Administrative Incentives Impact Crop-Residue Burning and Health in South Asia *

Gemma Dipoppa† Saad Gulzar‡

November 21, 2022

Abstract

Air pollution in South Asia is one of the largest public health emergencies on the planet, responsible for a million deaths annually. A third of the pollution is caused by crop-residue burning (CRB). While solutions have been proposed to reduce CRB, implementing any of them at scale requires government intervention. Here we examine how administrators' incentives shape the effective management of pollution in India and Pakistan. We causally identify variation in administrators' incentives to manage CRB by leveraging a decade of wind, fire, and health data from satellites and DHS surveys spanning over 17 million grid cells. Results show that CRB decreases when smoke pollutes the home district of the administrator, and increases when the externality is borne by neighboring jurisdictions, with substantial impacts on child mortality. These strategic patterns are five-times larger on the India-Pakistan border, where inter-jurisdictional cooperation is difficult. Our results indicate that administrators' accountability to their own constituents and lack of accountability to others citizens impact polluting behavior. Taken together, these findings offer hope that policies aimed at curbing environmental damages can substantially improve their effectiveness by incorporating incentives faced by administrators responsible for implementing policy solutions at scale.

*We thank Ali Shah, Anna Hudson, Diego Tocre, Annalisa Pezone and Max Vandervelden for outstanding research assistance. We also thank Fatiq Nadeem for helpful discussions and suggestions. Seminar and conference participants at Society for Institutional Organizational Economics Conference, Petralia Political Economy Conference, Princeton Development Seminar, Rochester, Columbia, Tweeds, UCSC, UPenn Environment Conference, and UPPER conference at SAIS provided valuable feedback and comments. This study is funded by grants from Stanford UPS Endowment Fund, Stanford Center for Innovation in Global Health, and the Stanford Sean N Parker Center for Allergy and Asthma Research.

†Department of Political Science, Brown University. gemma_dipoppa@brown.edu

‡Department of Politics and School of Public & International Affairs, Princeton University. gulzar@princeton.edu

1 Introduction

South Asia is home to one of the world's largest public health emergencies caused by air pollution. One in four humans on the planet live in this region, which has 42 of the 50 with the worst air pollution in the world (RCoHP South Asia, 2021): see Lahore, Pakistan in Figure 1. Bad air quality causes economic losses (Lu, 2020), increases the rate of climate change (Knippertz et al., 2015), and has severe health impacts on the hundreds of millions living in the region. By some estimates, bad air quality causes up to a million deaths every year in South Asia (Cusworth et al., 2018).[1]

Figure 1: Air Pollution in Lahore



About a third of air pollution in South Asia is caused by emissions from burning the residue that is left after crops like rice are harvested (often called "crop residue burning") (Abdurrahman, Chaki, and Saini, 2020). Curbing this practice is *the* archetypal environmental challenge: it is individually rational and relatively cost-free for farmers to burn stubble as the air pollution from burning is borne by those who are downwind from the field, and alternatives to burning, like mechanized harvesters, are expensive and not easily available (Singh, 2018).

Studies suggest that paying farmers to not burn (Jack et al., 2022), and shifting the timing of burning may help to limit the pollution impacts of crop-residue burning (Lan et al., 2022). However, the successful implementation of all policies at the vast scale at which crop-residue burning is experienced will require government intervention.

States in both India and Pakistan have adopted regulations that make stubble burning illegal, punishable by fines and jail terms. Unfortunately, the implementation of these regulations has been uneven for at least two reasons: first, it is difficult to monitor and sanction

[1]The state of the health emergency can be gauged by comparing the problem with Covid-19 that caused 6 million deaths across the planet in the first two years

thousands of dispersed farmers who are each individually deciding to pollute or not; and second, administrators who have the power to curb polluting behavior themselves are typically overloaded with work.

This paper examines state administrator incentives to both allow and curb crop-residue burning in the context of India and Pakistan. We argue that administrators curb residue burning when the practice pollutes their own jurisdiction. Using changes in wind direction, we estimate a model that shows that indeed the incidence of fires in satellite data between 2012-2021 (VIIRS and MODIS) decreases by 10-14 percent right after wind switches move a location from polluting mostly the neighboring jurisdiction to an administrator's own jurisdiction. Given that rice is the primary crop associated with residue burning that impacts air quality, we show that the reduction in polluting behavior is about four times as large for areas that produce rice and during the months in which rice is harvested.

Next, we present further evidence that patterns in the data conform to administrators' incentives to curb crop residue burning. First, we show that fires increase the closer we move towards district borders, but only when those are downwind, meaning that the pollution externality exits the home districts quickly. Instead, fires drop monotonically as we move closer to the upwind border, an area that would expose most of the home district to pollution. Our interpretation of these results is that moving closer to the border intensifies the externality incentives to allow or curb polluting behavior. Second, we show that as the size of the polluting problem increases – as proxied by the proportion of the district that cultivates rice – the strategic response of administrators increases, as before, in the case of both the upwind and the downwind border. Third, we show that the effect increases five-fold on the India-Pakistan border, where the ability to coordinate with neighboring bureaucrats should be absent. Finally, we show that the strategic reduction in fires is conditional on administrators' capacity and ability to observe and curb fires.

The final question we ask is whether this behavior impacts health outcomes. We calculate the non-linear trajectory that fires particles take across space using a meteorological model (HYSPLIT)[2] and combine these trajectories with DHS geo-localized data to estimate how many fires each child is exposed to in-utero. We estimate the impact of fires on child mortality holding fixed the location and time of birth, and comparing mortality outcomes for children born to the same mother but with different in-utero exposures. We document a large increase in child mortality and conduct a back-of-the-envelope calculation that suggests that if administrators were able to reduce residue burning such that all areas in their jurisdiction behaved the way those that pollute the jurisdiction do, mortality effects would decrease by about 25 percent with respect to the mean.

[2] `https://www.ready.noaa.gov/HYSPLIT.php`

Globally, scholars have started to document the perverse effects of poor air on health, productivity, and climate change in general (Deryugina et al., 2019; Graff Zivin and Neidell, 2012; Heft-Neal et al., 2019). However, nowhere is the problem of air pollution worse than in South Asia. To our knowledge, ours is the first systematic analysis of the political economy determinants and consequences of crop burning, combining evidence from a region where 500 million people live, and whose decisions potentially impact the air pollution of over 1.5 billion individuals.[3]

Understanding when public administrators exert control to curb pollution, or conversely, withdraw effort to allow polluting behavior, is important because these are the actors who will implement any environmental regulations at scale. Indeed, an examination of the political economy determinants of effective environmental management is as an open research question according to a recent review of the literature (Greenstone, Reguant, et al., 2021, p. 48). We join recent literature that examines how jurisdictional incentives affect pollution (Kahn, Li, and Zhao, 2015; Lipscomb and Mobarak, 2016; Morehouse and Rubin, 2021).

2 The Management of Air Pollution in South Asia

High levels of PM particles in the air have been shown to cause strokes, diabetes and respiratory diseases, leading to lower work productivity and educational attainment as well as to more hospitalizations and premature deaths (Deryugina et al., 2019; Lelieveld et al., 2015). The cost of poor air for governments and society is immense and estimated at 2.9 trillion US dollars, corresponding to 3.3 percent of the global GDP (Myllyvirta, 2020).

This burden is not equally distributed across countries: levels of particulate matter in the air are highest in the developing world, particularly in South Asia (Figure 2).

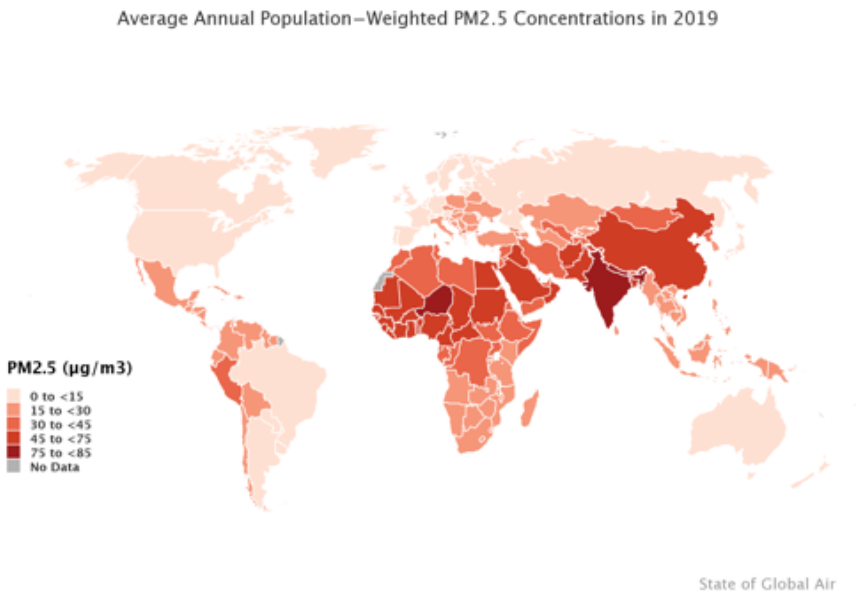
2.1 Crop Burning in South Asia

The largest contributor to PM2.5 is agricultural emissions (Lelieveld et al., 2015), and in particular crop residue burning. Used as a fast and inexpensive method for clearing plantations before the next cycle, crop burning has costs that far outweigh its convenience. Farmers echo this sentiment: "If I can clear my farm using a one-rupee matchbox, why will I spend thousands?" (Jitendra et al., 2017)

In India and Pakistan, this practice is used extensively. Several factors have contributed to the severity of the problem. First, the Green Revolution enabled farmers in many areas

[3]See also Greenstone and Hanna (2014) for an assessment of the effect of air pollution regulations on mortality in South Asia and Jayachandran, 2009 for an analysis of how wildfires impact early life mortality.

Figure 2: Population-weighted average PM2.5 concentrations, 2019 (Source: State of Global Air)



in the northern Indian sub-continent to cultivate multiple crops every year. This practice is called the 'rice-wheat cropping system' (Singh, 2018, p. 150), and is characterized by rapid harvest and sowing cycles so that crops can catch prime conditions for their growth. Given that both India and Pakistan have a large agricultural sector that focuses on rice and wheat production – which generate the largest crop residues (NPMCR, 2014) – the problem of stubble burning is acute. Second, crop burning has increased with mechanization technologies, which produce larger crop residues and increase incentives to use fire to clear the fields (Maue, 2021). An increase of labor wage – due to out-migration from rural areas among other factors over the last few decades – has further increased the use of these mechanized tools.[4]

2.2 The Role of the State in Curbing Crop Burning

In India, burning crop residue is a crime since 1981,[5] but states are ultimately responsible for "bringing suitable law, rules, or orders for prevention or banning of the practice of crop residue burning as per policy and priority of the State".[6] The states of Punjab, Haryana and Uttar Pradesh adopted laws to ban this practice in 2015 and the Supreme Court prohibited

[4]See Singh, 2018, p. 156-164 for more details.
[5]See section 188 of the Indian Penal Code and the Air and Pollution Control Act of 1981.
[6]National Policy for Management of Crop Residue, 2014.

the setting of fires in these states in 2019[7]. The practice is also illegal in the Pakistani Punjab since 2018.[8]

While states adopt rules against crop burning, their enforcement, the monitoring of fires, and the proposal of alternative methods to handle stubble is delegated to district level bureaucrats. Bureaucrats can first of all enforce already existing legislation punishing crop-burners with fines and even jail time, but they can also adopt other initiatives. Following, are examples from the press of how bureaucrats have handled reducing crop-residue burning in different contexts:[9]

> "A senior bureaucrat known for his tough approach to polluters, Pannu is leading the anti-burning efforts in Punjab. He has mostly carrots, not sticks, at his disposal. For the first time this year, India's central government earmarked money to help farmers buy machinery that turns the straw into mulch. It has also contributed to a large-scale awareness campaign, complete with songs on social media, television advertisements and village-by-village meetings all urging farmers not to burn."
>
> "Gurdial Kumar is one of the foot soldiers in the statewide effort. A 48-year-old employee of the Punjab agriculture department, he and his colleagues began visiting each of the 182 villages in their territory last month. On a recent afternoon, they arrived in Madan Heri, a village of 2,500 people west of Chandigarh, where about 20 farmers had gathered in the shade of a giant banyan tree. Kumar and his three colleagues put up posters illustrating the impact of air pollution, including images of diseased lungs and small children wheezing. They talked about the virtues of the new subsidized machines, including one called the Happy Seeder. They appealed to religious precepts; most of the farmers practice Sikhism, whose sacred text emphasizes the holiness of nature."
>
> " Conversations with farmers indicated that those in the Haryana district shun stubble burning... [because of] strict vigilance by the district administration... "Farmers were afraid of penalties", said Ajmer Singh, a farmer in Jaswanti village in Kaithal. "In our village, at least two farmers were fined for stubble fire. So, the fear of penalty also discouraged farmers from burning". Dr Karam Chand, the deputy director of Kaithal district's agriculture department, told Newslaundry the district administration had collected Rs 9.10 lakh as penalties from 402 farmers in the district. "

[7]Order in IA Nos.158128 and 158129 of 2019 In Writ Petition (C) No.13029 Of 1985.
[8]Section 144 of the Code of Criminal Procedure.
[9](Washington Post, October 2018)

(Newslaundry, 2022)

Notwithstanding increased public concern for air pollution, there continues to be large variation in where law enforcement and public initiatives against crop burning take place, indicating a possible role for bureaucratic discretion.

2.3 Bureaucrats and Farmers' Incentives

Bureaucrats face a trade-off when they decide whether or not to adopt initiatives against crop burning. One one hand, they might receive positive returns in terms of public and higher level bureaucrats' opinion if they can contribute to reducing smoke. On the other hand, they might face costs arising from farmers' complaints, especially if they use fines and jail to punish violators.

We do not expect bureaucrats to exert the same level of effort to reduce fires spatially. Instead, we hypothesize that they will act when the returns from avoiding a fire are highest, which is when a larger share of their district would be exposed to smoke from that fire. This is reasonable because bureaucrats have limited time and resources and must be strategic in where and when they exert effort (Dasgupta and Kapur, 2020). Depending on wind direction, a fire occurring in a certain location might or might not pollute a larger portion of the administrative unit, making the bureaucratic choice on whether to intervene to stop fires time-variant.

Correspondingly, farmers anticipating bureaucratic law enforcement against crop burning might avoid setting fires to clear their fields in the periods and places in which they expect bureaucrats to be most alert. Farmers might also try and burn crop stubble right before the fire season starts to avoid bureaucratic monitoring. In fact, given the limited time span between the harvesting period and the time at which new seeds need to be planted before the monsoon starts, the margin for anticipating bureaucratic action is very small. Additionally, if this became the equilibrium, bureaucrats would also ad just their monitoring behavior. Finally, a possible strategy of charging the higher cost from using non-polluting technologies to consumers is not available to farmers in India and Pakistan, as farmers cannot modify prices individually. The government buys agricultural produce at higher than market rate and only the government can centrally decide on modifying prices.

Based on the incentives bureaucrats and farmers face to reduce crop burning, we elaborate a set of hypotheses on where we should expect to observe stubble burning. Notice that, while in the rest of the paper we refer to "bureaucrats" for simplicity, the behavior we aim to capture arises from the interaction between bureaucratic action to reduce fires and farmers' response to this action or anticipated action to avoid punishment and to receive the benefits



Figure 3: (left) Total number of fires per month. (right) District-average actual and 10 year rolling average wind direction. All data refer to our sample of states in the period 2012-2021.

offered as incentives to adopt alternative clearing methods.

1. We expect a reduction in fires in places and times when smoke from a potential fire would affect more of the district (e.g. when the district area or population is mostly located downwind from the fire).
2. This effect should be larger in rice growing areas and rice harvest months.
3. The effect should increase as the polluting capacity of a location increases – for instance, as the area lies closer to the upwind border of the district.
4. If the effect is driven by bureaucratic action, it should exist where bureaucrats have the capacity to monitor and manage fires occurrence, as well as the incentives to do so.

3 Data

Fires: We gather data on fires presence and brightness from the Nasa satellites Modis (1km x 1km resolution) and Viirs (375m x 375m resolution) from 2012, the first year in which Viirs starts operating. We plot the total number of fires over time in Figure 3: each year, there are two peaks in the months of April-May and October-November, in correspondence with the crop burning season of wheat and rice (Singh and Kaskaoutis, 2014).

Crops: We gather data on average crop area, production and yield in 2004-2006 – the latest years in which micro-level crop data are available before the start of our analysis sample (2012-2021). These data are retrieved from FAO maps, which has a resolution of 5 x 5 arc/min, corresponding to approximately the size of a grid. The crop for which fires are mostly used to clear the field, rice and wheat, are mapped in Figure A.1.

Winds: Data on wind direction and speed is from Ecmwf and is precise at a 25km resolution. We proxy for expected wind direction of bureaucrats in a given month-year by

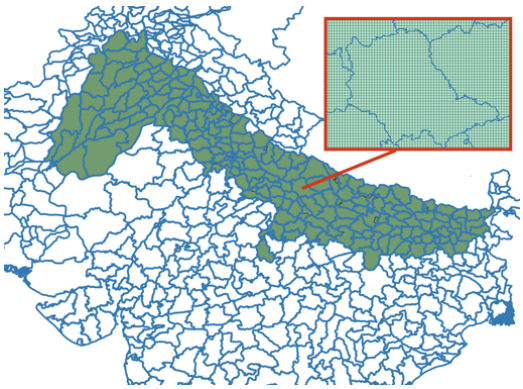


Figure 4: Illustration of our gridded 5 squared km dataset for Hardoi District in Uttar Pradesh

computing ten-year rolling averages of wind direction in each month and year by using wind data starting in the 2000s. We plot actual and rolling averages of wind directions in Figure 3, right panel. We can see that rolling averages follow closely actual wind direction. The two major wind switches happen in June and August in each year.

Administrative borders: We spatially merge satellite data with administrative units in India and Pakistan. We focus on districts as those are the administrations responsible for the enforcement of anti-crop burning laws and initiatives. In addition, our main sample of analyses is states where the problem of crop stubble burning is particularly salient and, consequently, where air pollution is particularly high. These states are Punjab in Pakistan, and Punjab, Haryana, Uttar Pradesh, Delhi and Bihar in India.[10]

Combining all the data at a 5 square km grid level: Wind and fire grids do not overlap precisely: to reduce errors in the assignment of fires to wind grids, we split the 25x25km wind grids into squares of 5 squared km (2.25km x 2.25km size). This guarantees that even if wind and fire grids do not perfectly overlap, data on fires are assigned to the wind grid which contains most of the fire grid. Splitting wind grids in smaller parts also allows us to address a second measurement concern: some grids might overlap across jurisdictions boundaries. By reducing their size, we minimize the amount of misattribution and we are able to assign each grid to the appropriate jurisdiction. Notice that any remaining error should bias our findings towards zero, as errors are uncorrelated with the dependent and independent variable. For crop data, we average the pre-period crop production and assign it to each 5 squared km grid. This enables us to classify what crops are produced in each grid location. Figure 4 illustrates the size of this gridded dataset for Hardoi District in Uttar Pradesh, where each square refers to a 5 squared km grid. Within the sample area, we

[10]These are the same states considered in other papers focusing on air pollution in South Asia, e.g. Liu et al. (2020).

are able to observe 178,352 grids in each month and year from 2012 to 2021, for a total of 17,623,560 observations.

Health outcomes: We use geo-localized individual-level health information from surveys collected by the Demographic and Health Survey in India (2014-15) and Pakistan (2016-17).

4 Administering Pollution When Costs Are Internalized

We being by analyzing when the incidence of fires decrease. First, we show that when location is more likely to pollute the district itself, the number of fires in that location decrease. Second, probing the location and timing of this phenomena, we show that it is more intensely observed for areas that produce rice, particularly in the harvest months, when the stubble has to be cleared quickly to sow the winter wheat crop.

4.1 Definition of Treatment: Downwind Areas

For each grid, we want to establish the degree to which a smoke from a fire will pollute its own jurisdiction. This influences how much district administrators should care about reducing the incidence of fires in that grid. Variation from changes in expected wind directions will affect the polluting potential of each grid cell over time. For each month-year, we measure administrators' expectations on wind directions by using 10 month-year rolling averages of wind direction. We then calculate the area of the jurisdiction that is downwind, as well as upwind, from each grid in each month using our measure of expected wind direction. We do this by sketching a 180 degree line that passes through the grid centroid and is perpendicular to the wind direction vector.

We define a grid as being treated in a month-year if the area of the jurisdiction downwind from the grid is larger than the area upwind.[11] Figure 6(a) provides a visual illustration of this treatment definition. On the left hand side is a treated grid. When wind blows from the top-left quadrant, smoke from a fire occurring in that grid will affect more than half the jurisdiction (district) where the grid is located. On the right hand side is a control grid: the area downwind from that grid is less than the area upwind from that grid. Therefore, this area is going to matter less than the first grid for the administrator, keeping wind direction

[11] Our method to define downwind areas follows the approach proposed by Morehouse and Rubin (2021). Splitting the jurisdiction into up and downwind areas by using the 180 degree line captures the expectations of smoke dispersal for district administrators.

fixed.

Figure 6(b) illustrates how treatment status varies over time. In time 1, when wind blows from the top-left corner, the top-left grids are treated: a fire taking place in those grids would affect most of their own jurisdiction. But as winds change direction in time 2 (right panel), those grids become control because smoke from those grids would no longer affect most of their jurisdiction.[12] Figure 6 illustrates the actual representation of treatment status variation in different months.

4.2 Empirical strategy

4.2.1 Main estimation strategy

We test if the number of fires is higher for treated grids, that is grids from which most of the jurisdiction is downwind. To do so, we use changes in wind direction to estimate a difference-in-differences model that leverages within grid changes in treatment status over time. For a 5 squared km grid i in jurisdiction (district) j and month-year t, we estimate:

$$Fires_{it} = \beta Downwind_{it} + \alpha_i + \gamma_j \times \theta_t + X_{it} + \epsilon_{it} \tag{1}$$

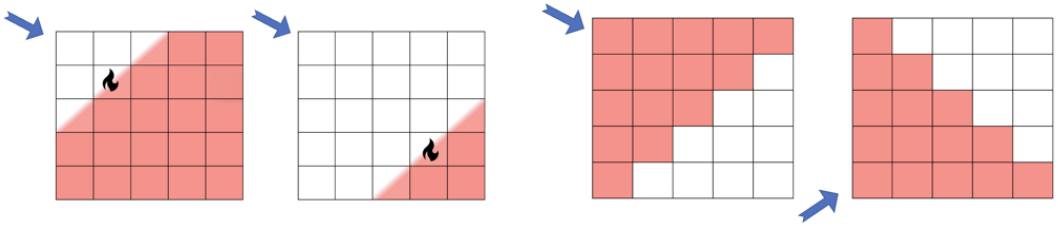
where our dependent variable is the number of fires in a grid-month. Standard errors are clustered at the grid and month-year times district level. The first accounts for temporal dependence of observations at the treatment status level (grid). The second accounts for potential spatial correlation in wind patterns. The coefficient of interest β captures the effect of having most of the jurisdiction exposed to smoke on the number of fires in that grid. In our main specification, we define *Downwind* as a dummy equal to 1 in periods in which the area downwind from the jurisdiction is larger than the area upwind. We show the robustness of results to alternative definitions of treatment.

The high-dimensionality of our data allows for the inclusion of fine-grained fixed effects. First, α_i are 5 squared km grid fixed effects that absorb time-invariant factors for over 178,000 grid locations in our data, such as crop suitability or economic characteristics of small locations. Second, $\gamma_j \times \theta_t$ are jurisdiction by month-year fixed effects that account for time-varying shocks at the district level, such as changes in administrative personnel. X is a vector of controls, including average wind speed and direction, such that we compare units with similar exposure to winds.

A typical concern in the standard difference-in-differences setting is that treatment might

[12] While in this example all treated grids become control grids in time 2, this is not necessarily the case in real jurisdictions, which have more irregular shapes. Some grids can stay treated or control for the entire period.

Figure 5: Illustration of treatment definition and treatment variation



((a)) Treatment definition ((b)) Shifting treatment

Note: The blue arrow represents wind direction, the square a district, the smaller squares grids. (a) In the left district, the area downwind from the fire (N=20.5 grids) is larger than the area upwind from the fire (4.5 grids). We define this grid as treated. In the right panel, the reverse is true: we define this grid as control. (b) When winds blow from North-West, the top-left grids are treated: a fire in these grids would produce smokes affecting most of the jurisdiction. Treated grids shift when winds blow from South-West.

be assigned based on pre-existing characteristics of the unit endogenous to the outcome – for example, poorer areas might be more likely to receive welfare interventions, or highly competitive constituencies might be reached by political campaigns at higher rates. Our setting has instead a unique feature: treatment – which area is more likely to pollute a jurisdiction – is assigned based on changes in wind direction, which are exogenous. This feature makes the parallel trends assumption – that treated grids would have displayed a similar trend in fires as the control grids had the wind not shifted, exposing a larger area of the jurisdiction to smoke – considerably more credible, as treatment is assigned quasi-randomly. We nevertheless provide evidence in favor of this assumption testing the event study specification of our main equation.

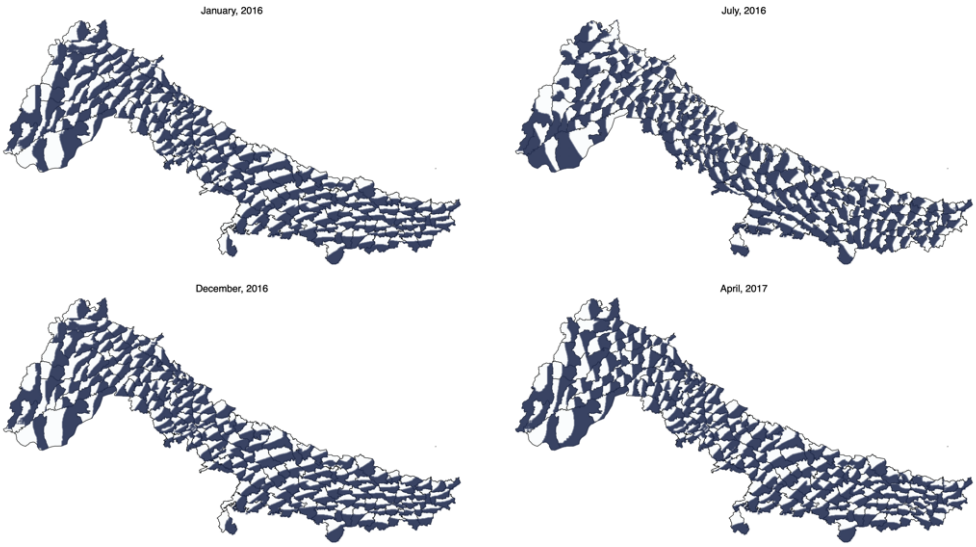
4.2.2 Event-study specification

To trace the behavior of fires around the time in which grids switch into treatment, we estimate the event-study equivalent of Equation 1. For relative time r from the switch to treatment D, we estimate:

$$Fires_{it} = \sum_{r=-1}^{z} \zeta D_{ir} + \alpha_i + \gamma_j \times \theta_t + X_{it} + \epsilon_{it} \tag{2}$$

where our coefficients of interest ζ represents the effect in each period before and after treatment, relative to the switching month-year (the omitted category). For example, if changes in wind direction convert grid i from control to treatment in May, relative time will take value 0 in May, +1 in June and values -1 in April, -2 in March, and so on.

Figure 6: Treated grids over time



Note: We display an example of the spatial distribution of our treatment and its temporal variation in four different months. In blue, treated grids – those from which most of the district would be exposed to smokes if a fire was to be set in that area. In white, control grids – those from which most of the district would be unexposed. The full monthly variation can be seen at this link (Dynamic treatment).

The event study differs from the main analyses in two ways: first, we drop grids that are treated in the first period in the sample, January 2012, because their relative time cannot be established; second, we restrict the analyses to the sample with a sufficient number of relative time periods, keeping only observations in a six-months window around treatment.[13]

4.3 Results

Table 1 shows the effect of a grid becoming treated on the number of fires. On average, when the district area downwind from the fire is larger than the upwind area, fire occurrence drops by 0.015 to 0.02 less fires per month per grid cell – a 10 to 14 percent reduction compared to the average number of fires in a grid-month. The finding is robust to a wide range of fixed effects and controls, as discussed above, always remaining significant at the 1% level.

To further establish the substantive bases of our argument, we check if the treatment effects we document are indeed driven by historically rice producing areas, the crop whose stubble clearing creates a major source of pollution. Table B.1, Columns 1 and 2 document that the incidence of fires is higher in historically rice producing areas, but that this effect is many times larger during the harvest season, when residue needs to be quickly cleared to

[13]For a similar approach, see for example, Dobkin et al., 2018. Figure A.2, SI, shows that the vast majority of our sample's density is in this period.

Table 1: Effect of district exposure to smokes on fire occurrence

	(1) N Fires	(2) N Fires	(3) N Fires	(4) N Fires
Down>Up	-0.0157***	-0.0202***	-0.0189***	-0.0145***
	(0.00571)	(0.00389)	(0.00296)	(0.00247)
Observations	17,623,560	17,623,560	17,623,560	17,623,560
Month year fe	N	Y	-	-
District fe	N	Y	-	-
Wind controls	N	N	Y	Y
District x Month fe	N	N	Y	Y
Grid fe	N	N	N	Y
Mean DV	0.148	0.148	0.148	0.148

Notes: The dependent variable is the number of fires in a 5 squared km grid and month. The treatment indicator is a dummy equal to one in months and grids from which most of the district area is exposed to smokes (i.e. when the area of the district downwind is larger than the area upwind from the grid). Controls include average wind direction and wind speed in each grid and month. Data include Punjab (Pakistan), Punjab (India), Haryana, Uttar Pradesh, Delhi, and Bihar from 2012 to 2021. Standard errors are clustered at the grid and month×district level. We report the mean of the dependent variable in the control group in table.

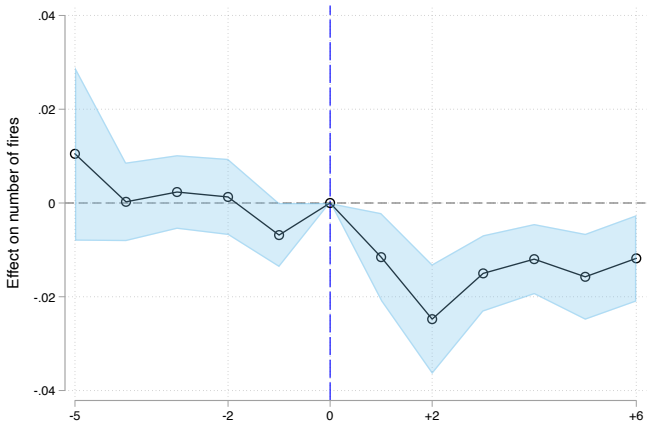
sow the winter wheat crop. In Column 3, we consider how treatment effect varies by rice producing areas and harvest times. First, the effect of treatment exists even in areas that do not produce rice. Second, there is no detectable differential effect of treatment in rice producing areas outside of harvest months, or in harvest months outside of rice producing areas. Third, in rice producing areas and during rice harvest months, treatment effect is about six times larger. Finally, summing all treatment coefficients, the overall reduction in fires caused by treatment in rice growing regions and harvest months is of 0.067, a 54% drop with respect to the average number of fires in the control group.

Overall, this first set of results shows that jurisdictional incentives matter for the timing and location of fires, and that these incentives matter particularly in places where the qualitative literature suggests the highest amount of residue burning should be located. This points to the role of district-level administrators in reducing fires when and where those would be most damaging for the constituents.

4.3.1 Event study results

In Figure 7, we plot the coefficients ζ, capturing the difference between treated and control units in each relative time period from when a grid switches to treatment. The flat pre-

Figure 7: Effect of district exposure to smokes on fire occurrence by relative time period from exposure



Notes: Event study as in Equation 2. We estimate the effect of treatment on the number of fires in each period before and after a grid switches to treatment. The reference category is t=0, when the switch takes place. We restrict the analyses to a 6 months window around the switch to maximize support.

period indicates that there is no difference in the number of fires between grids that will be treated and control grids. Instead, right after a grid switches to treatment, we observe a large and persistent drop in fires, reaching its lowest point two months after treatment and remaining stable in the following periods.

4.3.2 Robustness

We start by testing the robustness of our findings when we relax the classic difference-in-differences assumptions. We relax the treatment homogeneity assumption and adopt the estimator proposed by De Chaisemartin and d'Haultfoeuille (2020), which is robust to heterogeneous treatment effects (Table B.8). A second assumption is that there are no carry over effects, such that units can go in and out of treatment without biasing the estimates. We relax this assumption by estimating our model on a restricted sample in which the treatment is an absorbing state (that is, units remain treated, and we drop all periods from the end of the 1st switch to treatment onwards, Table B.8).

Second, we test whether findings are robust to adopting different specifications. First, we vary the definition of when a grid is considered most likely to expose the district to smoke (Table B.3): we restrict our attention to areas in which the downwind area is larger than the upwind area by one standard deviation or more, and we adopt two continuous definitions of treatment, the percentage difference between down and upwind area and the percentage change in downwind over upwind area. Second, we vary the definition of the dependent

variable, replacing the count of fires with a dummy equal to one for whether any fire took place and a continuous variable capturing the average brightness of fires in a grid-month (Table B.4). Third, findings are robust to using a Poisson – accounting for our depending variable being a count – and Negative Binomial distribution – accounting for the skew of the dependent variable towards zero (Table B.5).

Finally, we show that results are not dependent on the sample we use. First, we subset our findings by each of the states in the sample (Table B.6) and find that effects are significant and of similar magnitude in all states, except in Delhi, where the coefficient has similar magnitude but is insignificant, likely since this is the smallest and most urban state in our sample. Second, results are significant and of similar size in each of the years considered, although the magnitude increases in more recent periods, consistent with increased salience of stubble clearing via burning (Table B.7). Finally, results are also robust to adopting the restricted sample used in the event study (Table B.2).

5 How Pollution Responds to Administrators' Incentives

5.1 Administrative Effort is Proportional to Pollution of Constituents

The evidence presented so far shows that internal jurisdiction-specific considerations influence fire incidence within the district. We now consider how this internal calculus affects the incidence of fires and pollution in other jurisdictions. In this section, we ask whether fires increase as we move closer to downwind district borders – e.g. those areas in which smoke from fires would affect a neighboring jurisdiction. Consistent with previous work, the idea is to see if fires are monotonically more likely to occur as the externality is more likely borne by a neighboring jurisdiction, and vice versa progressively less likely to occur as the externality is more likely to be born by self.

Fire incidence is the result of two forces: administrators' action to reduce fires where they are most damaging, and administrators' inaction towards fires which would affect a neighboring jurisdiction. Considering the difference between upwind and downwind borders allows us to separate these two components and understand to what extent administrators' behavior could reduce fire incidence. We consider the following logic: if a fire was to take place close to an upwind border, it would direct pollution towards a large area of the district, leading district administrators to take action to avoid fires (or farmers to anticipate bureau-

cratic action to curb them). Conversely, if a fire was to take place close to the downwind border of the district, it would affect mostly the neighboring jurisdiction, creating incentives for district administrators to take no actions to reduce fires.

5.1.1 Empirical Strategy

For this analysis, we generate a new database at the level of the pair of a grid and each district bordering the district in which the grid is located. Each pair is observed in all time periods, such that the same pair switches from being up to downwind at different times. This means that we can account for district times month-year fixed effects, as well as fixed effects for each pair grid-border. We also calculate the distance of each grid to each of its corresponding border.

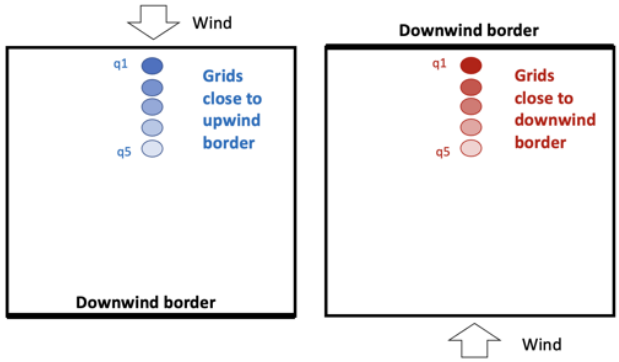
Since our treatment is time-varying, we can improve on previous work in two ways: First, we can estimate within unit changes over time, allowing us to hold time-invariant factors related to bordering areas constant. Second, we can estimate the effect of both the upwind and downwind borders, which allows us to understand if the effort bureaucrats exert to prevent fires in upwind borders is equivalent to the effort not exerted to prevent fires in downwind borders, holding fixed the attributes of the specific borders and specific locations.

Calling a pair of grid-border segment g, we estimate, for each segment g belonging to grid i and district j :

$$Fires_{gt} = \sum_{q=1}^{5} \gamma Distance_{gt} + \sum_{q=1}^{5} \lambda Distance \times DownwindNeighbor_{gt} + \zeta_g + \eta_j \times \theta_t + \mu_{gt} \quad (3)$$

where q are quintiles of the distance of each grid to the bordering district and we control for grid-border pair and district times month-year fixed effects. We use the fifth quintile (the farthest) as absorbed category. We are interested in two coefficients: γ captures the effect of proximity to the upwind border, an area where bureaucrats have the largest incentives to curb fires, as those would affect most of their own district. Instead, λ captures what happens as we approximate the downwind border, which might be the same location as an upwind area in other periods. Our expectation is that bureaucrats have little incentives to act on fires when those are located close to downwind borders, as smokes would exit their district. The logic of this estimation is further illustrated in Figure 8.

Figure 8: Graphical illustration of the estimation strategy in Equation 3



Note: (Left panel) When wind blows from the north, grids closer to the upwind border (closest quintile in darkest blue) are those in which bureaucrats have most incentives to curb fires as most of the district would be exposed to pollution. (Right panel) When wind shifts, the same locations mostly affect the neighboring district (the smallest self-exposure from the closest quintile in darkest red). Here incentives are reversed: the closest the grid is to the downwind border, the smallest the incentive for bureaucrats to reduce fires.

5.1.2 Results

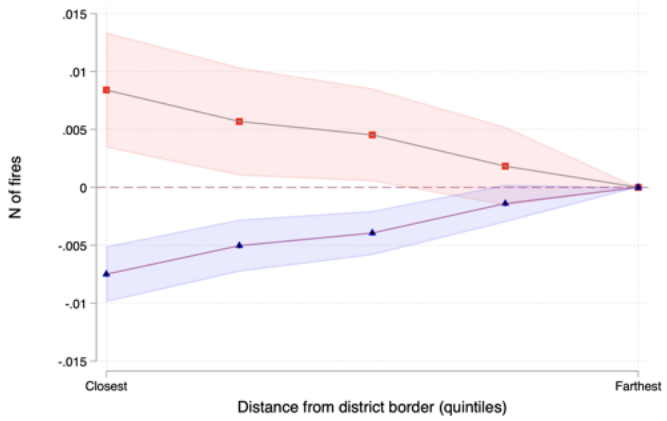
Figure 9 plots the results of this analysis. First, fires drop monotonically as we approximate the upwind border, decreasing more the closer we move towards the border, where most smokes would affect a larger share of the bureaucrats' own district (coefficients in blue). Second, fires monotonically increase as we approximate the downwind border (coefficients in red). Third, the reduction in fires close to the upwind border is about half the size of the increase in fires close to the downwind border (absolute difference between red and blue). This indicates that there is a lot of scope for bureaucrats to reduce total fires if they exert effort such that it produces outcomes similar to those in areas closest to the upwind border.[14]

5.2 The Strategic Allocation of Effort

Consider the case in which two administrators have the same resources but the problem of crop stubble burning is much larger in one versus the other. This is the case in districts where a large share of the area is covered by rice crops relative to districts where rice production is minimal. We ask if bureaucrats are even more strategic in targeting their action to reduce pollution where more resources would be needed to tackle the problem? In Figure C.2, we compare the effects of treatment by distance to the upwind and downwind border when bureaucrats administer a district covered by rice crops. The effects observed in the whole

[14]In Section 7.4 we do back of the envelope calculations to examine how a policy that brought this outcome would affect health outcomes.

Figure 9: Effect of distance to the upwind (in blue) and downwind (in red) border on fires



Note: (Red/squares) Effect of switching to being in a grid facing a downwind district border, where smoke would travel to a neighboring district. (Blue/triangles) Effect of switching to being a grid facing an upwind district border, where smokes would affect the district in which the grid is located. All effects are plotted by quintiles of distance to the border, and the 5th quintile is used as absorbed category. The regression includes fixed effects for grid-border pairs and district times month-year and controls for wind speed and direction. Standard errors are clustered at the grid and district times month-year level.

sample (the same plotted in Figure 9, here in dotted lines) are even more amplified: fires are allowed at higher rates in areas that would not damage the district (e.g. at the downwind border), while they are reduced even more in areas that would (e.g. at the upwind border). This result is consistent with a strategic reallocation of bureaucratic effort in face of larger constraints: effort is concentrated more in the most important areas of the district and redistributed away from less important areas.

5.3 Cost of Coordination: the India-Pakistan Border

Our evidence suggests that, on average, the reduction in fires is largest when in grids closest to the upwind border of the jurisdiction (where smoke from fires would pollute the entire district) and smallest in grids close to the downwind border (where smoke would travel elsewhere, Figure 9). If district-level bureaucrats drive this effect, we should expect this effect to change depending on who is on the other side of the border. In particular, at the border between India and Pakistan, coordination between bureaucracy should collapse, particularly for air pollution for which there exist no treaties or fora for coordination.[15] We test this hypothesis using the same specification used in Equation 2 but adding an interaction with a dummy for whether a grid-border pair is also a border across countries. Given the

[15]In contrast, the Indus Water Treaty of 1965 helps govern the river system across the two countries.

time varying nature of our treatment, we are able to identify if the change in fires in the same grid after a switch in wind makes it a source of pollution exposure for its own district or for the other country, and compare it with another grid exposing its own district or a neighboring district in the same country. The increase in fires is almost 5 times larger when we approximate the India-Pakistan border compared to when we approximate any district border in the same country (see Figure C.1).[16] This finding, indicating that treatment effects respond to variation in costs of coordination among administrations, provides further evidence that district-level bureaucrats drive fires patterns.

6 Capacity Constrains and the Administration of Pollution

Can the decrease in fires under treatment, where grids are likely to pollute their own jurisdiction, be tied to bureaucratic behavior? We do not observe bureaucratic behavior directly as the area our paper covers is very large. The approach we take in this section triangulates evidence to show that treatment effects are only present or larger when: (1) administrators have the administrative capacity to observe and curb fires; (2) administrators have not just turned over; and (3) pollution is highly visible due to atmospheric conditions, increasing the incentives to tackle it.

6.1 Legibility of Polluting Behavior

If administrative action drives the reduction in fires when one's own jurisdiction is most exposed to smoke, we should expect efforts to be targeted towards areas where there is sufficient legibility over fire occurrence, that is, where the administrator or someone else besides the farmer can observe fire-setting behavior. Areas close to roads are such areas.[17] We calculate the distance between each 5 squared km grid and the nearest main road in 1992.[18] We run our main analysis by 0.5km segments of distance to a main road. Notice that this is not a heterogeneous treatment effects exercise as, with grid fixed effects, we have identification with changes in wind direction over time. Figure D.1 shows that treatment effects are concentrated in the area at a short (less than 4km) distance to a main road.

[16] Notice also that the effect of being closer to an upwind border is larger, possibly relating to the larger awareness about pollution spillovers in districts that border other high-polluting states.

[17] This was suggested to the authors in an interview with a Pakistan Administrative Services bureaucrat on Mar 22, 2022.

[18] Data on roads comes from DCW-GMT which maps roads up to 1992, before the start of our database and before attention to air pollution was high. This feature allows us to exclude that roads are endogenous to the outcome as they are unlikely to have been built to reach specific locations where fires are common.

This finding is in line with what we would expect from a bureaucrat-driven reduction in fire occurrence, existing only thanks to heightened state legibility.

6.2 Bureaucratic Experience

If the reduction in pollution is driven by bureaucrats' behavior, then we should expect this strategic behavior to reduced when bureaucrats are inexperienced. In the months in which the district collector has just changed, there is an automatic reduction in bureaucratic experience, as the new administrator settles in the new position. We collect original data on all district collectors/magistrates tenure in our sample in Pakistan and India and build a time series of district officer turnovers. We use this data to estimate treatment effects in each period from bureaucratic turnover on the number of fires in an event study setting. We adopt the same specification as in Equation 2 but we replace periods from treatment with periods from bureaucratic turnover interacted with a dummy equal to 1 when a grid is treated. Each coefficient in Figure D.2 represents the treatment effect for each month since the change in the district officer. As we should expect, there is no particular trend in fire occurrence in treated grids in the periods before a turnover. Instead, in the month right after the turnover, there is a large relative increase in fire occurrence. Our interpretation is, in initial period where bureaucrats are likely building capacity to administer, the strategic treatment effect we observe in other periods reverses. That is, bureaucrats exhibit less control in reducing the most damaging fires for their district. This anomaly, however, quickly corrects as fires return to displaying strategic bureaucratic behavior, in line with pre-period levels, 4 months after the new bureaucrat takes office.

6.3 High Visibility of Pollution

If bureaucrats drive the reduction in fires rather than farmers or other phenomena, we should expect that they exert more effort in times in which their lack of action would be most criticized. We use the fact that during winter months pollution is exogenously more visible as a result of atmospheric conditions.[19] Figure 3 shows that there are actually two peaks in terms of fire activity every year. This correlates with the two crops in which crop stubble burning is most used, rice and wheat, the first harvested during winter and the second during summer months. A heterogeneity analysis reveals that, indeed, treatment effects are only present for the winter-harvested months and is absent for the summer-harvested ones (see Table D.1).

[19]The colder air in the winter makes the pollution particles more visible in the winter months, leading to a thick layer of smog.

7 The Health Impacts of Crop Stubble Burning

This section examines the consequences of crop burning-induced air pollution for a population at particularly high risk from exposure, infants in-utero. We provide the first causal evidence of the direct health costs of crop burning in South Asia. These estimates allow us to do a back-of-the envelope calculation to retrieve how much the strategic behavior of bureaucrats contributes to (or could reduce) health deterioration in this region.

7.1 Linking fires to downstream recipients of pollution

Whether a child is exposed in-utero to the smoke from any given fire depends on a number of conditions: the distance from the fire, the direction the wind blows, its speed, the atmospheric pressure, the temperature, and so on. To link fire occurrence in a particular place and time with whether a person in another location is exposed to smoke from that fire, we use a meteorological model which factors in all these, as well as other weather conditions. The HYSPLIT[20] – an algorithm used widely by meteorologists and scientists – models the non-linear trajectory taken by PM2.5 particles released by each fire in our database. A visual illustration of the resulting patterns is presented in Figure 11(a), which plots the non-linear trajectories of all fires occurring in one month in our sample. We calculate all trajectories from all fires occurring in from 2012 to 2021.

Next, we overlay these trajectories to the DHS health data, which are organized in clusters. The overlay of one fire trajectory with the DHS clusters it crosses is illustrated in Figure 11(b). We count how many fires, coming from any location, affect any given DHS cluster (in the example in Figure 11(b), the pollution from the fire passes through 53 DHS clusters over a 24 hour period). This process allows us to link each fire to pregnancies receiving the smokes from those fires, thereby generating a precise measure of the expected air pollution they receive while gestating.

For each pregnancy (N=179,000 across 28,000 DHS clusters over 7 years), we build a database spanning the 9 months before the expected delivery and merge each mother-location-month to our measure of air pollution exposure. We then collapse the data at the pregnancy level, calculating the sum of the total pollution received by each child while in-utero. The temporal variation in pregnancies creates within-DHS cluster variation in exposure to air pollution.

[20] `https://www.ready.noaa.gov/HYSPLIT.php`

Figure 10: Illustration of HYSPLIT calculation



((a)) Trajectories from May 2019 fires (HYSPLIT)

((b)) Particulate path over DHS clusters for one fire

7.2 Estimation strategy

Our quantity of interest is the impact of the number of fires a child is exposed to while in utero on mortality. This quantity is generally confounded by characteristics of the location, period of birth and of the family which might be correlated with both air pollution and correlates of bad health outcomes. We can address all of these concerns in our setting. The degree to which each child is exposed to fire-related air pollution while in-utero depends on the timing as well as the location of the mother while pregnant. Since we observe multiple mothers in the same location and month, and since some mothers have more than one child, we can estimate the effect of the number of fires a child is exposed to using fixed effects for time (t), location (c) and even for the same mother (m), ruling out genetic and behavioral confounders of mortality. We estimate:

$$Mortality_{mct} = \zeta Fires_{ct} + \theta_t + \kappa_c + \gamma_m + \nu_{mct} \tag{4}$$

where our coefficient of interest, ζ, identifies the effect of exposure to fire-related air pollution in-utero on child and infant mortality, accounting for time, location and family fixed characteristics. Child mortality is defined as the number of deaths in the first 5 years of birth per 1000, while infant mortality is deaths within the first year of birth per 1000.

7.3 Effect of fires on infant and child mortality

Table 2 presents the results. In-utero exposure to air pollution from any number of fires above zero increases child by between 17 and 30%, and infant mortality by 16-25%. Considering a continuous definition of exposure, receiving 1% more fires while in utero increases both child

mortality and infant mortality by between 2 and 10%. The size of the effect becomes larger as we move towards the most restrictive specification. This effect is robust to modelling fire exposure as the inverse hyperbolic sine transformation of the number of fires a pregnant person is exposed to, or as the terciles of fire exposure (Table E.1). As a reference point to understand these magnitudes, exposure to fires while in-utero move the average child mortality from the levels of countries like Namibia and Senegal (value of 39, in line with world averages of 37) to the levels of Gambia and Madagascar (50).[21]

Table 2: Effect of In-Utero Exposure to Fires on Infant and Child mortality

	Panel A: Child Mortality					
	(1)	(2)	(3)	(4)	(5)	(6)
Any fire	16.83***		28.76*		29.89**	
	(1.217)		(16.90)		(12.16)	
Ln fires		2.714***		1.084		9.504**
		(0.209)		(1.930)		(4.290)
Observations	179,194	179,194	177,480	177,480	67,102	67,102
R-squared	0.002	0.002	0.149	0.149	0.500	0.500
Time fe	Y	Y	N	N	Y	Y
DHS cell fe	N	N	Y	Y	Y	Y
Mother fe	N	N	N	N	Y	Y
Mean DV	39.341	39.341	39.341	39.341	39.341	39.341
	Panel B: Infant Mortality					
	(1)	(2)	(3)	(4)	(5)	(6)
Any fire	15.96***		28.76*		24.74**	
	(1.177)		(16.90)		(11.84)	
Ln fires		2.529***		2.181		10.18**
		(0.202)		(1.875)		(4.194)
Observations	178,652	178,652	176,932	176,932	66,582	66,582
Time fe	Y	Y	N	N	Y	Y
DHS cell fe	N	N	Y	Y	Y	Y
Mother fe	N	N	N	N	Y	Y
Mean DV	36.717	36.717	36.717	36.717	36.717	36.717

Note: Effect of in-utero exposure to any / Ln number of fires on child (Panel A) and infant mortality (Panel B). Fixed effects for birth date (time), location (DHS cell) and for having the same mother are all included in the most restrictive specifications (Columns 3 and 6). Standard errors are clustered at the DHS cell level.

[21]The range is from 2 (Iceland) to 115 (Somalia) (World Bank, child mortality data 2020).

7.4 Back of the envelope calculation

We have shown that (1) bureaucratic strategic behavior impacts the incidence of fires, and that (2) fires have a substantial impact on child and infant mortality. We now connect these two findings to assess how much of the increase in mortality determined by fires can be attributed to jurisdictional incentives. We do a back of the envelope calculation to understand what would happen if bureaucrats had the same incentives and capacity across the entire district area, and acted everywhere in the same way as where their own jurisdiction bears the brunt of the pollution. Of course, a problem with bureaucratic action is the lack of capacity, but this scenario illustrates an upper bound on the effects that could be obtained by creating slack in bureaucratic capacity through the prioritization of environmental management. Under this scenario, fires in downwind grids could be reduced by the amount of our treatment effect (as shown in Table 1): that is, there could be 15 fewer fires than we observe (a 14.5% drop with respect to the dependent variable). In downwind grids located in rice producing areas and in the harvest months, the strategic reduction is 67 fewer fires (or a 55% reduction with respect to the mean dependent variable) (Table B.1).

Now consider what effect this reduction in pollution would have on child mortality: we know that a 1% increase in fires over the 9 months of the pregnancy raises child mortality by approximately 0.1 in a 1,000 children (Table 2). If bureaucrats had the same incentives and capacity to reduce pollution in the entire district, they would contribute to as much as a 5.5 deaths reduction in 1,000 children, moving child mortality from current levels in the area (39) to levels of countries like South Africa (33).

8 Conclusion

States have an increasingly important role in fighting climate change and environmental degradation. Public opinion and media attention is generally focused on international agreements aimed at reaching new shared legislation. However, on a daily basis, states are in-charge of enforcing already existing legislation whose enforcement or failure to enforce has enormous effects on carbon footprint and health outcomes. In this paper, we focus on one such case, air pollution from stubble burning in South Asia. Even though legislation exists to curb this practice, uneven enforcement continues to make this phenomenon the largest health emergency on the planet.

We show that crop stubble burning patterns – as detected by observations of fires from space – respond to internal political economy dynamics within the state: district-level incentives to be accountable to one's own constituents and unaccountable to constituents in other

districts. We demonstrate that jurisdictionally strategic enforcement of pollution regulations has a large effect on fire incidence as well as on child and infant mortality. Indeed, quoting the policy-experts on this topic Jitendra et al., 2017: "What is threatening is the fact that pollution does not recognise political boundary. Pollution from this large-scale residue burning ultimately affects the larger air shed, impacting all of us, irrespective of whether we are in Delhi or Harda."

One view is that jurisdictional incentives exacerbate the problem. Another way to view our results is to show that jurisdictions create the precise incentives that enable bureaucrats to take action where needed. How can these incentives be strengthened so that administrators are able to further curb the problem? Jitendra et al., 2017 attributes China's success in reducing crop burning exactly to administrators incentives: "China... banned stubble burning in 1999 and penalises officials under whose jurisdictional area the burning is detected. The fines go up to USD 80,000. In 2015 in Henan county government collected USD 37 million in such penalties."

References

Abdurrahman, Muhammad Isa, Sukalpaa Chaki, and Gaurav Saini (2020). "Stubble burning: Effects on health & environment, regulations and management practices". In: *Environmental Advances* 2, p. 100011.

Cusworth, Daniel H et al. (2018). "Quantifying the influence of agricultural fires in northwest India on urban air pollution in Delhi, India". In: *Environmental Research Letters* 13.4, p. 044018.

Dasgupta, Aditya and Devesh Kapur (2020). "The political economy of bureaucratic overload: Evidence from rural development officials in India". In: *American Political Science Review* 114.4, pp. 1316–1334.

De Chaisemartin, Clément and Xavier d'Haultfoeuille (2020). "Two-way fixed effects estimators with heterogeneous treatment effects". In: *American Economic Review* 110.9, pp. 2964–96.

Deryugina, Tatyana et al. (2019). "The mortality and medical costs of air pollution: Evidence from changes in wind direction". In: *American Economic Review* 109.12, pp. 4178–4219.

Dobkin, Carlos et al. (2018). "The economic consequences of hospital admissions". In: *American Economic Review* 108.2, pp. 308–52.

Graff Zivin, Joshua and Matthew Neidell (2012). "The impact of pollution on worker productivity". In: *American Economic Review* 102.7, pp. 3652–73.

Greenstone, Michael and Rema Hanna (2014). "Environmental regulations, air and water pollution, and infant mortality in India". In: *American Economic Review* 104.10, pp. 3038–72.

Greenstone, Michael, Mar Reguant, et al. (2021). "Evidence paper: Engergy & Environment". In: URL: `https://www.theigc.org/wp-content/uploads/2019/12/IGC-Energy-and-Environment-evidence-paper-3003.pdf`.

Heft-Neal, Sam et al. (2019). *Air pollution and infant mortality: Evidence from Saharan dust*. Tech. rep. National Bureau of Economic Research.

Jack, B Kelsey et al. (2022). "Money (Not) to Burn: Payments for Ecosystem Services to Reduce Crop Residue Burning". In.

Jayachandran, Seema (2009). "Air quality and early-life mortality evidence from Indonesia's wildfires". In: *Journal of Human resources* 44.4, pp. 916–954.

Jitendra, Shreeshan Venkatesh et al. (2017). *India's burning issue of crop burning takes a new turn*. URL: `https://www.downtoearth.org.in/coverage/agriculture/river-of-fire-57924` (visited on 08/08/2022).

Kahn, Matthew E, Pei Li, and Daxuan Zhao (2015). "Water pollution progress at borders: the role of changes in China's political promotion incentives". In: *American Economic Journal: Economic Policy* 7.4, pp. 223–42.

Knippertz, Peter et al. (2015). "The possible role of local air pollution in climate change in West Africa". In: *Nature Climate Change* 5.9, pp. 815–822.

Lan, Ruoyu et al. (2022). "Air quality impacts of crop residue burning in India and mitigation alternatives". In: *Nature Communications* 13.1, pp. 1–13.

Lelieveld, Jos et al. (2015). "The contribution of outdoor air pollution sources to premature mortality on a global scale". In: *Nature* 525.7569, pp. 367–371.

Lipscomb, Molly and Ahmed Mushfiq Mobarak (2016). "Decentralization and pollution spillovers: evidence from the re-drawing of county borders in Brazil". In: *The Review of Economic Studies* 84.1, pp. 464–502.

Liu, Tianjia et al. (2020). "Crop residue burning practices across north India inferred from household survey data: bridging gaps in satellite observations". In: *Atmospheric Environment: X* 8, p. 100091.

Lu, Jackson G (2020). "Air pollution: A systematic review of its psychological, economic, and social effects". In: *Current opinion in psychology* 32, pp. 52–65.

Maue, Casey Campbell (2021). *Climate, Agriculture, and Economic Development: Theories and Methods for Agricultural Productivity Analysis*. Stanford University.

Morehouse, John and Edward Rubin (2021). "Downwind and out: The strategic dispersion of power plants and their pollution". In: *Available at SSRN 3915247*.

Myllyvirta, Lauri (2020). "Quantifying the economic costs of air pollution from fossil fuels". In: *Centre for Research on Energy and Clean Air*.

Singh, Ramesh P and Dimitris G Kaskaoutis (2014). "Crop residue burning: a threat to South Asian air quality". In: *Eos, Transactions American Geophysical Union* 95.37, pp. 333–334.

Singh, Siddharth (2018). *The great smog of India*. Penguin Random House India Private Limited.

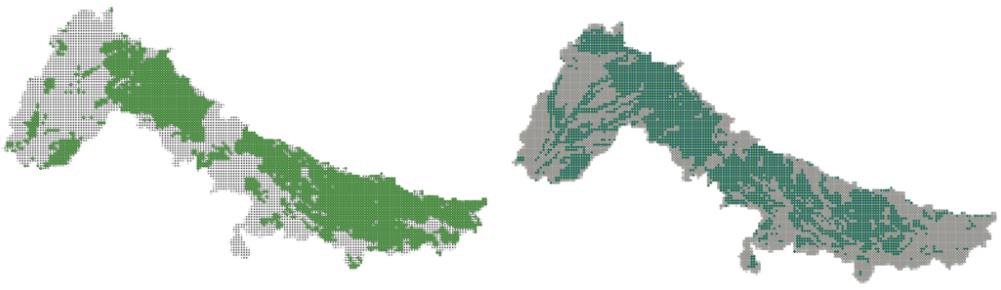
SUPPLEMENTARY INFORMATION
— For Online Publication —

A Descriptives

Table A.1: Summary statistics

Variable	Mean	Std. Dev.	N
N Fires	0.14	1.042	17623560
Fire brightness	333.848	10.967	849260
Any fire	0.048	0.214	17623560
Down>Up	0.51	0.5	17623560
Down-Up	100,515,909	4,206,624,325	17,623,560
Down>Up by 1std	0.126	0.332	17623560
Downwind over total area	50.538	29.183	17623560
Down-Up over total area	1.076	58.366	17623560
Wind Speed	1.311	0.562	17623560
Avg wind speed	1.331	0.424	17623560
Wind direction	-11.68	68.298	17623560
Avg wind direction	-10.969	59.074	17623560

Figure A.1: Rice and Wheat Areas in Northern Indian Subcontinent



((a)) Rice crops ((b)) Wheat crops

Figure A.2: Number of observations per relative time period from treatment

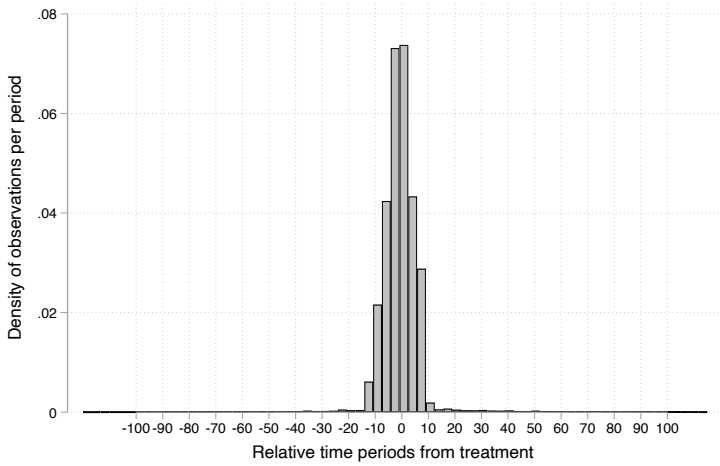
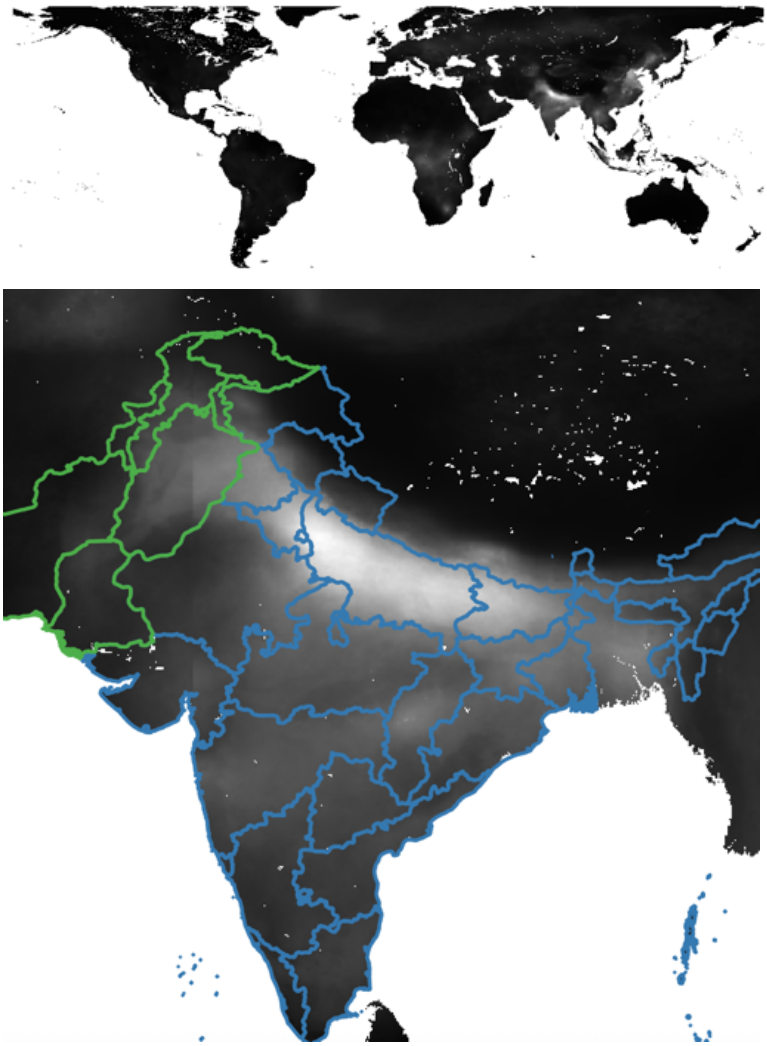


Figure A.3: PM Particle on average 2016 (SDIS)



B Additional Results - Internalized Costs

Table B.1: Effect of rice producing areas and treatment status

	(1) N Fires	(2) N Fires	(3) N Fires
Rice	0.00899***		
	(0.00307)		
Rice x Harvest		0.0863***	0.120***
		(0.00708)	(0.0131)
Down>Up			-0.0102***
			(0.00138)
Rice x Down>Up			-0.00005
			(0.00327)
Harvest x Down>Up			0.000869
			(0.00584)
Rice x Harvest x Down>Up			-0.0668***
			(0.0209)
Observations	17,623,560	17,623,560	17,623,560
District x Month fe	Y	Y	Y
Grid fe	Y	Y	Y

Notes: The table shows the effect of rice producing areas on fire occurrence depending on whether we are in the rice harvest months (October-November) and on treatment status. High rice presence is defined as having a larger than median area of the grid devoted to rice crop production. All regressions include fixed-effects as in the main table, control for average wind direction and wind speed in each grid and month and standard errors are clustered at the grid and month×province level.

Table B.2: Effect of district exposure to smokes on fire occurrence, restricting to eventstudy sample

	(1) N fires
Down>Up	-0.0139***
	(0.00237)
Observations	12,196,891
District x Month fe	Y
Grid fe	Y

Notes: The table reproduces the results of Table 1 restricting the analyses to the sample used for the event-study design, which drops observations treated further than 7 periods from the date in which a unit switches to being mostly downwind.

Table B.3: Effect of district exposure to smokes on fire occurrence, robustness by treatment indicator

	(1) N fires	(2) N fires	(3) N fires
Down>Up	-0.0145***		
	(0.00247)		
Down−Up		-0.000170***	
		(2.53e-05)	
Down−Up percent			-0.000341***
			(5.07e-05)
Observations	17,623,560	17,623,560	17,623,560
District x Month fe	Y	Y	Y
Grid fe	Y	Y	Y

Notes: The table reproduces the results of Table 1 considering different treatment indicators: a dummy for whether the downwind area is larger than the upwind area (as in the main analysis) in Column 1; the difference between down and upwind area in Column 2; the percentage difference in Column 3.

Table B.4: Effect of district exposure to smokes on fire occurrence, robustness by dependent variable

	(1) Any fire	(2) Mean brightness
Down>Up	-0.00458***	-1.535***
	(0.000660)	(0.220)
Observations	17,623,560	17,622,199
District x Month fe	Y	Y
Grid fe	Y	Y
Mean DV	0.0509	16.95

Notes: The table reproduces the results of Table 1 using alternative dependent variables: a dummy for whether any fire took place (Column 1) and the average brightness of fires in a grid-month (Column 2).

Table B.5: Poisson and Negative Binomial

	(1) N fires	(2) N fires
Down>Up	-0.0934***	-0.0561***
	(0.0194)	(0.00271)
Observations	9,904,028	11,784,360
Grid fe	Y	Y
Month-year fe	Y	Y
Controls	Y	Y
District x Monthyear fe	Y	N

Notes: The table reproduces the results of Table 1 using alternative specifications: a poisson to accomodate a count dependent variable, and a negative binomial to accomodate for the high skew of the dependent variable towards zero.

Table B.6: Effect of district exposure to smokes on fire occurrence, robustness by state

	(1) Bihar	(2) Haryana	(3) Delhi	(4) Punjab (In)	(5) Punjab (Pak)	(6) Uttar Pradesh
Down>Up	-0.00706**	-0.0304**	-0.00533	-0.0431**	-0.0108***	-0.00751***
	(0.00299)	(0.0137)	(0.00771)	(0.0169)	(0.00290)	(0.00221)
Observations	2,541,600	1,230,000	39,840	1,424,040	5,816,400	6,571,680
District x Month fe	Y	Y	Y	Y	Y	Y
Grid fe	Y	Y	Y	Y	Y	Y
Mean DV	0.0516	0.203	0.0671	0.889	0.100	0.0604

Notes: The table reproduces the results of Table 1 subsetting the analysis by each of the states in our sample: Bihar, Haryana, Delhi, Punjab (India), Punjab (Pakistan), Uttar Pradesh.

Table B.7: Heterogeneity by year

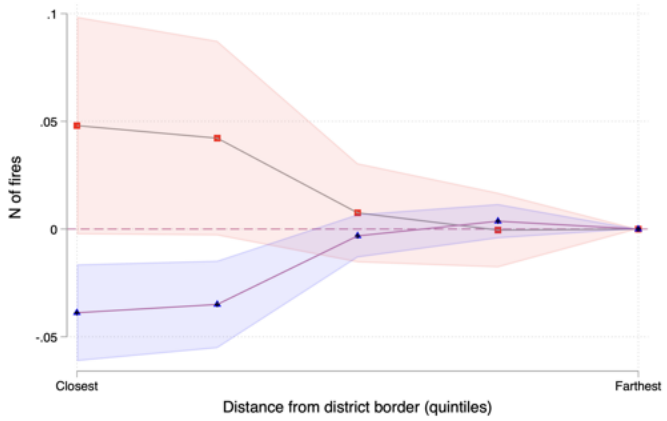
	(1) 2012	(2) 2013	(3) 2014	(4) 2015	(5) 2016	(6) 2017	(7) 2018	(8) 2019	(9) 2020	(10) 2021
Down>Up	-0.0145***	-0.0132***	-0.0157***	-0.0117***	-0.0131**	-0.0145***	-0.00963***	-0.0133***	-0.0237***	-0.0198***
	(0.00406)	(0.00430)	(0.00411)	(0.00354)	(0.00513)	(0.00446)	(0.00344)	(0.00360)	(0.00695)	(0.00519)
Obs	1,762,356	1,762,356	1,762,356	1,762,356	1,762,356	1,762,356	1,762,356	1,762,356	1,762,356	1,762,356
District	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
x Month fe	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Grid fe	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table B.8: Effect of district exposure to smokes on fire occurrence, relaxing DID assumptions

	(1) Treatment Heterogeneity Punjab (In)	(2) Treatment Heterogeneity Haryana	(3) Absorbing Treatment
Down>Up	-0.0256***	-0.0271***	-0.0108***
	(0.0099)	(0.0055)	(0.00339)
Observations	1,412,173	1,219,750	5,788,764
N Switchers	272,189	188,234	-
District x Month fe	Y	Y	Y
Grid fe	Y	Y	Y
Mean DV	0.889		0.0952

Notes: The table reproduces the results of Table 1 relaxing the assumptions of the DID. In Column 1 and 2 we test robustness using the (De Chaisemartin and d'Haultfoeuille, 2020) estimator on the two smallest states in our sample (we are in the process of estimating the same effect in other states, but the size of the sample makes it computationally challenging). In Column 3, we consider treatment as an absorbing state, e.g. we drop all periods from the second time a unit switches to treatment onwards and only consider the periods before and after the first switch to treatment.

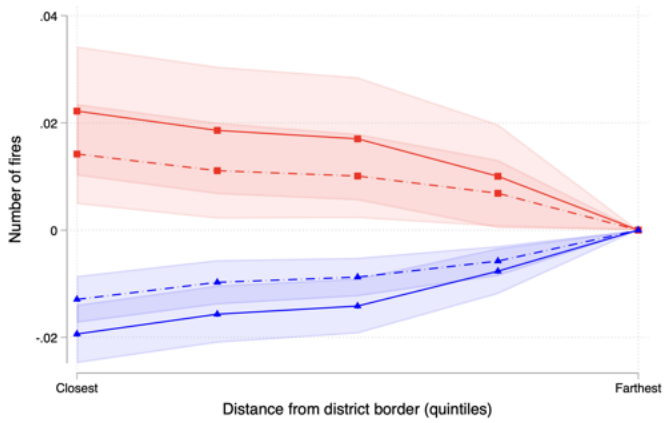
Figure C.1: Effect of distance to the downwind border across countries on fires



Note: (Red) Effect of switching to being in a grid facing the downwind border with another country, where smokes would travel to the other country. (Blue) Effect of switching to being a grid facing an upwind border with another country, where smokes would affect the district in which the grid is located rather than the other country. All effects are plotted by quintiles of distance to the border. The regression includes fixed effects for grid-border pairs and district times month-year and controls for wind speed and direction. Standard errors are clustered at the grid and district times month-year level.

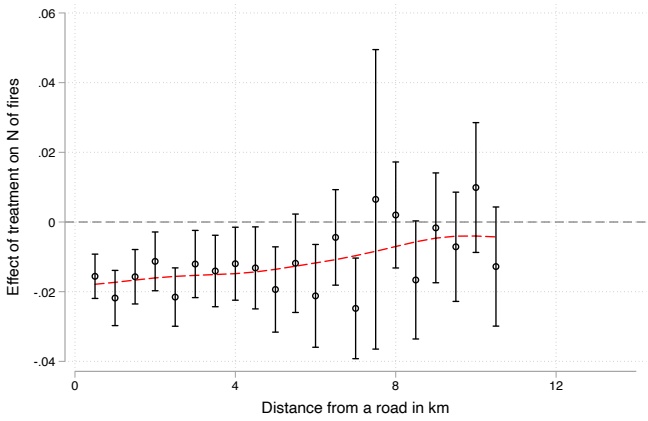
C Additional Results - Administrators Incentives

Figure C.2: Effect of distance to the downwind border across countries on fires in rice producing districts



Note: Coefficients from Figure 9 connected by dotted lines, and coefficients from the same regression but focusing only on rice producing areas connected by solid lines.

Figure D.1: Treatment effect by distance to nearest main road



Notes: Treatment effect by 0.5km distance between the grid centroid and the nearest main road. All regression include grid and district by month fixed effects and control for wind speed and direction. Standard errors are clustered at the grid and month by district level.

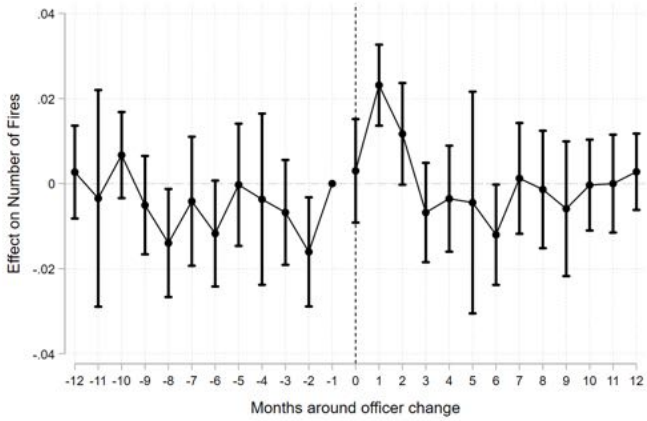
D Additional Results - Administrators Costs

Table D.1: Visibility: effect by winter and summer crops

	(1) N fires	(2) N fires	(3) N fires	(4) N fires	(5) N fires	(6) N fires
Rice	0.00899***					
	(0.00307)					
Rice x Harvest Winter		0.0863***	0.120***			
		(0.00708)	(0.0131)			
Down>Up			-0.0102***			-0.00636***
			(0.00138)			(0.00192)
Rice x Downup			-5.95e-05			
			(0.00327)			
Harvest Winter x Downup			0.000869			
			(0.00584)			
Rice x Harvest Winter x Downup			-0.0668***			
			(0.0209)			
Wheat				0.0196***		
				(0.00231)		
Wheat x Harvest Summer					0.0192***	0.0161**
					(0.00554)	(0.00779)
Wheat x Downup						-0.0108**
						(0.00439)
Harvest Summer x Downup						-0.0213***
						(0.00726)
Wheat x Harvest Summer x Downup						0.00525
						(0.00997)
Observations	17,623,560	17,623,560	17,623,560	17,623,560	17,623,560	17,623,560
District x Month fe	Y	Y	Y	Y	Y	Y
Grid fe	N	N	N	N	N	N
Controls	Y	Y	Y	Y	Y	Y

Notes: The table shows the effect of rice (Columns 1-3) and wheat (Columns 4-6) producing areas on fire occurrence depending on whether we are in the harvest months (October-November for rice, April-May for wheat) and on treatment status. High crop presence is defined as having a larger than median area of the grid devoted to that crop production. All regressions include fixed-effects as in the main table, control for average wind direction and wind speed in each grid and month and standard errors are clustered at the grid and month×province level.

Figure D.2: Treatment effect in the month of administrative turnover



Notes: Event study specification as in Equation 2 considering the effect of months from a bureaucratic turnover depending on whether a grid is treated (is a source of exposure to pollution for the district). Data on bureaucratic turnover in India is collected by the authors. We exclude districts for which we do not possess the full list of bureaucrats (27 districts in Uttar Pradesh, 1 in Haryana and all of Delhi). All regression include grid and district times month-year fixed effects, and control for wind speed and direction. Standard errors are clustered at the grid and month by district level.

E Health Effects, Robustness

Table E.1: Effect of in-utero exposure to fires on infant and child mortality

	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A: Child Mortality			
Fires 2nd tercile	17.21***		28.76*		30.39**	
	(1.611)		(16.90)		(12.15)	
Fires 3rd tercile	16.42***		29.08*		37.48**	
	(1.610)		(17.35)		(14.69)	
IHS Fires		2.462***		1.108		9.528**
		(0.188)		(1.904)		(4.267)
Observations	179,194	179,194	177,480	177,480	67,102	67,102
Time fe	Y	Y	N	N	Y	Y
DHS cell fe	N	N	Y	Y	Y	Y
Mother fe	N	N	N	N	Y	Y
Mean DV	39.34	39.34	39.34	39.34	39.34	39.34
	(1)	(2)	(3)	(4)	(5)	(6)
	(1)	(2)	(3)	(4)	(5)	(6)
			Panel B: Infant Mortality			
Fires 2nd tercile	17.31***		28.76*		25.12**	
	(1.559)		(16.90)		(11.83)	
Fires 3rd tercile	14.54***		28.86*		30.61**	
	(1.553)		(17.32)		(14.28)	
IHS Fires		2.299***		2.170		10.15**
		(0.182)		(1.850)		(4.172)
Observations	178,652	178,652	176,932	176,932	66,582	66,582
Time fe	Y	Y	N	N	Y	Y
DHS cell fe	N	N	Y	Y	Y	Y
Mother fe	N	N	N	N	Y	Y
Mean DV	36.72	36.72	36.72	36.72	36.72	36.72

Note: Effect of in-utero exposure to fires (terciles in odd columns, inverse hyperbolic sine transformation in even columns) on child (Panel A) and infant mortality (Panel B). Fixed effects for birth date (time), location (DHS cell) and for having the same mother are all included in the most restrictive specifications (Columns 3 and 6). Standard errors are clustered at the DHS cell level.