affect areas near mines and could offset some of the negative effects of pollution also experienced by these areas. For example, increased wealth could be invested into inputs such as fertilizer or other more productive technologies that may counteract the effects of mining-induced pollution on crop health. Alternatively, increased wealth could be used to treat health issues that may arise due to polluted air or water. This pollution offsetting effect of wealth may result in a smaller estimated NDVI reduction in the near vs. far DID, thus making making the denominator in the share smaller and making the share of the overall effect of mining on yields attributed to pollution appear

larger.

Is it realistic for increased wealth from mining activity to offset the effects of increased pollution? Existing work has documented large and statistically significant wealth increases from mining, with household wealth in the medium run rising by about 0.3 standard deviations for those in the vicinity of a mine [von der Goltz and Barnwal, 2019]. The authors argue that this effect is comparable to having an electricity connection or living in a dwelling with finished flooring in the case of Peru in the year 2000, and to the effect of owning a motorbike or mobile phone in the case of Burkina Faso, in the year 2010.

A thorough investigation of the magnitude and the mechanisms through which pollution can offset wealth would require detailed geo-located survey data that covers household agricultural input use and health behaviors, before and after a mine opening. While some data sources that cover this information do exist, such as the LSMS-ISA, their spatial and temporal coverage does not align with the timing and location of mine openings in my sample. As a result, examining the mechanisms through which wealth may offset pollution effects is beyond the scope of my paper.

9.2 Translating NDVI Effect to Yield Effect Second, how can the effect on yields measured by NDVI be translated to yields measured in kilograms per hectare? Ideally, we would observe farm-level yields around mining areas over time. However, this data is not available in Sub-Saharan Africa, making the use of NDVI as a proxy for yields more appealing. Like many remotely-sensed variables, NDVI is an imperfect measure of the true outcome of interest and subject to different types of measurement error, some of which may be non-classical. Work by Proctor et al. (2023) suggests that measurement error in remotely sensed dependent variables would most likely result in attenuation of estimated coefficients, as well as possible underestimation of the standard errors on these coefficients. While some methods exist to correct for measurement error in remotely-sensed variables, such as the multiple imputation method proposed by Proctor et al. (2023), there is insufficient proximal data on ground-based yields to implement these methods in my setting. I discuss NDVI measurement error issues, as well as the data limitations that prevent me from addressing it, in Appendix 10.14.

To translate estimated NDVI effects to yield effects, I conduct the following quantification exercise using data on plot-level yields. Data on actual maize yields, defined in kilograms of maize produced per hectare, are obtained from Aramburu-Merlos et al. [2024]. This dataset covers over 13,000 smallholder farms across 7 countries: Burundi, Uganda, Nigeria, Tanzania, Kenya, Rwanda and Zambia. It is a repeated-cross section, where measurements are taken over the 2016-2022 time period from different plots located in similar areas. The data from Aramburu-Merlos et al. [2024] was collected by One Acre Fund, an NGO that provides smallholder farmers with access to credit, training, crop insurance and farming supplies. Maize yields for

each farmer at the time of harvest were measured in two randomly spaced boxes of 36 square meters, avoiding field edges. While farm size is not provided in the public data of Aramburu-Merlos et al. [2024], African agriculture is dominated by small farms, typically defined as less than 1 acre [Carletto et al., 2015].

To link plot-level yields to NDVI, I define a grid of 1km x 1km cells over the study areas in Aramburu-Merlos et al. [2024]. Each plot is linked to a grid-cell using plot GPS coordinates. I opt for the 1km x 1km grid cell to limit the influence of inaccuracies in plot coordinates [Jin et al., 2017]. While I do not observe the same plots over time, most grid cells can be observed for at least 2 different years. For each grid cell, I calculate the mean NDVI by year, within each of the following seasons: planting, early growing, late growing, harvest and non-farm. Mean seasonal NDVI at the cell-year level is then linked to mean actual yields in kg/ha, which are calculated by taking the average of observed yields across all plots in each cell-year.

To link remotely sensed yields to observed yields, I run the following regression:

$$Yield_{ct} = \alpha_c + \lambda_t + \beta_{NDVI}NDVI_{ct} + \Gamma X_{pct} + \epsilon_{pct} \quad (7)$$

where $Yield_{ct}$ is mean maize yields in kg/ha across all plots p in cell c in year t and $NDVI_{ct}$ is mean NDVI in one of the five seasons. I include grid-cell and year fixed effects, as well as cell-level controls for the average number of maize growing degree days, average precipitation over the maize season and average temperature over the maize season from Aramburu-Merlos et al. [2024], across all plots within a cell-year. I also include cell-level weather variables such as mean cloud cover and evapotranspiration from Harris et al. (2016). $NDVI_{ct}$ is scaled so that β_{NDVI} reflects the effect of a 0.01 unit increase in NDVI on maize yields. Inclusion of weather controls and fixed effects help limit omitted variable bias in estimation of β_{NDVI} due to NDVI detection errors. These errors can be driven by a variety of factors, most notably cloud cover, but also surface reflectance, canopy thickness, the level of atmospheric aerosols and sensor errors. Detection errors make NDVI observed from satellite data an imperfect measure of NDVI in the presence of no detection errors. While some types of detection errors, such as satellite sensor errors, are likely uncorrelated with plot-level yields, detection errors driven by cloud cover or aerosol loading could be correlated with plot-level yields through the direct effect of atmospheric determinants on crop health.

Table 5 shows that on average, a 0.01 unit increase in NDVI is associated with a statistically significant 52.10 kg/ha increase in maize yields during the early growing season, with smaller increases in other seasons. This is supported by existing literature, which also finds that NDVI is a strong predictor of yields specifically during the early stages of growing [Panek and Gozdowski, 2021]. Linking this result to the overall effect of mining on NDVI of -0.00691 suggests that on average, industrial mining reduces actual yields for a smallholder farmer by about 36 kilograms per hectare. This corresponds to about a 1.2 % decrease in actual

yields relative to the mean of 3000 kg/ha, which is a fifth of the magnitude of the reduction in yields caused by an additional dust storm in Iran [Birjandi-Feriz and Yousefi, 2017] or the reduction in yields caused by the Acid Rain Program in the USA [Sanders and Barreca, 2022]. Importantly, these other estimates were generated using administrative data on yields, so may not suffer from attenuation bias due to measurement error common to remotely sensed variables (discussed in more detail in Appendix 10.14).

Table 5: Relationship between Cell-level NDVI and Cell-level Yields, cell-size = 1000m

	(1)	(2)	(3)	(4)	(5)
	Planting	Early Growing	Late Growing	Harvest	Nonfarm
Mean NDVI	11.83 (16.42)	52.10** (22.04)	20.25 (21.11)	24.43 (17.69)	9.937 (14.62)
Year FE	Yes	Yes	Yes	Yes	Yes
Grid cell FE	Yes	Yes	Yes	Yes	Yes
Include weather controls	Yes	Yes	Yes	Yes	Yes
Mean Yields	3005.643	3005.643	3005.643	3005.643	3005.643
Obs.	1127	1153	1153	1141	1153
R-sq	0.631	0.638	0.632	0.634	0.632

Each column reports the results of an OLS regression. The unit of observation is a cell-year. The dependent variable is mean plot-level yields in kg/ha across all plots falling within a cell, in a given year. *Mean NDVI* is the average NDVI in the cell of size 1000m containing the plot, across days in a particular season during the year that maize on the plot was harvested. The columns indicate the season for which NDVI is calculated during the year the plot was harvested. *Mean NDVI* is scaled so that one unit represents a 0.01 increase in NDVI. Each regression includes linear controls for cell-level averages of growing degree days, temperature and precipitation during the maize season across plots falling within the cell, as well as cell-level controls for mean temperature, precipitation, vapor pressure, cloud cover, evapotranspiration and wet days. Year and cell-level fixed effects are included in the regressions.

Table 5 also shows that the NDVI model at the cell-level can explain over 50% of the variation in plot-level yields, which is similar to the R^2 found in other work linking the One Acre Fund data to remotely sensed yields (Burke and Lobell [2017]; Jin et al. [2017]). Conditional on observing a sufficiently high R^2 , in Appendix 10.14 I discuss whether the estimated relationship between NDVI and local yields from Equation 7 may be biased in my setting. Weather controls and fixed effects limit the potential bias induced by measurement error in the estimation of Equation 7. Furthermore, Appendix Tables 20 - 22 demonstrate robustness to varying the size of the grid cell used in the quantification exercise, showing that NDVI during the growing season is most strongly correlated with observed plot-level yields.

References

- Donald W. K. Andrews. Tests for Parameter Instability and Structural Change With Unknown Change Point. *Econometrica*, 61(4):821–856, 1993. ISSN 0012-9682. doi: 10.2307/2951764. URL <https://www.jstor.org/stable/2951764>. Publisher: [Wiley, Econometric Society].
- Donald W. K. Andrews and Werner Ploberger. Optimal Tests when a Nuisance Parameter is Present Only

- Under the Alternative. *Econometrica*, 62(6):1383–1414, 1994. ISSN 0012-9682. doi: 10.2307/2951753. URL <https://www.jstor.org/stable/2951753>. Publisher: [Wiley, Econometric Society].
- Fernando M. Aragón and Juan Pablo Rud. Natural Resources and Local Communities: Evidence from a Peruvian Gold Mine. *American Economic Journal: Economic Policy*, 5(2):1–25, May 2013. ISSN 1945-7731. doi: 10.1257/pol.5.2.1. URL <https://www.aeaweb.org/articles?id=10.1257/pol.5.2.1>.
- Fernando M. Aragón and Juan Pablo Rud. Polluting Industries and Agricultural Productivity: Evidence from Mining in Ghana. *The Economic Journal*, 126(597):1980–2011, 2016. ISSN 1468-0297. doi: 10.1111/eoj.12244. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/eoj.12244>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/eoj.12244>.
- Fernando Aramburu-Merlos, Fatima A. M. Tenorio, Nester Mashingaidze, Alex Sananka, Stephen Aston, Jonathan J. Ojeda, and Patricio Grassini. Adopting yield-improving practices to meet maize demand in Sub-Saharan Africa without cropland expansion. *Nature Communications*, 15(1):4492, May 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-48859-0. URL <https://www.nature.com/articles/s41467-024-48859-0>. Publisher: Nature Publishing Group.
- Sam Asher and Paul Novosad. Rent-Seeking and Criminal Politicians: Evidence from Mining Booms | The Review of Economics and Statistics | MIT Press. *The Review of Economics and Statistics*, 2023. URL <https://direct.mit.edu/rest/article/105/1/20/102838/Rent-Seeking-and-Criminal-Politicians-Evidence>.
- Sebastian Axbard, Anja Benschaul-Tolonen, and Jonas Poulsen. Natural resource wealth and crime: The role of international price shocks and public policy. *Journal of Environmental Economics and Management*, 110:102527, October 2021. ISSN 0095-0696. doi: 10.1016/j.jeem.2021.102527. URL <https://www.sciencedirect.com/science/article/pii/S0095069621000905>.
- Anja Benschaul-Tolonen. Local Industrial Shocks and Infant Mortality. *The Economic Journal*, n/a(n/a), 2020. ISSN 1468-0297. doi: 10.1111/eoj.12625. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/eoj.12625>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/eoj.12625>.
- Nicolas Berman, Mathieu Couppenier, Dominic Rohner, and Mathias Thoenig. This Mine is Mine! How Minerals Fuel Conflicts in Africa. *American Economic Review*, 107(6):1564–1610, June 2017. ISSN 0002-8282. doi: 10.1257/aer.20150774. URL <https://pubs.aeaweb.org/doi/10.1257/aer.20150774>.
- Maliheh Birjandi-Feriz and Kowsar Yousefi. When the Dust Settles: Productivity and Economic Losses Following Dust Storms, November 2017. URL <https://papers.ssrn.com/abstract=3230265>.

- Leonardo Bonilla Mejía. Mining and human capital accumulation: Evidence from the Colombian gold rush. *Journal of Development Economics*, 145:102471, June 2020. ISSN 0304-3878. doi: 10.1016/j.jdeveco.2020.102471. URL <https://www.sciencedirect.com/science/article/pii/S0304387820300468>.
- Marshall Burke and David B. Lobell. Satellite-based assessment of yield variation and its determinants in smallholder African systems. *Proceedings of the National Academy of Sciences*, 114(9):2189–2194, February 2017. doi: 10.1073/pnas.1616919114. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1616919114>. Publisher: Proceedings of the National Academy of Sciences.
- Calogero Carletto, Sydney Gourlay, and Paul Winters. From Guesstimates to GPStimates: Land Area Measurement and Implications for Agricultural Analysis. *Journal of African Economies*, 24(5):593–628, November 2015. ISSN 0963-8024, 1464-3723. doi: 10.1093/jae/ejv011. URL <https://academic.oup.com/jae/article-lookup/doi/10.1093/jae/ejv011>.
- Francesco Caselli and Guy Michaels. Do Oil Windfalls Improve Living Standards? Evidence from Brazil. *American Economic Journal: Applied Economics*, 5(1):208–238, January 2013. ISSN 1945-7782. doi: 10.1257/app.5.1.208. URL <https://www.aeaweb.org/articles?id=10.1257/app.5.1.208>.
- Burhan U. Choudhury, Akbar Malang, Richard Webster, Kamal P. Mohapatra, Bibhash C. Verma, Manoj Kumar, Anup Das, Mokidul Islam, and Samarendra Hazarika. Acid drainage from coal mining: Effect on paddy soil and productivity of rice. *The Science of the Total Environment*, 583:344–351, April 2017. ISSN 1879-1026. doi: 10.1016/j.scitotenv.2017.01.074.
- James Cust and Albert Zeufack. Africa’s Resource Future. 2023. URL <https://documents.worldbank.org/en/publication/documents-reports/documentdetail/099080123145011993/p16722906c03ca09409ace06cb32991395b>.
- Dominik Dietler, Andrea Farnham, Georg Loss, Günther Fink, and Mirko S. Winkler. Impact of mining projects on water and sanitation infrastructures and associated child health outcomes: a multi-country analysis of Demographic and Health Surveys (DHS) in sub-Saharan Africa. *Globalization and Health*, 17(1):70, June 2021. ISSN 1744-8603. doi: 10.1186/s12992-021-00723-2.
- Stanislaw Dudka and Domy C. Adriano. Environmental Impacts of Metal Ore Mining and Processing: A Review. *Journal of Environmental Quality*, 26(3):590–602, 1997. ISSN 1537-2537. doi: 10.2134/jeq1997.00472425002600030003x. URL <https://onlinelibrary.wiley.com/doi/abs/10.2134/jeq1997.00472425002600030003x>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.2134/jeq1997.00472425002600030003x>.

- Zhaozhong Feng and Kazuhiko Kobayashi. Assessing the impacts of current and future concentrations of surface ozone on crop yield with meta-analysis. *Atmospheric Environment*, 43(8):1510–1519, March 2009. ISSN 1352-2310. doi: 10.1016/j.atmosenv.2008.11.033. URL <https://www.sciencedirect.com/science/article/pii/S1352231008010972>.
- Nicolas Gendron-Carrier, Marco Gonzalez-Navarro, Stefano Polloni, and Matthew A. Turner. Subways and Urban Air Pollution. *American Economic Journal: Applied Economics*, 14(1):164–196, January 2022. ISSN 1945-7782. doi: 10.1257/app.20180168. URL <https://www.aeaweb.org/articles?id=10.1257/app.20180168>.
- M. K. Ghose and S. R. Majee. Assessment of the Status of Work Zone Air Environment Due to Opencast Coal Mining. *Environmental Monitoring and Assessment*, 77(1):51–60, July 2002. ISSN 1573-2959. doi: 10.1023/A:1015719625745. URL <https://doi.org/10.1023/A:1015719625745>.
- I. Harris, P.D Jones, T.J. Osborn, , and D.H. Lister. Updated high-resolution grids of monthly climatic observations—cru ts3.26 dataset. *International Journal of Climatology*, 2014.
- Matias Heino, Pekka Kinnunen, Weston Anderson, Deepak K. Ray, Michael J. Puma, Olli Varis, Stefan Siebert, and Matti Kummu. Increased probability of hot and dry weather extremes during the growing season threatens global crop yields. *Scientific Reports*, 13(1):3583, March 2023. ISSN 2045-2322. doi: 10.1038/s41598-023-29378-2. URL <https://www.nature.com/articles/s41598-023-29378-2>. Publisher: Nature Publishing Group.
- Qing Huang, Victoria Wenxin Xie, and Wei You. Resource Rents, Urbanization, and Structural Transformation. *SSRN Electronic Journal*, 2023. ISSN 1556-5068. doi: 10.2139/ssrn.4543449. URL <https://www.ssrn.com/abstract=4543449>.
- Zhenong Jin, George Azzari, Marshall Burke, Stephen Aston, and David B. Lobell. Mapping Smallholder Yield Heterogeneity at Multiple Scales in Eastern Africa. *Remote Sensing*, 9(9):931, September 2017. ISSN 2072-4292. doi: 10.3390/rs9090931. URL <https://www.mdpi.com/2072-4292/9/9/931>. Number: 9 Publisher: Multidisciplinary Digital Publishing Institute.
- Daniel Kaufmann, Aart Kraay, and Massimo Mastruzzi. The Worldwide Governance Indicators: Methodology and Analytical Issues, September 2010. URL <https://papers.ssrn.com/abstract=1682130>.
- Andreas Kotsadam and Anja Tolonen. African Mining, Gender, and Local Employment. *World Development*, 83:325–339, July 2016. ISSN 0305-750X. doi: 10.1016/j.worlddev.2016.01.007. URL <https://www.sciencedirect.com/science/article/pii/S0305750X1600005X>.

- R.C. Levy, S. Mattoo, L.A. Munchak, L.A. Remer, A.M. Sayer, F. Patadia, and N.C. Hsu. The collection 6 modis aerosol products over land and ocean. *Atmospheric Measurement Techniques*, 2013.
- Peng Liu, Qiumei Wu, Wenyong Hu, Kang Tian, Biao Huang, and Yongcun Zhao. Effects of atmospheric deposition on heavy metals accumulation in agricultural soils: Evidence from field monitoring and Pb isotope analysis. *Environmental Pollution*, 330:121740, August 2023. ISSN 0269-7491. doi: 10.1016/j.envpol.2023.121740. URL <https://www.sciencedirect.com/science/article/pii/S026974912300742X>.
- Xiang Liu and Ankur R. Desai. Significant Reductions in Crop Yields From Air Pollution and Heat Stress in the United States. *Earth's Future*, 9(8):e2021EF002000, 2021. ISSN 2328-4277. doi: 10.1029/2021EF002000. URL <https://onlinelibrary.wiley.com/doi/abs/10.1029/2021EF002000>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2021EF002000>.
- Davis Lobell, Stefania Di Tomasso, and Jennifer Burney. Globally ubiquitous negative effects of nitrogen dioxide on crop growth. *Science Advances*, 2022. doi: 10.1126/sciadv.abm9909. URL <https://www.science.org/doi/10.1126/sciadv.abm9909>.
- R. Maggs, A. Wahid, S. R. A. Shamsi, and M. R. Ashmore. Effects of ambient air pollution on wheat and rice yield in Pakistan. *Water, Air, and Soil Pollution*, 85(3):1311–1316, December 1995. ISSN 1573-2932. doi: 10.1007/BF00477163. URL <https://doi.org/10.1007/BF00477163>.
- Phenny Mwaanga, Mathews Silondwa, George Kasali, and Paul M. Banda. Preliminary review of mine air pollution in Zambia. *Heliyon*, 5(9):e02485, September 2019. ISSN 2405-8440. doi: 10.1016/j.heliyon.2019.e02485. URL <https://www.sciencedirect.com/science/article/pii/S2405844019361456>.
- Ewa Panek and Dariusz Gozdowski. Relationship between MODIS Derived NDVI and Yield of Cereals for Selected European Countries. *Agronomy*, 11(2):340, February 2021. ISSN 2073-4395. doi: 10.3390/agronomy11020340. URL <https://www.mdpi.com/2073-4395/11/2/340>. Number: 2 Publisher: Multidisciplinary Digital Publishing Institute.
- Aditya Kumar Patra, Sneha Gautam, and Prashant Kumar. Emissions and human health impact of particulate matter from surface mining operation—A review. *Environmental Technology & Innovation*, 5: 233–249, April 2016. ISSN 2352-1864. doi: 10.1016/j.eti.2016.04.002. URL <https://www.sciencedirect.com/science/article/pii/S2352186416300153>.
- E. Petavratzi, S. Kingman, and I. Lowndes. Particulates from mining operations: A review of sources, effects and regulations. *Minerals Engineering*, 18(12):1183–1199, October 2005. ISSN 0892-6875.

doi: 10.1016/j.mineng.2005.06.017. URL <https://www.sciencedirect.com/science/article/pii/S0892687505002050>.

Tomás Rau, Sergio Urzúa, and Loreto Reyes. Early Exposure to Hazardous Waste and Academic Achievement: Evidence from a Case of Environmental Negligence. *Journal of the Association of Environmental and Resource Economists*, 2(4):527–563, December 2015. ISSN 2333-5955. doi: 10.1086/683112. URL <https://www.journals.uchicago.edu/doi/full/10.1086/683112>. Publisher: The University of Chicago Press.

William J. Sacks, Delphine Deryng, Jonathan A. Foley, and Navin Ramankutty. Crop planting dates: an analysis of global patterns. *Global Ecology and Biogeography*, 19(5):607–620, 2010. ISSN 1466-8238. doi: 10.1111/j.1466-8238.2010.00551.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1466-8238.2010.00551.x>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1466-8238.2010.00551.x>.

H. B. Sahu, N. Prakash, and S. Jayanthu. Underground Mining for Meeting Environmental Concerns – A Strategic Approach for Sustainable Mining in Future. *Procedia Earth and Planetary Science*, 11:232–241, January 2015. ISSN 1878-5220. doi: 10.1016/j.proeps.2015.06.030. URL <https://www.sciencedirect.com/science/article/pii/S1878522015000818>.

Nicholas J. Sanders and Alan I. Barreca. Adaptation to Environmental Change: Agriculture and the Unexpected Incidence of the Acid Rain Program. *American Economic Journal: Economic Policy*, 14(1): 373–401, February 2022. ISSN 1945-7731. doi: 10.1257/pol.20190060. URL <https://www.aeaweb.org/articles?id=10.1257/pol.20190060>.

Luke D. Schiferl, Colette L. Heald, and David Kelly. Resource and physiological constraints on global crop production enhancements from atmospheric particulate matter and nitrogen deposition. *Biogeosciences*, 15(14):4301–4315, July 2018. ISSN 1726-4170. doi: 10.5194/bg-15-4301-2018. URL <https://bg.copernicus.org/articles/15/4301/2018/>. Publisher: Copernicus GmbH.

Pyae Sone Soe, Win Thiri Kyaw, Koji Arizono, Yasuhiro Ishibashi, and Tetsuro Agusa. Mercury Pollution from Artisanal and Small-Scale Gold Mining in Myanmar and Other Southeast Asian Countries. *International Journal of Environmental Research and Public Health*, 19(10):6290, May 2022. ISSN 1660-4601. doi: 10.3390/ijerph19106290.

Jan von der Goltz and Prabhat Barnwal. Mines: The local wealth and health effects of mineral mining in developing countries. *Journal of Development Economics*, 139:1–16, June 2019. ISSN 0304-

3878. doi: 10.1016/j.jdeveco.2018.05.005. URL <https://www.sciencedirect.com/science/article/pii/S0304387818304875>.

10 Appendix

10.1 Data Construction

10.1.1 Water Pollution To examine the effect of mining on yields through the water pollution channel, I select only mines that are within 1 kilometer of a river. I use the HydroRIVERS data to identify rivers located in Africa. This dataset is a vectorized line network of rivers that have a catchment area of at least 10 km^2 , an average river flow of at least $0.1 \text{ m}^3/\text{sec}$, or both. Each observation in this dataset is a river segment, where segment end points are defined by splitting rivers or streams at nodes where they fork. Each mine within 1 km of a river in this network is “snapped” to the closest river segment, then the upstream and downstream river segments from the mine river segment are identified. The upstream and downstream segments are buffered by 1-kilometer on either side of the line to identify land area around the rivers that would likely be affected by acid mine drainage either through seepage or irrigation. The “side” of a mine in the water pollution analysis refers to the upstream or downstream buffer.

10.2 Air Pollution The algorithm detecting AOD can perform poorly in certain cases, which generates missing values. Levy et al. [2013] find that the algorithm performs poorly over light surfaces, such as desert or snowy regions. To address this limitation, I exclude mining areas that are adjacent to the desert, such as parts of Burkina Faso and Mali that neighbor the Sahel. In addition, the MODIS instruments can only detect AOD on cloud-free days, so missing data in certain regions or on certain days tend to reflect high levels of cloud cover. This results in seasonality of missingness in the MODIS data; I manage this issue by using mine-year-month fixed effects in my model, which account for mine-specific monthly variations in weather patterns correlated with satellite errors. Furthermore, in regressions with AOD as the outcome variable, I control for the number of non-missing pixels used to construct mean monthly AOD. I opt for a linear control of non-missing pixel count over interpolating missing AOD as interpolation algorithms may introduce measurement error that is difficult to interpret.

10.3 NDVI

10.4 Agricultural Seasons For each mining area, I observe 4-5 agricultural seasons. Like Sacks et al. [2010], I define the seasons for planting and harvest using start and end dates, and define the growing season as the days in between the planting end date and the harvest start date. As the biological sciences literature suggests that pollution can affect crops more strongly during the reproductive phase of the plant,

which occurs early in the growing season, I divide this growing season in half to make early and late growing seasons Liu and Desai [2021]. Finally, I define the non-farming season as the time between the harvest end date of one year and the planting start date of the next year. In mining areas where the harvesting period of one year overlaps with the planting period of the subsequent year, I define the harvest season as ending at the start of next year’s planting season. I assign each mining area the modal season dates across all cells falling within the buffer linked to that mine.

Table 6: Relationship between Total Suspended Solids from Ground-based Sampling and Remotely-Sensed Turbidity

	(1)	(2)
Normalized Difference Turbidity Index	940.9*** (341.4)	887.1** (364.7)
Number of sampling sites	178	178
Include weather controls	No	Yes
TSS mean	147.012	147.012
NDTI mean	.007	.007

Each column reports the results of a linear regression. The unit of analysis is a sampling site on the date that the water sampling occurred. The dependent variable is the level of total suspended solids identified from ground-based water sampling and testing at the sampling site, on a specific date. The independent variable is the mean normalized difference turbidity index for the river segment falling within a 500 meter buffer around the sampling site, in the month corresponding to the sampling date. All models include linear and quadratic controls for mean temperature, precipitation, vapor pressure, wet days, evapotranspiration and cloud cover. Robust standard errors are reported in parentheses

10.5 NDTI vs. ground-based turbidity measurements

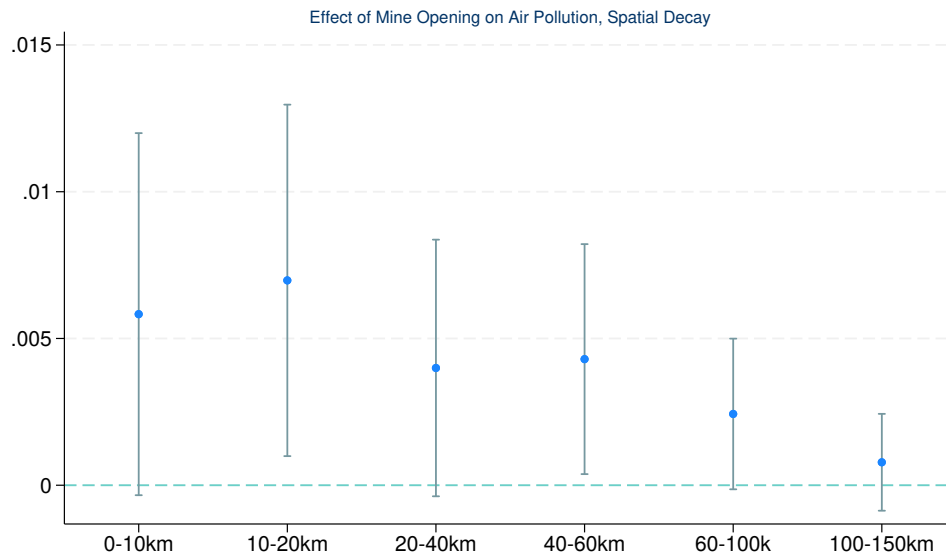


Figure 10: Mean AOD Spatial Decay Model

10.6 AOD Spatial Lag Model

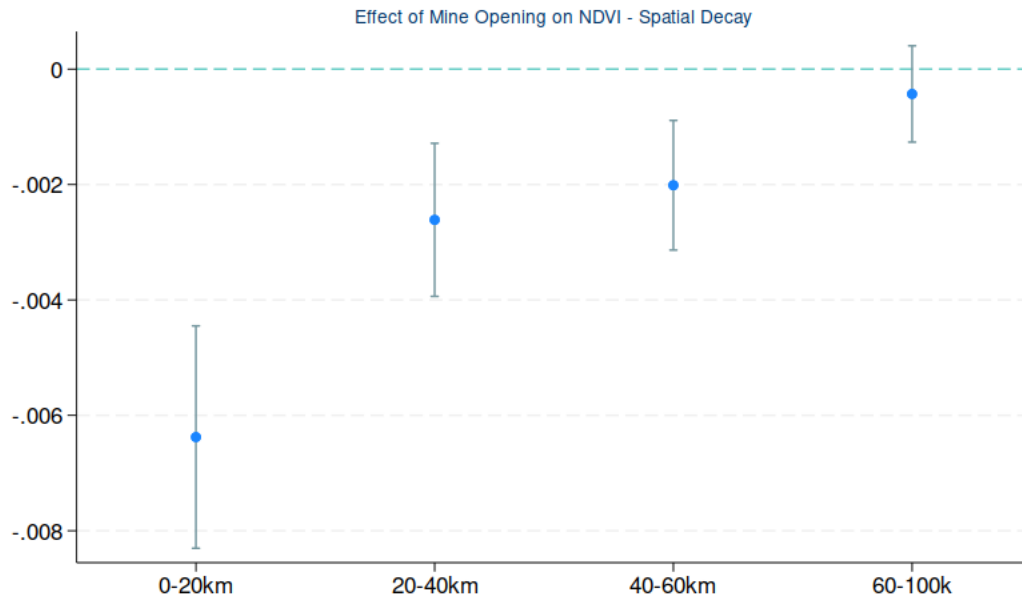


Figure 11: Mean NDVI Spatial Decay Model

10.7 NDVI Spatial Lag Model

10.8 Distributed

Effect of Cumulative Air Pollution Exposure on NDVI

Number of Months Prior	$\sum_p^P \delta_p$	$\frac{\sum_p^P \delta_p}{\text{Mean NDVI}}$	Standard error	p
1	0.0003	0.0656	0.0011	0.7956
2	0.0006	0.1307	0.0012	0.6483
3	0.0005	0.1263	0.0014	0.6890
4	0.0007	0.1501	0.0015	0.6616
5	0.0007	0.1581	0.0016	0.6672
6	0.0005	0.1240	0.0017	0.7488
7	0.0004	0.0869	0.0018	0.8297
8	0.0002	0.0405	0.0018	0.9220
9	0.0001	0.0222	0.0018	0.9579

10	0.0001	0.0260	0.0018	0.9509
11	0.0001	0.0302	0.0018	0.9432
12	0.0001	0.0285	0.0018	0.9464
13	0.0001	0.0293	0.0018	0.9451
14	0.0001	0.0285	0.0019	0.9468
15	0.0001	0.0291	0.0019	0.9460
16	0.0001	0.0231	0.0019	0.9574
17	0.0001	0.0268	0.0019	0.9506
18	0.0002	0.0348	0.0019	0.9363
19	0.0002	0.0345	0.0019	0.9371
20	0.0001	0.0330	0.0019	0.9401
21	0.0002	0.0374	0.0019	0.9326
22	0.0001	0.0290	0.0019	0.9478
23	0.0001	0.0214	0.0019	0.9614
24	0.0001	0.0148	0.0019	0.9734




Figure 12: Average effect of cumulative mining-induced air pollution exposure on NDVI - growing vs. non-growing seasons

Table 8: Heterogeneous Treatment Effects - Economic Factors

	Base	Above Med. Mineral Rents	Above Med. GDP - 2000	Above Med. Pop - 2000
Near \times Post	-0.00638*** (0.000980)	-0.00534*** (0.00108)	-0.00692*** (0.00157)	-0.00709*** (0.00160)
Near \times Post \times Z		-0.00568** (0.00242)	0.00107 (0.00195)	0.00142 (0.00195)
Number of mines	307	307	307	307
Mean NDVI (t-1)	.476	.476	.476	.476

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Heterogeneous Treatment Effects - Mine Characteristics

	Base	Open Pit	Gold	Coal	Near Other Mine	Near Town	Old Mine
Near \times Post	-0.00638*** (0.000980)	-0.00615*** (0.00174)	-0.00639*** (0.00133)	-0.00675*** (0.00113)	-0.00416*** (0.00123)	-0.00611*** (0.00150)	-0.00615*** (0.00159)
Near \times Post \times Z		-0.000376 (0.00210)	0.00000921 (0.00178)	0.00247 (0.00151)	-0.00449** (0.00194)	-0.000573 (0.00193)	-0.000362 (0.00201)
Number of mines	307	307	307	307	307	307	307
Mean NDVI (t-1)	.476	.476	.476	.476	.476	.476	.476

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Heterogeneous Treatment Effects - Environmental Conditions

	Base	Above Med. NDVI - 2000	Above Med. Precipitation - 2001	Growing Season
Near \times Post	-0.00638*** (0.000980)	-0.00526*** (0.00152)	-0.00762*** (0.00163)	-0.00560*** (0.00113)
Near \times Post \times Z		-0.00226 (0.00195)	0.00247 (0.00194)	-0.00162 (0.00114)
Number of mines	307	307	307	307
Mean NDVI (t-1)	.476	.476	.476	.476

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

10.9 Heterogeneity - Near vs. Far

10.10 Heterogeneity - Contemporaneous Air Pollution

Table 11

	Base	EITI	Control of Corruption	Govt. Effectiveness	Rule of Law	Reg. Quali
Wind	-0.00123 (0.00224)	-0.00261 (0.00412)	-0.00190 (0.00333)	-0.000730 (0.00382)	-0.00280 (0.00337)	-0.000526 (0.00375)
Wind \times (Post - 3)	0.000277 (0.00146)	0.000621 (0.00202)	-0.00240 (0.00189)	-0.00122 (0.00166)	-0.000435 (0.00188)	-0.00169 (0.00172)
Wind \times (Post - 3) \times Z		-0.000554 (0.00284)	0.00402 (0.00267)	0.00246 (0.00271)	0.00106 (0.00271)	0.00321 (0.00268)
Number of mines	102	102	102	102	102	102
Mean NDVI (t-1)	.492	.492	.492	.492	.492	.492

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 12

	Base	Above Med. Mineral Rents	Above Med. GDP - 2000	Above Med. Pop - 2000
Wind	-0.00123 (0.00224)	-0.00283 (0.00265)	-0.00102 (0.00282)	-0.000486 (0.00341)
Wind \times (Post - 3)	0.000277 (0.00146)	0.0000825 (0.00175)	0.000147 (0.00193)	0.00193 (0.00209)
Wind \times (Post - 3) \times Z		0.000671 (0.00274)	0.000372 (0.00285)	-0.00370 (0.00274)
Number of mines	102	102	102	102
Mean NDVI (t-1)	.492	.492	.492	.492

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 13

	Base	Open Pit	Gold	Coal	Copper	Near Other Mine	Near Town	O
Wind	-0.00123 (0.00224)	-0.00688 (0.00555)	0.00243 (0.00333)	-0.00407* (0.00235)	-0.00174 (0.00256)	-0.00227 (0.00328)	-0.00156 (0.00329)	-0
Wind \times (Post - 3)	0.000277 (0.00146)	0.00779** (0.00300)	0.00125 (0.00204)	0.000267 (0.00164)	0.00104 (0.00143)	0.00229 (0.00186)	-0.0000280 (0.00217)	0
Wind \times (Post - 3) \times Z		-0.0104*** (0.00338)	-0.00252 (0.00253)	-0.000114 (0.00369)	-0.00381 (0.00439)	-0.00412 (0.00281)	0.000698 (0.00278)	
Number of mines	102	102	102	102	102	102	102	
Mean NDVI (t-1)	.492	.492	.492	.492	.492	.492	.492	

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 14

	Base	Above Med. AOD - 2003	Above Med. NDVI - 2003	Above Med. Precipitation -
Wind	-0.00123 (0.00224)	0.000823 (0.00463)	-0.00151 (0.00338)	-0.00132 (0.00461)
Wind \times (Post - 3)	0.000277 (0.00146)	-0.000160 (0.00151)	-0.00162 (0.00160)	0.000426 (0.00216)
Wind \times (Post - 3) \times Z		0.000682 (0.00265)	0.00298 (0.00258)	-0.000231 (0.00292)
Number of mines	102	102	102	102
Mean NDVI (t-1)	.492	.492	.492	.492

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$




Figure 13: Heterogeneity by control of corruption

10.11 Heterogeneity - Distributed Lag




Figure 14: Heterogeneity by rule of law




Figure 15: Heterogeneity by government effectiveness




Figure 16: Heterogeneity by regulatory quality




Figure 17: Heterogeneity by voice and accountability




Figure 18: Heterogeneity by mine type




Figure 19: Heterogeneity by mine age




Figure 20: Heterogeneity by distance to nearest mine




Figure 21: Heterogeneity by distance to nearest town




Figure 22: Heterogeneity by commodity: gold




Figure 23: Heterogeneity by commodity: coal




Figure 24: Heterogeneity by commodity: copper




Figure 25: Heterogeneity by mineral rents as share of GDP - 2000




Figure 26: Heterogeneity by GDP - 2000




Figure 27: Heterogeneity by population - 2000




Figure 28: Heterogeneity by initial pollution levels (AOD 2003)




Figure 29: Heterogeneity by initial yields (NDVI 2001)

Effect of Cumulative Air Pollution Exposure on NDVI

Number of Months Prior	$\sum_p^P \delta_p$	$\frac{\sum_p^P \delta_p}{\text{Mean NDVI}}$	Standard error	p
1	0.0003	0.0656	0.0011	0.7956
2	0.0006	0.1307	0.0012	0.6483
3	0.0005	0.1263	0.0014	0.6890
4	0.0007	0.1501	0.0015	0.6616
5	0.0007	0.1581	0.0016	0.6672
6	0.0005	0.1240	0.0017	0.7488
7	0.0004	0.0869	0.0018	0.8297
8	0.0002	0.0405	0.0018	0.9220
9	0.0001	0.0222	0.0018	0.9579
10	0.0001	0.0260	0.0018	0.9509
11	0.0001	0.0302	0.0018	0.9432
12	0.0001	0.0285	0.0018	0.9464
13	0.0001	0.0293	0.0018	0.9451
14	0.0001	0.0285	0.0019	0.9468
15	0.0001	0.0291	0.0019	0.9460
16	0.0001	0.0231	0.0019	0.9574
17	0.0001	0.0268	0.0019	0.9506
18	0.0002	0.0348	0.0019	0.9363
19	0.0002	0.0345	0.0019	0.9371
20	0.0001	0.0330	0.0019	0.9401
21	0.0002	0.0374	0.0019	0.9326
22	0.0001	0.0290	0.0019	0.9478
23	0.0001	0.0214	0.0019	0.9614
24	0.0001	0.0148	0.0019	0.9734

Table 16

	Base	EITI	Control of Corruption	Govt. Effectiveness	Rule of Law	Reg. Qua
Downstream \times Post	-0.0140** (0.00587)	-0.0129 (0.00911)	-0.0131** (0.00600)	-0.0164* (0.00840)	-0.0126** (0.00487)	-0.0117 (0.0065)
Downstream \times Post \times Z		-0.00203 (0.0119)	-0.00185 (0.0119)	0.00527 (0.0115)	-0.00330 (0.0126)	-0.0051 (0.0121)
Number of mines	38	38	38	38	38	38
Mean NDVI (t-1)	.467	.467	.467	.467	.467	.467

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 17

	Base	Open Pit	Gold	Coal	Near Other Mine	Near Town	Old Mine
Downstream \times Post	-0.0140** (0.00587)	-0.0162 (0.0109)	-0.00849 (0.00511)	-0.0132* (0.00706)	-0.000109 (0.00625)	-0.0189*** (0.00687)	-0.0234** (0.0100)
Downstream \times Post \times Z		0.00333 (0.0129)	-0.0166 (0.0147)	-0.00432 (0.00961)	-0.0285** (0.0108)	0.00937 (0.0115)	0.0151 (0.0122)
Number of mines	38	38	38	38	38	38	38
Mean NDVI (t-1)	.467	.467	.467	.467	.467	.467	.467

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 18

	Base	Above Med. Mineral Rents	Above Med. GDP - 2000	Above Med. Pop - 2000
Downstream \times Post	-0.0140** (0.00587)	-0.0180*** (0.00553)	-0.0121 (0.00948)	-0.0193** (0.00878)
Downstream \times Post \times Z		0.0129 (0.0151)	-0.00406 (0.0115)	0.0108 (0.0115)
Number of mines	38	38	38	38
Mean NDVI (t-1)	.467	.467	.467	.467

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

10.12 Heterogeneity - Water

Table 19

	Base	Above Median Turbidity - 2000	Above Med. NDVI - 2000	Above Med. Precipi
Downstream \times Post	-0.0140** (0.00587)	-0.0252*** (0.00856)	-0.0107** (0.00450)	-0.0214* (0.00975)
Downstream \times Post \times Z		0.0199* (0.0113)	-0.00633 (0.0112)	0.0146 (0.0116)
Number of mines	38	38	38	38
Mean NDVI (t-1)	.467	.467	.467	.467

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

10.13 Structural Breaks Model For each mine, I construct total luminosity for each year from 1992-2012 as the sum of nighttime light intensity across all pixels that fall within the 20-km disk around the mine centroid, in a given year. The nightlight data is provided by the DMSP OLS at a 1-km resolution. However, the DMSP-OLS data is not available to the public after 2013. From 2013 onwards, there was a switch to the Visible Infrared Imaging Radiometer Suite (VIIRS) instrument as the source of nightlights. Given that the spatial and radiometric resolution from VIIRS is higher than DMSP-OLS, Li et al. (2020) created a temporally harmonized global nighttime lights dataset from 1992-2019. Unfortunately, the temporal consistency of their harmonized data does not perform well in areas with low levels of luminosity (pixels with DN values greater than 7). Time series of nighttime light intensity for mining areas reveal a sharp jump in luminosity in 2013, when there was a switch to VIIRS. Given that the average pixel in a mining area has a DN value of 6.5, this suggests that the harmonized data is not temporally consistent for mining areas. As a result, I only estimate a structural break in nightlights for mines that opened before 2013, to avoid confounding due to the switch from DMSP to VIIRS.

I use the methods of Andrews (1993) and Andrews and Ploberger (1994) to detect a structural break in the mean of nightlights for each mine. Given a single proposed break point, the nightlights time series is split into two bins. The estimation method fits a regression of nightlights on an intercept for the data in each bin and calculates an F statistic based on the null hypothesis that the mean in nightlights is the same between the two bins. This step is repeated for all possible break points in the data to yield a time series of F-statistics. The structural break is identified as the year which yields the largest F-statistic. To test whether this structural break is statistically significant, I use the critical values for the F-statistics identified by Andrews (1993) and Andrews and Ploberger (1994), which are computed under the null hypothesis of no structural change such that the asymptotic probability that the supremum of the time series of F-statistics exceeds this critical value is $\alpha = 0.05$. In other words, if the F-statistic of the structural break is higher than the critical value, we can reject the null of no structural breaks in nightlights at the 0.05 level.

I am successfully able to estimate a structural break for 100 mines in my sample that opened prior to 2013. For the median mine, a structural break in nightlights occurs 3 years prior to the listed S&P start date, with approximately 60% of the mines having the break year between 0 and 3 years prior to the S&P defined date. This finding aligns with that of Benschaul-Tolonen [2020], who uses a local polynomial regression to show that there is a break in the trend for nightlights within 10 kilometers of a mine that occurs roughly 2 years prior to the date of a mine opening. She refers to this break as the start of the “investment” phase in a mine’s life cycle. Figure 30 plots the histogram of the difference between the S&P year and the break year.




Figure 30: Difference between S&P Year and Structural Break Year

Furthermore, for most mines the year identified as the structural break visually aligns with the timing of a sharp increase in nighttime lights. Figure 31 provides an example from Goedgevonden coal mine in South Africa.

To investigate the robustness of this method for detecting structural breaks, I examine the correlation between the F-statistic from the break year and the start-date “error”, defined as the difference between the S&P defined year and the break year. One might be concerned if the method is more likely to reject the null of no structural break in cases where there is a large difference between the S&P defined year and the break year, suggesting that it would not perform well in detecting breaks that occur close to the listed date.

These concerns do not seem to be borne out: in Figure 32 we see that across the domain of the “error”, the sup F statistic ranges between 0-150 in most cases. In fact, the few outliers with very high supF statistics occur for cases where the S&P year and the break year are very close.




Figure 31: Time series of nightlights for Goedgevonden coal mine, South Africa




Figure 32: Magnitude of supF statistic vs. the difference between S&P defined start year and break year

10.14 Linking NDVI to Actual Yields




Figure 33: All plots pooled, growing season only




Figure 34: Aggregate to means by cell-year, growing season only

10.15 Scatterplots of actual yields vs. NDVI

10.16 Effects of Measurement Error on DID Estimates: Ideally, we would be able to observe farm-level yields around mining areas over time. However, this data is not available in Sub-Saharan Africa, making the use of NDVI as a proxy for yields more appealing. As NDVI is an imperfect measure of yields, I specify the following linear measurement error model (Keogh et al., 2020) to help understand the types of measurement error present in NDVI:

$$NDVI_{smt} = \theta + \mu Yield_{*smt} + v_{smt} \quad (8)$$

where NDVI on side s of mine m at time t is observed, while total yields across all farms on side s of mine m at time t is the true measure of interest and v_{smt} is noise. I assume v to be mean 0 and $cov(Yield_{*}, v) = 0$. With classical measurement error, assuming v_{smt} is uncorrelated with downwind or downstream exposure and uncorrelated with the timing of a mine opening, estimates of the DID regressions using NDVI as the outcome of interest will be unbiased, though standard errors will be larger. Classical measurement error corresponds to the special case of 8 where $\theta = 0$, $\mu = 1$ and $cov(Yield, v) = cov(X, v) = 0$, where X corresponds to the main independent variables of my DID specification. However, Proctor et al. (2023) show that assumptions of classical measurement error are unlikely to hold in regressions with remotely sensed dependent variables. First, they demonstrate that mean-reverting measurement error (negative correlation between errors in one variable and itself) is common with remotely-sensed measures, by highlighting that extreme values are systematically underestimated in remotely sensed predictions (Bound and Krueger, 1991, Ratledge et al., 2022). Mean-reverting measurement error corresponds to the case of Equation 8 where $\mu < 1$. Second, they show that differential measurement error, correlations between errors in one variable and levels of another variable, can also occur, by demonstrating non-zero covariance between errors in one variable and levels of other variables.

Proctor et al. (2023) suggests that the case with only mean-reverting error and no differential measurement error, a model using uncorrected NDVI with measurement error will tend to have attenuated coefficients. However, with differential measurement error, biases of the coefficient can arise in either direction. Furthermore, they show that while mean-reverting measurement error is most common with remotely-sensed variables, differential measurement error explains most of the coefficient bias in the case of a mismeasured dependent variable. Importantly, both mean-reverting and differential measurement error can also lead to bias in standard errors of coefficients. In general, the reduced variance in the outcome variable in the presence of mean-reverting measurement error in the dependent variable could lower the sum of squared errors and underestimate standard errors on the DID coefficients of my main specification. However, when differential measurement error is also present, the direction of bias in the uncertainty parameters is theoretically

ambiguous (Carroll et al., 2006).

Unfortunately, it is infeasible in my setting to implement the multiple imputation correction outlined by Proctor et al. (2023) to address issues of measurement error bias in NDVI. This is because potential data sources for ground-based yields (One Acre Fund data, World Bank LSMS-ISA) have an insufficient number of observations that are observed for at least one year pre and post-mine openings, in both the near and far groups. Proctor et al. (2023) highlight that the calibration set should have at least 500 observations, otherwise bias can be even higher in the corrected model. Furthermore, while plot-level yields from data outside of Sub-Saharan Africa could be used for the calibration set, bias in the corrected model can be higher than in the uncorrected model if calibration locations are greater than 200km away from the locations used in the main analysis (Proctor et al., 2023).

10.17 Effects of Measurement Error on Quantification Exercise: Given plot-level yields, the “ideal” measure of NDVI for each plot, $NDVI_p^*$, would be NDVI based on information detected by the satellite without any errors or interference and calculated over the exact boundaries of plot p . However, the NDVI measure used in Equation 7 differs from $NDVI^*$ in two ways. Dropping the t subscript for simplicity, I define the observed NDVI measure for plot p , which falls in cell c , as follows:

$$NDVI_{p(c)} = \xi NDVI_p^* + d_p + a_p$$

where d_p corresponds to deviations from “true” plot-level NDVI caused by satellite-detection errors and a_p corresponds to deviations for plot p NDVI from the average of all plots in c . NDVI detection errors, d_p , are driven by a variety of factors, most notably cloud cover, but also surface reflectance, canopy thickness, the level of atmospheric aerosols and sensor errors (MODIS article). These detection errors make NDVI observed from satellite data an imperfect measure of NDVI in the presence of no detection errors. Indeed, low pixel quality due to detection errors can cause errors in vegetation indices to increase by 0.04-0.1. While some types of detection errors, such as satellite sensor errors, are likely uncorrelated with plot-level yields, detection errors driven by cloud cover or aerosol loading could be correlated with plot-level yields through the direct effect of atmospheric determinants on crop health. In these cases, we would have $Cov(Yield_p, d_p) \neq 0$, suggesting that d_p could introduce bias from non-classical measurement error.

Additionally, aggregation errors, a_p , are generated by the limited 500m resolution of the MODIS data, which makes it infeasible to obtain precise plot-level NDVI measures.¹⁰ GPS-based measures of plot area from household surveys in four African countries show that over 50% of the fields in these countries are

¹⁰While higher resolution data is available from sources like Sentinel-2 or PlanetLabs, this data does not go far enough back to cover mines that opened prior to 2014.

smaller than 1 acre, suggesting that the grid-cells used in my analysis likely cover multiple plots (Carletto et al., 2015). If yields are highly heterogeneous across plots within a grid-cell, averaging NDVI over the grid cells may not only weaken the correlation between grid-cell NDVI and plot-level yields, but also introduce aggregation bias into estimation of Equation 7.

I address measurement error in cell-level NDVI as follows. First, I limit potential bias from detection errors by calculating NDVI using only high quality pixels free of clouds and aerosols, as well as controlling for local weather variables such as cloud cover, vapor pressure and precipitation. Second, I estimate an instrumental variables version of Equation 7, shown in Table ???. I instrument for mean NDVI at the cell-year level calculated over the even days in a given season with the same measure calculated over the odd days in the season. By instrumenting one noisy measure of NDVI with another noisy measure, I can ideally purge noise that is common to both. Third, I estimate (1) a version of Equation 7 with smaller grid cells to investigate whether β_{NDVI} changes with the degree of aggregation and (2) a version of Equation 7 that aggregates plot-level yields to the grid-cell used to aggregate NDVI, so that spatial scales are consistent between the dependent and independent variables.

Table 20: Relationship between Cell-level NDVI and Cell-level Yields, cell-size = 500m

	(1)	(2)	(3)	(4)	(5)
	Planting	Early Growing	Late Growing	Harvest	Nonfarm
Mean NDVI	45.68 (32.85)	38.76 (25.83)	40.76 (33.81)	30.06 (22.08)	7.269 (25.08)
Year FE	Yes	Yes	Yes	Yes	Yes
Grid cell FE	Yes	Yes	Yes	Yes	Yes
Include weather controls	Yes	Yes	Yes	Yes	Yes
Mean Yields	3005.643	3005.643	3005.643	3005.643	3005.643
Obs.	340	467	421	524	445
R-sq	0.657	0.673	0.673	0.672	0.650

Each column reports the results of an OLS regression. The unit of observation is a cell-year. The dependent variable is mean plot-level yields in kg/ha across all plots falling within a cell, in a given year. *Mean NDVI* is the average NDVI in the cell of size 500m containing the plot, across days in a particular season during the year that maize on the plot was harvested. The columns indicate the season for which NDVI is calculated during the year the plot was harvested. *Mean NDVI* is scaled so that one unit represents a 0.01 increase in NDVI. Each regression includes linear controls for cell-level averages of growing degree days, temperature and precipitation during the maize season across plots falling within the cell, as well as cell-level controls for mean temperature, precipitation, vapor pressure, cloud cover, evapotranspiration and wet days. Year and cell-level fixed effects are included in the regressions.

Table 21: Relationship between Cell-level NDVI and Cell-level Yields, cell-size = 2500m

	(1)	(2)	(3)	(4)	(5)
	Planting	Early Growing	Late Growing	Harvest	Nonfarm
Mean NDVI	21.96* (12.06)	41.64** (18.23)	0.756 (16.61)	29.74** (12.99)	27.23** (10.77)
Year FE	Yes	Yes	Yes	Yes	Yes
Grid cell FE	Yes	Yes	Yes	Yes	Yes
Include weather controls	Yes	Yes	Yes	Yes	Yes
Mean Yields	3005.643	3005.643	3005.643	3005.643	3005.643
Obs.	2038	2070	2071	2057	2069
R-sq	0.588	0.585	0.583	0.585	0.586

Each column reports the results of an OLS regression. The unit of observation is a cell-year. The dependent variable is mean plot-level yields in kg/ha across all plots falling within a cell, in a given year. *Mean NDVI* is the average NDVI in the cell of size 2500m containing the plot, across days in a particular season during the year that maize on the plot was harvested. The columns indicate the season for which NDVI is calculated during the year the plot was harvested. *Mean NDVI* is scaled so that one unit represents a 0.01 increase in NDVI. Each regression includes linear controls for cell-level averages of growing degree days, temperature and precipitation during the maize season across plots falling within the cell, as well as cell-level controls for mean temperature, precipitation, vapor pressure, cloud cover, evapotranspiration and wet days. Year and cell-level fixed effects are included in the regressions.

10.18 Addressing measurement error

Table 22: Relationship between Cell-level NDVI and Cell-level Yields, cell-size = 5000m

	(1)	(2)	(3)	(4)	(5)
	Planting	Early Growing	Late Growing	Harvest	Nonfarm
Mean NDVI	11.93 (11.11)	54.57*** (15.96)	18.00 (14.47)	46.21*** (11.52)	34.77*** (10.03)
Year FE	Yes	Yes	Yes	Yes	Yes
Grid cell FE	Yes	Yes	Yes	Yes	Yes
Include weather controls	Yes	Yes	Yes	Yes	Yes
Mean Yields	3005.643	3005.643	3005.643	3005.643	3005.643
Obs.	2016	2040	2038	2036	2040
R-sq	0.551	0.559	0.552	0.557	0.553

Each column reports the results of an OLS regression. The unit of observation is a cell-year. The dependent variable is mean plot-level yields in kg/ha across all plots falling within a cell, in a given year. *Mean NDVI* is the average NDVI in the cell of size 5000m containing the plot, across days in a particular season during the year that maize on the plot was harvested. The columns indicate the season for which NDVI is calculated during the year the plot was harvested. *Mean NDVI* is scaled so that one unit represents a 0.01 increase in NDVI. Each regression includes linear controls for cell-level averages of growing degree days, temperature and precipitation during the maize season across plots falling within the cell, as well as cell-level controls for mean temperature, precipitation, vapor pressure, cloud cover, evapotranspiration and wet days. Year and cell-level fixed effects are included in the regressions.

10.19 Air Pollution - Instrumental Variables

10.19.1 Reduced form

$$NDVI_{smt} = \alpha_{sm} + \lambda_{mt} + \delta Wind_{smt} + \eta Post_{mt} + \beta_A Wind_{smt} \times Post_{mt} + \mathbf{X}'_{smt} \boldsymbol{\Gamma} + \epsilon_{smt} \quad (9)$$

10.19.2 First stage

$$AOD_{smt} = \alpha_{sm} + \lambda_{mt} + \delta Wind_{smt} + \eta Post_{mt} + \beta_A Wind_{smt} \times Post_{mt} + \mathbf{X}'_{smt} \boldsymbol{\Gamma} + \epsilon_{smt} \quad (10)$$

10.19.3 Second stage

$$NDVI_{smt} = \theta AOD_{smt} + \alpha_{sm} + \lambda_{mt} + \mathbf{X}'_{smt} \boldsymbol{\Gamma} + \epsilon_{smt} \quad (11)$$

10.20 Mine-specific DID

Table 23: Effect of Mining Air Pollution on NDVI - IV

	(1)	(2)	(3)	(4)
	No window restriction, post opening	No window restriction, post investment	5 pre, 5 post, post opening	5 pre, 5 post, post investment
Mean AOD	-0.475** (0.195)	-0.477** (0.195)	-0.567 (0.605)	-0.327 (0.410)
First stage F	15.94	15.65	2.18	3.96
Num mines	533	533	101	101
Num obs.	378376	378376	71888	71888

Each column reports the results from an instrumental variables regression. The unit of analysis is a mine-side-month. I instrument for mean AOD with the share of days a side is downwind from a mine in a given month as well as this share interacted with the post dummy. Both the first and second stage include mine-side and mine-year-month fixed effects, as well as linear and quadratic controls for mean temperature, precipitation, evapotranspiration, wet days, and cloud cover. Standard errors in parentheses are clustered at the mine-level. Column 1 reports the results for the IV regression estimated over all mines that have non-missing AOD and NDVI data on all four sides, for at least 6 months of every year, for each year from 2003-2017. Column 2 reports the results estimated over the same sample as before but replaces the post dummy with a dummy equal to 1 for all years after the investment phase, which is defined as occurring 3 years before the mine opens. Column 3 reports the results estimated over the sample of mines that have non-missing AOD and NDVI data on all four sides, for at least 6 months of every year, and are observed 5 years pre- and post-mine opening. Column 4 reports the results estimated over the sample as in Column 3, but replaces the post dummy with a dummy equal to 1 for all years after the investment phase.

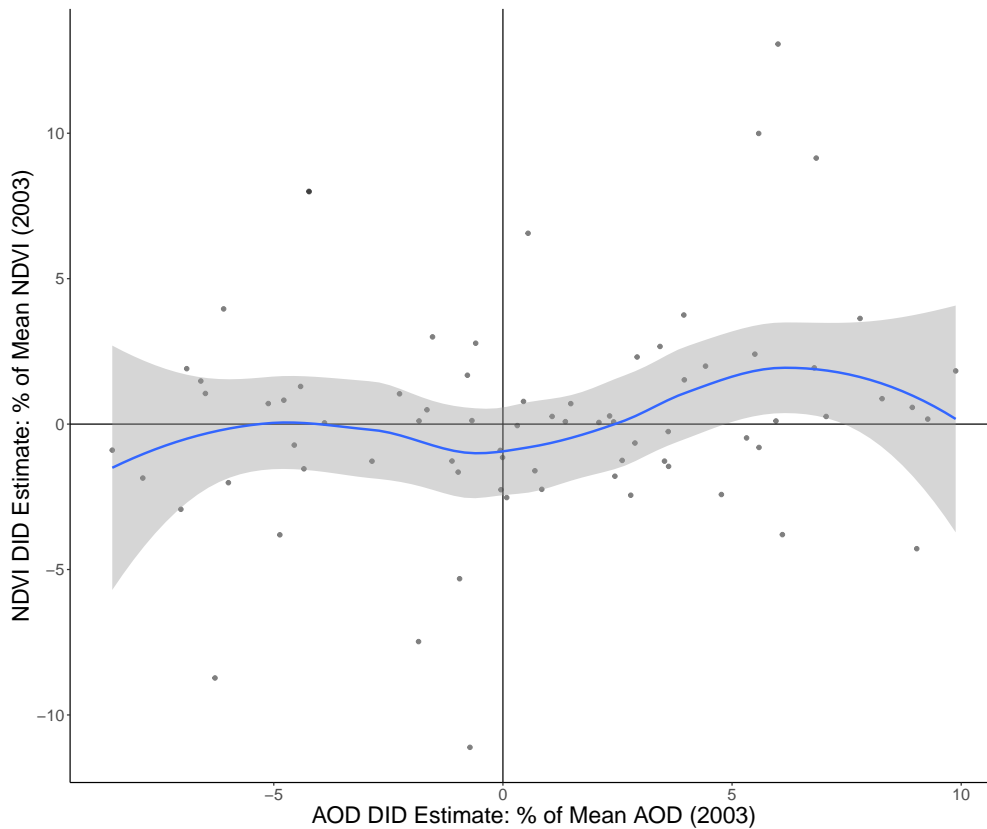


Figure 35: Mine-specific AOD DID estimates vs. NDVI DID estimates




Figure 36: Most Important Variables from Best Random Forest Model: Classification