

The Role of Environmental Regulation: Evidence from India's Crop Burning Crisis

Giulia Tozzi

Working Paper No. 996

August 2026

ISSN1473-0278

School of Economics and Finance



The Role of Environmental Regulation: Evidence from India's Crop Burning Crisis

Giulia Tozzi*

Queen Mary University of London

August 3, 2026

[\[Please click here for latest version\]](#)

Abstract

This paper examines the impact of environmental regulation on crop-residue burning, air pollution, and health in Northern India. Exploiting a spatial regression discontinuity at the Punjab–Haryana border and satellite data from 2012–2023, I show that Haryana's implementation of crop-residue burning regulation reduced rice-season fires by about one-third and increased the probability that a location is fire-free by roughly 80% relative to adjacent Punjab. Conditioning $PM_{2.5}$ on local wind regimes, the induced fall in burning lowers burning-season $PM_{2.5}$ on the Haryana side by roughly $9 \mu g/m^3$, about a 10% reduction relative to the local counterfactual just across the border. Applying standard concentration–response functions to the resulting exposure change implies approximately 1,600–2,100 avoided premature deaths per year, of which about one in seven deaths averted occurs among under-five children. These results show that a credibly and consistently enforced regulatory bundle can generate large and persistent reductions in agricultural residue burning, air pollution, and associated mortality, adding an important piece to the policy puzzle of how to confront what has long been regarded as an unsolvable environmental problem.

Keywords: Air Pollution, Crop Burning, Environmental Regulation, Climate Mitigation, Agri-Food Systems, South Asia.

* School of Economics and Finance, Queen Mary University of London (QMUL), Mile End Road, London, E1 4NS. Email: g.tozzi@qmul.ac.uk,
Webpage: <https://sites.google.com/view/giuliatozzi/>

I would like to thank my supervisors Jonathan de Quidt, Lucie Gadenne, and François Gerard. I am grateful to Sebastian Axbard, Marco Manacorda, Simon Franklin, Andrea Tesei, Felipe González, Gabriel Chaves Bosch, Lorenzo Germinetti, Marion Leroutier, Andrea Smurra, Marco Castelluccio, and fellow researchers at the Institute for Fiscal Studies for their help and advice. Thanks to the participants of the QMUL doctoral and departmental seminars, and IFS Environment working group. I gratefully acknowledge funding towards my doctoral studies by QMUL and AMSE. All errors remain my own.

1 Introduction

South Asia, hosting a quarter of the global population, faces a severe air pollution crisis with significant global health implications (Dipoppa and Gulzar, 2024). The region contributes half of the world’s four million annual premature deaths related to air pollution, with 10–25% occurring in India alone (The World Bank, 2022; Lan et al., 2022; WHO, 2023). A central driver is post-harvest residue burning¹ in the dual rice-wheat system, where the turnaround between harvest and sowing can be as short as ten days, making fire the dominant low-cost clearance technology. Roughly one-third of regional air pollution is attributed to emissions from crop-residue burning after the rice and wheat harvests (Abdurrahman et al., 2020). Despite its low private cost, crop-residue burning imposes large transboundary externalities that disproportionately impact regions downwind of agricultural fires, exacerbating public health concerns and contributing to greenhouse gas (GHG) emissions. During peak burning, PM_{2.5} in New Delhi rises to roughly ten times the WHO guideline, and prevailing northwesterlies transport pollution across state and national borders into Pakistan, Nepal, and Bangladesh (Lan et al., 2022; Bikkina et al., 2019; Jethva et al., 2019; Sarkar et al., 2018). Statutory bans exist, but implementation is uneven, raising concerns that residue burning constitutes an *intractable* problem, “sparking a blame game every year” (Dipoppa and Gulzar, 2024; Srinivasan, 2023; Jalali, 2023). More broadly, India reflects a canonical LMIC pattern in which strong *de jure* environmental regulation is undermined by constrained regulatory capacity and weak *de facto* enforcement (Greenstone and Hanna, 2014; Duffo et al., 2013; Jagnani and Mahadevan, 2025). Managing crop-residue burning illustrates a broader agrifood-policy tension in LMICs: a trade-off between food security & agricultural production and local environmental quality. Whether subnational regulation can curb a salient, privately cheap, and spatially mobile source of pollution remains an open, policy-relevant question.

This paper assesses the effectiveness of environmental regulation in India, asking whether state-level policy can reduce crop-residue burning and improve air quality even when damages spill across jurisdictions. The analysis centres on the Punjab–Haryana border, where policy timelines and enforcement diverged in the late 2010s and early 2020s. Together, the two states contribute 60% of the country’s central rice grain pool, with Punjab (48–75%) and Haryana (7.8–14%) consistently among the top contributors to residue-burning-related premature deaths (Lan et al., 2022). At the central level, guidance to reduce crop-residue burning had been in place since 2014–2015, but meaningful technology-adoption incentives were introduced only in 2018, when machinery subsidies were deployed to promote adoption

¹Crop residue burning is an important driver of pollution also in China, Southeast Asia, and Africa. (Jack et al., 2025)

among farmers. Haryana complemented this with a multi-instrument state-level system—initiated in 2016 and substantially expanded in 2018—that combined regulatory orders, on-the-ground implementation, penalty mechanisms, monitoring through district and sub-divisional structures, mobile field squads, and state-operated subsidy programmes. This administrative architecture enabled Haryana to absorb and operationalise the central incentives, aligning the state’s regulatory capacity with the farmer-level technology-adoption incentives made available in 2018. By contrast, a National Green Tribunal (NGT) order dated 2017 recorded that Punjab had not yet filed the required action plan, indicating more limited state-level engagement until the release of its 2022 crop-residue management plan. These institutional differences yield a pre-specified prediction: if enforcement matters, a discontinuity in burning should emerge at the border beginning in 2018. The spatial RD design exploits this contrast, focusing not on the effect of individual instruments but on whether a coordinated regulatory bundle reduces burning when agronomic and climatic conditions remain broadly comparable across the border.

For this purpose, I assemble a new, high-resolution geospatial panel linking satellite-detected fires to 2 km $\times$ 2 km agricultural and air quality information from 2012 to 2023. Fire outcomes come from daily 375 m VIIRS detections summarised by agronomic season (Kharif/rice and Rabi/wheat). Air pollution is measured using two daily 1 km satellite-based PM_{2.5} products: the *Global High Air Pollutants* (GHAP) dataset and the *Long-term Gap-free High-resolution Air Pollutants* (LGHAP) dataset. Wind comes from ERA5-Land 10-metre zonal and meridional components at $\sim$ 9 km resolution. Using these components, I classify daily meteorological regimes (e.g., near-calm and border-parallel winds) under which plumes remain local and cross-border transport is limited. For each regime, I construct seasonal mean PM_{2.5} during the rice-burning months. To assess comparability across the border prior to the policy, I draw on micro-geographic information from the 2011 Socio-Economic and Caste Census (SECC) and FAO Global Agro-Ecological Zones (GAEZ) 2010 datasets.

The empirical strategy is a spatial regression discontinuity (sRDD) that compares locations immediately inside Haryana to locations immediately inside Punjab. Given that I have multiple dependent variables, I determine the optimal bandwidth for each following the methods outlined in Calonico et al. (2014): I use an optimal 15 km bandwidth for burning outcomes and a 40 km bandwidth for pollution outcomes. The specification includes side-specific polynomials in distance to the boundary, excludes border-straddling cells (a “donut”), and absorbs local geography with border-segment fixed effects. Uncertainty is quantified with two-way spatial clustering on two overlapping 100 km² grids (Burgess et al., 2023). I estimate effects at both the extensive (any fire) and intensive (number of fires) margins and allow treatment intensity to vary over time with Haryana’s regulatory rollout.

I then apply the same design to harmonised PM_{2.5} under different wind regimes and combine exposure reductions with concentration-response functions to quantify avoided deaths, disability-adjusted life years (DALYs), and monetary benefits.

Three main sets of findings emerge. First, on the behaviour of farmers, Haryana’s enforcement materially changes crop-residue management practices. During the 2018–2021 “Haryana treated” period, cells on the Haryana side of the border experience roughly 2 fewer rice-residue fires per Kharif season relative to adjacent Punjab cells, a reduction of about one-third relative to the Punjab-side mean in the same period. At the extensive margin, the probability that a cell is completely fire-free rises by about 14 percentage points from a 17% baseline on the Punjab side, implying roughly an 80% increase in the likelihood of zero burning. Second, conditioning PM_{2.5} on wind regimes that restrict cross-border transport, the decline in burning yields about a 9 $\mu\text{g}/\text{m}^3$ reduction in burning-season pollution on the Haryana side, roughly a 10% drop relative to the local counterfactual just across the border. Annualised, this corresponds to a long-run reduction of approximately 1.5 $\mu\text{g}/\text{m}^3$ in average annual exposure. Third, applying standard concentration-response functions (Ayres, 2010; Chen and Hoek, 2020; Dipoppa and Gulzar, 2024) to these exposure changes implies approximately 1,600–2,100 avoided premature deaths per year in Haryana, including about 187 avoided infant deaths and 272 avoided under-five deaths. Expressed in disability-adjusted life years, the gains correspond to about 0.7 additional healthy days per resident annually and roughly 10.4 and 3 additional healthy days per year for the average infant and under-five child, respectively. Valued using a conservative Value of a Statistical Life (VSL) for India, these mortality reductions amount to approximately 1.10–1.45 billion USD per year. Under DALY-based valuations, the gains imply aggregate benefits of approximately 110–340 million USD per year, corresponding to about 170–520 USD per acre of residue burning avoided under WHO–CHOICE benchmarks. Balance tests from SECC 2011 and FAO-GAEZ 2010 indicate no pre-policy discontinuities in socio-economic or agro-ecological characteristics. Placebo tests outside harvest months and in the wheat season show no systematic discontinuities. Together with the fire and pollution results, these findings document that Haryana’s regulatory effort produced measurable improvements in environmental and health outcomes despite substantial cross-jurisdiction externalities.

This paper contributes to three strands of research. First, this paper speaks to the literature on the effectiveness of environmental regulation in low- and middle-income countries. Existing work shows that regulatory reforms can, in some cases, reduce emissions and improve health, but it also documents many instances in which ambitious *de jure* rules translate into only modest or no *de facto* improvements because monitoring is weak, enforcement is corruptible, agency design limits oversight, and externalities spill across jurisdictions (Davis, 2008; Duflo et al., 2013; Greenstone and Hanna, 2014; Oliva, 2015; Lipscomb and

Mobarak, 2017; Fenske et al., 2023). I add to this literature by showing that a tightly targeted state-level regime in a lower-middle-income setting can, when actively enforced, generate large and precisely estimated reductions in a salient emission source, lower seasonal $\text{PM}_{2.5}$ by roughly 10% at the border, and avert substantial mortality and DALY losses—implying that weak *de facto* enforcement, rather than the absence of regulation, is a central constraint to cleaner air in this context.

Second, the analysis contributes to the growing research on agricultural residue burning. Prior studies document the scale and upward trend of crop-residue burning in South Asia, its contribution to extreme seasonal air pollution and health damage, and its links to mechanisation, labour-market frictions, and agricultural policy, while a recent literature finds that financial incentives and appeals to prosocial motives typically yield limited and context-specific reductions in burning (Abdurrahman et al., 2020; Lan et al., 2022; Liu et al., 2020; Garg et al., 2024; Behrer A. Patrick, 2023; He et al., 2020; Rangel and Vogl, 2019; Falcon et al., 2022; Nian, 2023; Jack et al., 2025; Jagnani and Mahadevan, 2025; Dipoppa and Gulzar, 2024; Khundrakpam and Sarmah, 2023; Venkatesh et al., 2017; The World Bank, 2022; Srinivasan, 2023). The closest related paper in this subliterature is Dipoppa and Gulzar (2024), who use wind-driven variation in pollution transport across districts in India and Pakistan to show that local administrators exert more effort against crop fires when smoke affects their own jurisdictions, thereby reducing burning and child mortality and highlighting the central role of bureaucratic incentives in shaping *de facto* enforcement. Their results highlight that existing bureaucrats, when properly incentivised, can improve environmental management and public health outcomes.

My analysis is complementary and focuses instead on the design of the regulatory regime itself. I provide quasi-experimental evidence that a comprehensive bundle of penalties, monitoring, and technology subsidies in Haryana induced a large and sustained decline in rice-residue fires, reduced wind-conditioned $\text{PM}_{2.5}$ exposure, and generated sizeable implied health gains, demonstrating that crop burning can be curbed at scale when state enforcement is credibly aligned with farmers’ incentives.

Third, methodologically, the paper contributes to the literature that exploits geographic borders to identify the causal effects of policy and institutional discontinuities. Existing “border designs” compare locations on different sides of national or subnational boundaries to estimate the impact of tax regimes, labour-market policies, natural-resource regulation, or national institutions—such as property-rights enforcement and state capacity (Holmes, 1998; Magruder, 2013; Chen et al., 2013; Michalopoulos and Papaioannou, 2014; Lipscomb and Mobarak, 2017; Pinkovskiy, 2017; Burgess et al., 2023). I extend this approach by developing a spatial regression-discontinuity design along a state border in a rural setting and combining it with high-resolution satellite data on fires and wind-conditioned grid-

ded PM_{2.5}, thereby isolating the contribution of subnational environmental regulation to burning, air quality, and health in a context where pollution and regulatory authority are sharply misaligned in space.

Taken together, the existing evidence points to several distinct margins along which the state can act on what is widely portrayed in the media and policy debate as an intractable problem. This paper adds a complementary piece to the literature portfolio by showing that well-targeted and actively enforced subnational environmental regulation can substantially reduce burning and pollution at scale in a core food-producing region. The combined evidence implies that even highly diffuse agricultural externalities are amenable to policy when instruments targeting farmers (Jack et al., 2025), bureaucrats (Dipoppa and Gulzar, 2024), and regulatory design are deployed in a coordinated way. By documenting that a credibly enforced regulatory bundle can deliver large and persistent reductions in agricultural residue burning and exposure, this paper adds a further piece to the policy puzzle of how to confront what has long been treated as an unsolvable environmental problem.

2 Context

2.1 Structural Drivers of Crop Residue Burning

The practice of crop residue burning in Northern India is deeply rooted in the economic and agronomic realities of the region. As one farmer summarised, “If I can clear my farm using a one-rupee matchbox, why will I spend thousands?” (Venkatesh et al., 2017). Dipoppa and Gulzar (2024) documents several structural factors that explain the persistence of crop residue burning in Northern India. First, the Green Revolution² reinforced an irrigated rice-wheat rotation and thereby compressed the harvest-sowing calendar. Rice is cultivated in the Kharif season (June–October) and wheat in the Rabi season (November–April), a sequence that exploits local agro-climatic conditions but leaves limited time between crops. Second, groundwater-conservation legislation further compressed the rice harvest period, increasing burning by nearly 40%. The narrow time window between rice harvesting and wheat sowing—averaging just 13 days—leaves farmers with limited alternatives to burning. In contrast, the longer gap between wheat harvesting and rice planting (46–48 days) results in significantly less residue burning—in Punjab, 90% of farmers report burning crop residues after the rice harvest, compared with 11% after the wheat harvest (Kumar and Joshi, 2015; Garg et al., 2024). Third, greater reliance on mechanised harvesting generated larger volumes of loose straw, raising the cost of non-burning residue-management

²Adoption of high-yielding varieties (HYVs) and complementary inputs was particularly rapid in the northern states of India, supported by canal irrigation and relatively uniform growing conditions that made early HYVs well suited to the region (Bharadwaj et al., 2020).

alternatives. A range of non-burning options is available in principle. In-situ management practices mulch or incorporate straw into the soil using implements such as disc ploughs, disc harrows, rotavators, zero-tillage seed drills and the Happy Seeder (Kumar and Joshi, 2015; Government of India and Ministry of Agriculture & Farmers Welfare, 2019, 2021). Other alternatives (ex-situ) include composting and the collection of straw for use as fodder or as feedstock for biomass and industrial uses. In the prevailing rice-wheat system, however, these practices require additional machinery, fuel and labour within an already compressed harvest-sowing window, so that open burning typically remains the privately cost-minimising choice for farmers. Finally, rising agricultural wages due to sustained rural out-migration further reinforce the adoption of mechanised tools and the resulting residue loads.

2.2 Government and State-Level Regulation

Efforts to curb crop residue burning in India reveal a marked divergence between the national regulatory framework and heterogeneous state implementation. Since 1991, agricultural burning has been criminalised under the Penal Code and the Air (Prevention and Control of Pollution) Act, and the 2014 National Policy of Crop Residue assigned states the responsibility to prevent and ban the practice. In December 2015, the National Green Tribunal (NGT) issued binding prohibitions on crop burning in Delhi, Uttar Pradesh, Rajasthan, Punjab, and Haryana, coupled with a penalty schedule (INR 2,500–15,000) based on landholding size. Enforcement, however, remained limited due to political resistance and scarce local monitoring capacity.

In 2018, the central government introduced a major technology-based intervention—the Central Sector Scheme on “Promotion of Agricultural Mechanization for In-Situ Management of Crop Residue”—which subsidized 50–80% of the cost of residue-management machinery for farmers and cooperatives.³ The 2018 Central Sector Scheme for In-Situ Crop Residue Management provided substantial farmer-level support through capital subsidies for machinery and the establishment of Custom Hiring Centres (CHCs) and represented the first substantial farmer-level incentive designed to accelerate technological adoption nationwide. According to Government of India and Ministry of Agriculture & Farmers Welfare (2019), for the 2018–19 financial year, Punjab received INR 269.38 crore (approximately USD 39 million) and Haryana INR 137.84 crore (approximately USD 20 million) in central funds.⁴ Figure 1 illustrates, respectively, the sequence of central policies and the timing of

³Government of India and Ministry of Agriculture & Farmers Welfare (2021)

⁴The allocations covered the establishment of CHCs (INR 176 crore in Punjab; INR 72 crore in Haryana), the distribution of in-situ machinery (INR 71.3 crore and INR 41.5 crore, respectively), information, education and communication campaigns (INR 16.8 crore and INR 21.64 crore), and administrative and

state-level actions in Haryana and Punjab.

State regulation in Haryana. Beginning in 2016, Haryana began operationalising the national framework and independently created a district-level enforcement architecture under the National Policy for the Management of Crop Residue.⁵ To ensure compliance with NGT and High Court directions, the state (i) constituted District Level Committees with authority to operationalise judicial directives; (ii) anticipated the 2018 Central Sector Scheme by introducing technology-based incentives covering machinery costs (free for < 2 acres; INR 5,000 for 2–5 acres; INR 15,000 for > 5 acres); and (iii) made machinery and equipment available at government cost, to be deployed at identified sites each growing season. The district-level structure established in 2016 was implemented that same year, whereas the analogous committee in Punjab was envisaged only in 2022 and had not been created at that time.

In 2018—the year in which substantial central subsidies became available—Haryana complemented the national scheme with an expansion of its administrative capacity.⁶ The state constituted a State Level Monitoring Committee (SLMC) under the Chief Secretary, bringing together senior officials from multiple departments to coordinate residue-management activities and assess district performance. At the same time, Sub-Divisional Magistrates (SDMs) were designated as nodal officers responsible for supervision, monitoring, and regulatory enforcement at the sub-division level, creating an intermediate administrative tier directly accountable for implementation. The state also introduced mobile squads in every tehsil/block, operating under SDM authority, to conduct field inspections, detect active burning, and initiate immediate action. This consolidation linked the central farmer-level machinery incentives with state-level oversight (SLMC), sub-division command (SDMs), and block-level monitoring (mobile squads), thereby creating an integrated enforcement structure capable of absorbing and operationalising the 2018 technological push.

From 2020 onward, Haryana launched the *Mera Pani Meri Virasat* scheme, providing INR 7,000 per acre to farmers who diversified more than 50% of their Kharif paddy area, thus altering incentives around the rice-wheat rotation. By 2021, the state complemented this with additional in-situ and ex-situ subsidy programmes.

State regulation in Punjab. Punjab’s policy trajectory contrasts with this timeline. The National Green Tribunal (2017) judgment of 18.12.2017 noted that “The State of Punjab has not even filed an action plan as contemplated under the orders of the Tribunal,” and that the state had instead submitted an affidavit committing to installing ambient air

flexi-fund components.

⁵Government of Haryana (2016), Order on Constitution of District Level Committee for Implementation of NGT/High Court Directions on Crop Residue Burning, Environment Department, 5 October 2016.

⁶Government of India and Ministry of Agriculture & Farmers Welfare (2019)

quality monitoring stations. Punjab’s district-level enforcement committee remained unimplemented as of 2022, despite its inclusion in the National Policy. The Punjab Pollution Control Board issued an order in 2018 concerning harvesters operating without Super-SMS equipment, and in September 2020 authorised prosecutions under the Air (Prevention & Control of Pollution) Act. These measures remained limited in scope and did not establish an operational framework comparable to Haryana’s multi-tier system. Punjab’s first comprehensive Crop Residue Management Plan was introduced in May 2022, almost six years after Haryana had instituted its district-level administrative architecture.

While the central government provided a uniform regulatory framework, state-level implementation evolved differently. Haryana translated the *de jure* mandate into a multi-instrument enforcement system from 2016 onward, with a major administrative escalation in 2018 that complemented the introduction of farmer-level technology-adoption incentives at the central level. From 2020–2021, additional state incentives complemented this architecture. Punjab, by contrast, had not operationalised the required enforcement structures by 2022, as indicated by the NGT, and undertook more comprehensive actions only in that year. This divergence generates a sharp quasi-experimental setting for assessing the extent to which state-level environmental regulation can reduce pollution even in the presence of cross-border externalities.

3 Data

3.1 Active Fire

The primary source of fire data used in this study is the VIIRS 375-m Active Fire dataset, spanning from 2012 to the present. This dataset provides daily observations with a spatial resolution of 375 metres and identifies fires by measuring how much hotter each pixel is relative to its surroundings. The sensor uses two heat-detecting infrared channels: one that is very sensitive to extreme temperatures (such as open flames) and one that captures normal land surface temperature. A pixel is classified as an active fire when it is significantly hotter than the natural variation observed in neighbouring pixels, after removing cloud-covered areas, water bodies, and other potential sources of false signals. This contextual comparison makes the VIIRS product effective at identifying small and short-duration agricultural fires typical of crop residue burning. In comparison to the MODIS dataset, frequently used in the literature (for example by Dipoppa and Gulzar (2024)), the higher resolution of VIIRS (375 m vs 1 km) enhances the detection of fires over relatively small areas, improves the mapping of large fire perimeters, and demonstrates superior performance in nighttime fire detection. These features make VIIRS a robust tool for analysing fire events with precision.

Figure 2 plots the spatial distribution of fires during rice-residue burning seasons within a 30 km band of the border.

3.2 Satellite-Based Measures of Air Pollution and Wind

Two complementary remote-sensing datasets are employed to measure surface-level fine particulate matter (PM_{2.5}): the *Global High Air Pollutants* (GHAP) dataset (Wei et al., 2023) and the *Long-term Gap-free High-resolution Air Pollutants* (LGHAP) dataset (Bai et al., 2024). These sources provide global daily estimates of PM_{2.5} at a 1 km spatial resolution and share a common methodological foundation linking satellite retrievals of atmospheric aerosols with ground-level particulate concentrations. The empirical strategy in this paper is a spatial regression discontinuity (RD) design that compares adjacent 2×2 km cells straddling the Punjab–Haryana administrative border during periods of intense agricultural residue burning. Identification therefore relies on short-run, highly localised spikes in PM_{2.5} that can differ discontinuously across the state line and dissipate within a few days. This creates two core data requirements: (i) high temporal resolution, in order to align the timing of burning events with contemporaneous pollution; and (ii) high effective spatial resolution, in order to detect pollution discontinuities across borders a few kilometres apart. The pollution data described below meet both requirements by providing daily PM_{2.5} at 1 km resolution, with a modelling approach designed to capture short-term local plumes.

Global High Air Pollutants (GHAP). The GHAP dataset combines multi-satellite measurements of *Aerosol Optical Depth* (AOD)—from MODIS, MISR, and VIIRS sensors—with physical information from chemical transport and meteorological models (e.g., GEOS-Chem, MERRA-2). AOD captures the total amount of aerosol in the vertical atmospheric column rather than near-surface pollution. GHAP converts AOD to ground-level PM_{2.5} using model-based vertical profiles, humidity corrections, and boundary-layer height estimates, followed by a machine-learning bias correction trained on ground monitoring stations worldwide. The result is a globally harmonised, daily, gap-free estimate of PM_{2.5} concentrations available for the period 2017–2022 that performs particularly well in heavily polluted regions such as Northern India.

Long-term Gap-free High-resolution Air Pollutants (LGHAP). LGHAP follows a similar fusion strategy but emphasises temporal continuity and the imputation of missing satellite observations caused by cloud cover, high surface reflectance, or scan gaps. A spatio-temporal data fusion algorithm reconstructs missing AOD observations using auxiliary meteorological variables (temperature, humidity, wind speed, boundary-layer height)

and statistical interpolation. LGHAP thus guarantees full spatial and temporal coverage at 1 km for 1998–2021. Bias corrections are also applied using ground-based measurements, though with greater emphasis on maintaining a smooth, internally consistent long-term time series rather than optimising short-run accuracy.

Comparative advantages and integration. In the overlapping period 2017–2021, both datasets provide high-quality $\text{PM}_{2.5}$ information. GHAP is optimised for precision in recent years, while LGHAP ensures continuous coverage over the longer horizon. To exploit these strengths, the analysis uses GHAP as the main source for 2017–2022 and LGHAP for 2015–2016. A harmonisation procedure aligns LGHAP with GHAP’s scale over the overlap period (2017–2021), ensuring that all pollution data are expressed in comparable units throughout 2015–2022.

ERA5 Post-Processed Daily Statistics on Single Levels. To measure surface wind conditions, I rely on the *ERA5-Land Derived Daily Statistics on Single Levels* dataset produced by the European Centre for Medium-Range Weather Forecasts (ECMWF) through the Copernicus Climate Data Store (Copernicus Climate Change Service, 2024). ERA5 is a global meteorological reanalysis that assimilates a wide range of satellite and ground-based observations into a physically consistent numerical weather model, providing hourly estimates of atmospheric and land-surface variables from 1940 onward, at a spatial resolution of 9 km. The derived daily product used here is a post-processed version of the core ERA5 reanalysis that aggregates hourly data into daily means, minima, and maxima at selected “single levels” near the Earth’s surface. Unlike variables defined on pressure levels (which describe the vertical structure of the atmosphere), single-level products represent quantities measured at fixed reference heights or directly at the surface, such as temperature at 2 m, wind at 10 m, or surface pressure.

The 10 m wind variables are particularly suited for this study because they describe the horizontal wind field immediately above the ground layer where crop burning and pollution dispersion occur. ERA5 provides two orthogonal wind components: the zonal component (u_{10}) representing east-west air movement and the meridional component (v_{10}) representing north-south movement. From these, I compute daily wind speed and direction at each grid point and merge the data with the high-resolution $\text{PM}_{2.5}$ grid used in the spatial regression discontinuity design.

This dataset is appropriate for my setting for three main reasons. First, its daily temporal resolution matches the frequency of burning events and short-lived pollution peaks, allowing an exact alignment between wind and $\text{PM}_{2.5}$ observations. Second, its ~ 9 km spatial resolution balances meteorological realism with computational tractability while remaining

fine enough to characterise local flow patterns around the Punjab–Haryana border. Third, the ERA5 reanalysis provides physically coherent and bias-corrected estimates even where local meteorological monitoring is sparse, which is crucial in agricultural regions of Northern India.

3.3 Agricultural Data

To benchmark agricultural fundamentals and pre-policy conditions around the Punjab–Haryana border, I draw on two complementary sources. The first is the 2011 Socio-Economic and Caste Census (SECC) (Ministry of Rural Development, Government of India, 2011; Asher et al., 2021), which provides a complete enumeration of household characteristics and assets and therefore speaks to the socio-economic conditions of agricultural households—land ownership, income sources, and access to equipment. The second is FAO’s Global Agro-Ecological Zones (GAEZ) “Guides” layers for rice and wheat circa 2010, which summarise the agronomic environment and production structure through model-based suitability indices and downscaled measures of harvested area and implied yields. Taken together, the SECC captures the wealth, production capacity, and technology available to households, while GAEZ characterises the agro-ecological backdrop and the scale of rice–wheat cultivation at local level immediately prior to my study period (2012–2023). Throughout, I summarise covariates on either side of the border within the optimal bandwidths selected following Calonico et al. (2014)—40 km and 50 km for SECC and FAO-GAEZ respectively—and I report baseline tables computed within those border corridors.

Socio-Economic Census Data and Agricultural Covariates The SECC is a complete enumeration of households with variables on land ownership, primary income source, dwelling characteristics, and ownership of agricultural and irrigation equipment. I use: (i) the share of households that *own land* (landholding structure); (ii) the share whose *primary income is cultivation* (reliance on farming); (iii) the *mean number of rooms* per dwelling (a coarse wealth proxy); (iv) *land with assured irrigation for two crops per household* (feasibility of the rice-to-wheat cycle); and (v) the shares owning *mechanized agricultural equipment*, *irrigation equipment*, and *general agricultural equipment* (capital intensity and potential substitutes for burning).

SECC variables are reported for irregular “SHRID” polygons that do not align with my analysis grid. Using area-weighted interpolation, I overlay SHRIDs with the 2×2 km grid and allocate counts/totals to cells via area-weighted shares; for ratios or means, I reconstruct cell numerators and denominators and re-compute the ratio to preserve additivity. The

resulting cell-level covariates are then averaged within the RD border corridors on each side. The baseline balance table based on these SECC variables is reported in Table 3.

FAO Global Agro-Ecological Zones (GAEZ) For rice and wheat circa 2010, I use three GAEZ layers at 9 km spatial resolution. First, *crop suitability indices* (0–10,000) produced by biophysical crop models that combine climate, soils, terrain, and standardized input/water assumptions. Second, *harvested area* (km²), obtained by statistically down-scaling FAOSTAT harvested-area totals to grids. Third, *implied average yields* (tons/km²), computed as downscaled production divided by downscaled harvested area. These layers represent realised production totals at higher administrative levels but their within-unit spatial allocation is model-based.

Unlike SECC, GAEZ is already on a regular raster. I therefore use the raw grid values without further interpolation. For rice, I use mean harvested area and implied yield; I construct analogous measures for wheat. I also include the rice and wheat suitability indices to characterise the agro-ecological backdrop. Harvested area is reported in km²; yields in tons per km²; suitability is unitless (0–10,000). The baseline summary for these variables appears in Table 4.

Limitations and role in identification. Because GAEZ suitability and the spatial allocation of “actual” production are model-generated, these layers should vary smoothly across space and are not expected to exhibit sharp discontinuities exactly at an administrative border. I therefore rely on SECC covariates to test pre-policy balance in *observed* household conditions and capital, and I use GAEZ variables principally to document fundamentals and control for slow-moving geography. A further limitation is temporal: both SECC (2011) and GAEZ (2010/11) pre-date the policy period by five years, and updated census-style information is not available for a later date. While this prevents observing potential small drifts in socio-economic or agro-ecological conditions immediately before enforcement, these data remain, to the best of my knowledge, the only high-resolution pre-period sources available for structural conditions. Taken together, this division of roles ensures that any post-2016 discontinuities in crop-residue burning and downwind PM_{2.5} are not artefacts of pre-existing differences in agricultural potential or household socio-economic structure, while avoiding over-interpretation of model-based or temporally limited raster products for cross-border breaks.

4 Effects of Environmental Regulation on Crop Burning Practices

This section examines how Haryana’s regulatory effort affected crop-residue burning relative to Punjab across the 2012–2023 period. The Punjab–Haryana border offers a sharp institutional contrast embedded within a largely homogeneous agro-ecological setting.

The empirical analysis relies on a high-resolution geographic panel in which the unit of observation is a 2×2 km cell (4 km^2), observed from 2012 to 2023. Each cell is observed annually during the *rice-residue burning season*, defined as September through November, when agricultural fires stem from post-harvest residue clearance in the rice–wheat system. Agronomically, the most intense burning occurs between October and November (Kumar and Joshi, 2015; Garg et al., 2024). For every cell c and year t , I construct two outcomes: the total number of VIIRS-detected fires in that season (the intensive margin) and an indicator equal to one if the cell records zero fires (the extensive margin).

Figure 2 visualizes burning within a 30 km band of the border, revealing sharp cross-border differences in fire intensity following Haryana’s regulatory ramp-up. What differentiates the two states over time is not their agricultural systems but their regulatory environments. Across the late 2010s and early 2020s, Haryana and Punjab followed different institutional trajectories. Critically, the timing of policy adoption is staggered. Haryana introduced regulatory orders, monitoring infrastructure, and district-level committees beginning in 2016, and scaled up enforcement and machinery subsidies after 2018, aligning state capacity with the subsidy-based central 2018 technology-adoption program. Punjab, by contrast, delayed the adoption of a coordinated residue-management plan until 2022. These sequences, detailed in Figure 1, motivate a pre-specified temporal segmentation of the analysis.

1. **Pre-treatment (2012–2015):** a period in which national prohibitions exist but neither state enforces them in practice.
2. **Early Haryana enforcement (2016–2017):** Haryana begins implementing district-level committees and monitoring and regulatory orders, while Punjab’s enforcement remains minimal.
3. **Haryana treated (2018–2021):** Haryana scales up penalties, monitoring, and machinery subsidies, whereas Punjab continues to lag.
4. **Both states treated (2022–2023):** Punjab adopts and starts implementing a statewide crop-residue management plan, narrowing the regulatory gap.

The combination of year-specific estimates and policy-period averages allows the analysis to capture both the timing and the consolidation of regulatory effects. The year-by-year design provides a transparent test of when discontinuities first arise, while the period-specific means quantify the average effect during distinct regulatory regimes.

The empirical strategy therefore distinguishes between *year-by-year* identification and *policy-period* aggregation. In the first step, I estimate a separate spatial regression discontinuity coefficient β_t for each year. This yields a dynamic series of border discontinuities that reveals when, if at all, the impact of Haryana’s regulatory regime emerges. The second step moves the attention from the yearly coefficients to the four policy periods that correspond to the institutional chronology documented above.

4.1 Empirical Strategy: Spatial Regression Discontinuity Design

To isolate the effect of Haryana’s regulatory policy, I exploit the administrative boundary dividing Haryana and Punjab and estimate a spatial regression discontinuity design (sRDD). The empirical premise, consistent with the balance analysis in Tables 3 and 4 and continuity checks in Figures 10 and 11, is that agronomic and socio-economic conditions vary smoothly across the border, while regulatory enforcement does not. Under these conditions, differences in outcomes at the boundary can be attributed to differences in state-level policy regimes.

Formally, for each year t , the following specification is estimated:

$$Y_{c,t} = \alpha_t + \beta_t \text{Haryana}_c + \lambda_{1t} D_c + \lambda_{2t} D_c \times \text{Haryana}_c + \gamma_{s(c)} + \varepsilon_{c,t}, \quad (1)$$

where $Y_{c,t}$ denotes the burning outcome for cell c in year t ; Haryana_c is an indicator equal to one if the cell lies in Haryana; D_c is the signed distance (in km) to the border (normalised to zero at the boundary and positive into Haryana); the interaction $D_c \times \text{Haryana}_c$ allows for different distance trends on the two sides, implementing a side-specific local-linear polynomial in the sense of Gelman and Imbens (2019). Cells intersecting with the border are excluded via a donut-hole correction; and $\gamma_{s(c)}$ are border-segment fixed effects controlling for fine-grained geographic heterogeneity along the boundary. The estimation window uses the optimal bandwidth selected following Calonico et al. (2014)—this yields $h_t = 15$ km on each side of the border in the baseline specification. Regressions use a uniform kernel, assigning equal weight to all observations within the estimation window.⁷ Standard errors are clustered using two-way overlapping spatial clustering on two $10 \text{ km} \times 10 \text{ km}$ grids to

⁷I use a uniform kernel to assign equal weight to all cells within the narrow estimation window. Given the discrete spatial grid and the exclusion of border-intersecting cells, only a limited number of observations lie near the cutoff; uniform weighting avoids placing disproportionate leverage on a handful of near-boundary cells. However, as shown in Section 4.2, results are robust to using a triangular kernel.

allow for flexible spatial correlation structures (Cameron et al., 2011; Burgess et al., 2023). Practically, following Burgess et al. (2023), I construct two partially offset grids—each of $10 \text{ km} \times 10 \text{ km}$ —that tessellate the study region. Each grid defines a set of spatial clusters, allowing residuals to be arbitrarily correlated within any cell of either grid but only weakly correlated across the two grids. The offset ensures that inference does not hinge on an arbitrary tiling of space and that clusters cross the border in different ways, mitigating concerns that the clustering scheme could mechanically align with the treatment boundary. This approach is well suited to spatial RD settings in which treatment varies sharply at the boundary while unobservables co-move smoothly in space and within seasons.

This design yields a sequence of yearly discontinuities β_t , which trace the evolution of the Haryana–Punjab gap as their institutional contexts diverge. To summarise policy effects, I hold the specification fixed and re-estimate the same RD regression on period-specific horizons corresponding to the regulatory phases defined earlier: pre-policy (2012–2015), transition (2016–2017), Haryana treated (2018–2021), and both states treated (2022). Specifically, for each phase, I restrict the sample to the cell-year observations falling within the relevant calendar years and run the identical border-RD specification. Across both approaches, identification leverages local comparisons, ensuring that any discontinuity at the border is consistent with institutional differences at the border, conditional on smooth structural or agronomic outcomes in D_c .

4.2 Main Results

4.2.1 Graphical Evidence

The graphical evidence uses both binned RD plots for selected years (Figure 3) and an event-study-style plot of the annual RD coefficients (Figures 4 and 5). The year-by-year plots provide a visual assessment of continuity before treatment, the onset of divergence when Haryana’s enforcement becomes operational, and the subsequent evolution of the discontinuity as Punjab begins to implement its residue-management plan.

The graphical analysis begins with the binned RD plots in Figure 3, which illustrate the spatial discontinuity in burning intensity across the Punjab–Haryana border for four representative years—2016, 2018, 2020, and 2022—corresponding to distinct phases of the policy timeline. In all plots, the running variable is D_c , the signed distance to the Punjab–Haryana boundary, defined as zero at the border and increasing positively into Haryana (south of Punjab). These plots provide a non-parametric view of how burning activity evolves in response to changes in the regulatory environment. In 2016, when both states were effectively under the same policy environment, the fitted lines on either side of the border are nearly collinear and there is no discontinuity. Average burning levels remain above 5 fires

per 4 km² on both sides of the boundary, and the Haryana–Punjab difference is small and visually indistinguishable from zero. This confirms the absence of pre-policy divergence and supports the validity of the RD design. In 2018, coinciding with Haryana’s expansion of enforcement capacity, penalties, and technology-based support, the plot reveals a clear downward break. Punjab-side cells adjacent to the border continue to exhibit mean fire counts between 4 and 6, whereas predicted levels on the Haryana side fall between 2 and 4. The visual discontinuity is on the order of 1 to 2 fewer fires per 4 km² cell, indicating the onset of differential enforcement. The discontinuity becomes largest in 2020, during the mature enforcement period. Punjab cells at the border display mean fire counts near 6 to 7, while fitted values on the Haryana side drop below 2. The magnitude of the break is visually close to 4 fires per cell, representing a substantial suppression of burning on the Haryana side. The slope of the fitted line within Haryana also flattens, reflecting a broad-based reduction in burning rather than a localised suppression at the boundary. By 2022, the first year in which Punjab begins implementing its own crop-residue management plan, the discontinuity attenuates. Punjab-side means fall sharply, while Haryana’s predicted levels are slightly lower than those observed in 2020–2021. The break in this year is closer to 2 fires per cell, consistent with the narrowing of the regulatory gap between the two states. These RD plots reveal not only the existence of a discontinuity but also its timing and magnitude in correspondence with the evolution of state-level institutional capacity.

The year-by-year evolution of the RD coefficient is summarised in Figure 4, an event-study-style plot reporting the estimated discontinuity for every year from 2012 to 2023. This figure aligns each coefficient with the corresponding policy environment: the “Same Policy Environment—No Treatment” period (2012–2015), the “Haryana Starts to Be Treated” phase (2016–2017), the “Haryana Treated” years (2018–2021), and the “Haryana & Punjab Both Treated” period (2022–2023). The plot shows a sequence of near-zero coefficients during 2012–2015, consistent with the absence of differential enforcement. This pattern indicates that, prior to Haryana’s policy rollout, there were no systematic differences in burning at the boundary, and if anything the number of fires was slightly higher on the Haryana side of the border. In 2016–2017, the coefficients remain small, reflecting the modest scale of Haryana’s initial district-level efforts. Starting in 2018, the coefficients shift sharply downward in a manner that is both abrupt and temporally aligned with the state’s expanded regulatory architecture. In 2019, the magnitude increases to roughly minus 1.5 fires, and in 2020 the coefficient reaches its peak magnitude of approximately minus 3.5 to minus 4 fires per cell. The discontinuity remains markedly negative through 2021 and then attenuates in 2022–2023, coinciding with Punjab’s delayed but substantive entry into coordinated enforcement. The per-capita RD estimates,⁸ shown in Figure 5, mirror

⁸Population count data at 100m × 100m spatial resolution are obtained from Stevens et al. (2015);

this temporal pattern, reinforcing that the observed discontinuity is not an artifact of differential settlement patterns but corresponds to spatial variation in burning activity induced by differing regulatory systems. Together, these graphical patterns demonstrate visually what the regression results later quantify: regulation-induced divergence emerges when the institutional environment changes, grows during the period of asymmetric enforcement, and narrows again once Punjab adopts comparable measures.

4.2.2 Policy-Period Estimates: Intensive and Extensive Margins

Tables 1 and 2 formalise the graphical patterns by averaging the yearly RD coefficients within the four policy periods identified earlier. These estimates derive from the canonical spatial RD specification with border-segment fixed effects and the optimal bandwidth. This approach aligns the temporal structure of the identification strategy with the documented policy timeline, yielding average effects for periods of stable institutional treatment. The table reports estimates from the canonical sRDD specification, with Haryana = 1 capturing the discontinuity at the state boundary during each of the four periods: “Pre 2016” (2012–2015), “2016–2017,” “2018–2021,” and “Post 2022.” The intensive and extensive margins are examined separately to distinguish between reductions in the number of fires and shifts toward complete compliance and elimination of burning within a cell-season.

The intensive-margin results in Table 1, estimated using the number of fires as the dependent variable, show a contrast across periods. In the pre-2016 years, the coefficient on the Haryana indicator is approximately 0.87, statistically indistinguishable from zero and small relative to the Punjab mean of about 5 fires, indicating that, if anything, crop-residue events were more frequent on the Haryana side of the border. In 2016–2017, the coefficient remains close to zero and statistically insignificant, reinforcing the interpretation that early regulatory activity in Haryana had limited operational bite. In 2018–2021, the coefficient becomes roughly -2 and statistically significant at the one percent level. Relative to the Punjab-side mean of about 6 fires in this period, this corresponds to a reduction of roughly one-third in burning intensity. In the post-2022 period, the estimated discontinuity is approximately -1.7 and significant at the one percent level, indicating sustained but attenuated divergence once Punjab begins implementation of its residue-management plan. The distance-to-border (D_c) term is small and negative across periods, while the Haryana-distance interaction remains close to zero and statistically insignificant, as expected in a design that exploits abrupt policy discontinuities at an administrative boundary.

The extensive-margin results in Table 2 reflect parallel patterns. In the pre-2016 and 2016–2017 periods, the coefficients are small and statistically insignificant. In 2018–2021, the

Bondarenko et al. (2025) for the period 2015–2023. Cells within 15 km from the Haryana–Punjab border have on average about 2,300 inhabitants in 2015.

coefficient rises to approximately 0.14 and is statistically significant at the five percent level. Relative to the Punjab-side probability of 0.17, this implies an increase of about 80% in the likelihood that a Haryana-side cell experiences no fires in a given season. This indicates that the regulation not only reduces the number of fires but also induces a shift into complete compliance. In the post-2022 period, the coefficient increases slightly to about 0.17, statistically significant at the one percent level, consistent with continued improvements but diminished divergence as Punjab ramps up enforcement. As with the intensive margin, the main distance term is positive and significant in multiple periods, reflecting mild gradients deeper into Punjab, while the Haryana-distance interaction is always small and statistically insignificant.

4.3 Robustness and Seasonality

A series of robustness exercises assesses the stability of the results.

Seasonality. The seasonality of burning activity provides a placebo environment in which to test the identification strategy. The wheat harvest (April–May) and the non-harvest months (December–March and June–August) present windows in which incentives to burn are reduced due to the longer period available for residue management. The RD plots for these periods, reported in Figure 6 and Figures 8 and 9, show no systematic border discontinuities. This absence of effects supports the claim that the main-season discontinuities reflect policy-driven changes in the incentives associated with burning.

Agro and socio-economic conditions. Pre-policy balance is assessed using both observed socio-economic characteristics from the 2011 SECC and agro-climatic fundamentals from FAO–GAEZ 2010/2011. As shown in Table 3 and Appendix Figure 10, SECC variables—including the share of households with cultivation as the primary income source, landownership patterns, agricultural capital, and housing characteristics—display smooth spatial profiles with no statistically significant discontinuities at the Punjab–Haryana border. Small visual drifts in a few panels (e.g., cultivation share and landownership) are economically minor; moreover, they run counter to the direction of later treatment effects, since Haryana’s side of the border exhibited slightly higher burning than Punjab in the pre-2016 years.

FAO–GAEZ layers, summarised in Table 4 and plotted in Appendix Figure 11, similarly show broad spatial continuity in wheat harvested area, yields, and suitability indices. Two rice-related rasters display small downward shifts on the Haryana side, yet these differences are negligible relative to period means and should be interpreted cautiously given that

GAEZ variables are partly model-based and smooth by construction, rather than direct observations of farm-level production. As discussed in Section 3, the GAEZ data are therefore used primarily to characterise slow-moving geography, not to validate cross-border balance.

Taken together, the SECC and GAEZ evidence indicates that neither observed socioeconomic conditions nor underlying agro-climatic fundamentals exhibit meaningful discontinuities at the border prior to the policy. This, and the absence of pre-trends, supports the identifying assumptions of the spatial RD design and confirms that the post-2018 discontinuities in burning and $\text{PM}_{2.5}$ reflect policy-induced changes rather than pre-existing differences in structural conditions.

Robustness to alternative inference and weighting choices. To assess the sensitivity of inference to spatial dependence and to the weighting scheme used within the RD window, I conduct two complementary robustness checks. First, I re-estimate the baseline specification using a coarser, one-way spatial clustering scheme based on 20×20 km grid cells. As shown in robustness Tables 5 and 6, this alternative clustering scheme yields similar standard errors and does not materially affect inference, suggesting that the baseline approach captures the relevant correlation structure. Second, I re-estimate the fire-count RD using a triangular kernel rather than the uniform kernel adopted in the main tables. Table 7 shows that the point estimates and standard errors are essentially unchanged, implying that the results are not sensitive to the choice of kernel (and hence to the relative weighting of observations within the 15 km window).

5 Effects of Environmental Regulation on $\text{PM}_{2.5}$ Pollution

Demonstrating that environmental regulation can curb agricultural burning practices raises a question of central policy relevance: to what extent do these practice changes translate into improvements in ambient $\text{PM}_{2.5}$, the principal indicator of air-pollution exposure with direct implications for climate mitigation and public health? This section examines how the observed policy discontinuities and reductions in crop-residue burning map into measurable changes in pollution concentrations and documents the magnitude, timing, and health implications of these $\text{PM}_{2.5}$ effects.

5.1 Wind Data and Conditioning Framework

To ensure that differences in $\text{PM}_{2.5}$ reflect local emissions rather than atmospheric transport, I condition pollution data on wind direction and intensity around the Punjab–Haryana

border. Daily wind vectors at 10 m height are obtained from the ERA5 reanalysis dataset and decomposed into zonal (u) and meridional (v) components. For each grid cell c and day d , the wind vector $w_{cd} = (u_{cd}, v_{cd})$ is projected onto two orthogonal unit vectors: a tangent vector t_c running along the border, and a normal vector n_c pointing from Punjab into Haryana. The across-border and along-border components and the wind speed are then measured as:

$$w_{cd}^{\perp} = w_{cd} \cdot n_c, \quad w_{cd}^{\parallel} = w_{cd} \cdot t_c, \quad s_{cd} = \|w_{cd}\|.$$

where $w_{cd}^{\perp}$ denotes the *across-border* wind component, capturing the portion of daily wind flow that blows directly from Punjab into Haryana (or vice versa); $w_{cd}^{\parallel}$ denotes the *along-border* wind component, measuring the share of wind flow that moves parallel to the administrative boundary and thus generates minimal cross-border advection; and s_{cd} is the scalar wind speed, representing the overall intensity of air movement independent of direction. The ratio $|w_{cd}^{\perp}|/s_{cd}$ measures the angular deviation of the wind from the border orientation. Winds are classified as *parallel* when this ratio is below $\sin(15^\circ) \approx 0.26$, meaning the wind direction lies within 15 degrees of the border line.

5.2 Beaufort Scale and Wind-Speed Thresholds

To distinguish between different intensities of air movement, wind speeds are categorised using the Beaufort wind force scale—the internationally recognised system used by the U.S. National Weather Service and the World Meteorological Organization to describe surface wind conditions over land (World Meteorological Organization, 1970). The following table summarises the full Beaufort classification; the categories used in this study are shown in bold.

Beaufort scale for land-based wind speed (m/s).

Beaufort number	Description	Wind speed (m/s)
0	Calm: smoke rises vertically	<0.3
1	Light air: smoke drifts, flags barely move	0.3–1.5
2	Light breeze: wind felt on face, leaves rustle	1.6–3.3
3	Gentle breeze: twigs in constant motion	3.4–5.4
4	Moderate breeze: dust and loose paper raised	5.5–7.9
5	Fresh breeze: small trees sway	8.0–10.7
6	Strong breeze: large branches in motion	10.8–13.8
7	Near gale: whole trees in motion	13.9–17.1
8	Gale: twigs break off trees	17.2–20.7
9	Strong gale: slight structural damage occurs	20.8–24.4
10	Storm: trees uprooted, considerable structural damage	24.5–28.4
11	Violent storm: widespread damage	28.5–32.6
12	Hurricane force: extreme destruction	≥ 32.7

5.3 Wind-Based Classification of Meteorological Regimes

Using the angular and speed measures defined above, I classify each day using a set of meteorological regimes that vary in the extent to which they isolate pollution locally around the Punjab–Haryana border. These regimes impose progressively weaker restrictions on cross-border atmospheric transport and form the basis for the wind-conditioned $\text{PM}_{2.5}$ outcomes used in the RD analysis.

Stringent regime. The primary identification sample consists of days on which (i) winds blow within 15 degrees of the border line ($|w_{cd}^\perp|/s_d \leq 0.26$), and (ii) wind speed is below 1.5 m/s (Beaufort 0 and 1, “calm” and “light air”). These conditions imply near-calm, border-parallel flows under which pollutants remain local and cross-border advection is minimal. This regime provides the cleanest isolation of emission-driven differences in air quality across the border.

Less stringent regime. The less stringent regime maintains the same parallel-wind requirement but relaxes the speed threshold to less than 3.3 m/s, thus including Beaufort 2, “light breeze”. Winds in this range may induce some mixing but remain sufficiently weak that most pollution dispersal occurs along, rather than across, the border. This regime provides a precision-robustness trade-off relative to the stringent sample.

Parallel-only regime. The next regime removes the speed restriction entirely and con-

ditions only on wind direction. All days with winds within 15 degrees of the border are retained, regardless of speed. This sample reflects meteorological conditions under which winds are aligned with the border but stronger dispersion may occur. It is therefore a useful test of whether the RD results persist when conditioning exclusively on directional alignment.

All-days regime. Finally, the all-days regime imposes no wind restrictions and simply averages $\text{PM}_{2.5}$ over the burning season. This specification provides a baseline measure of pollution but is expected to exhibit the weakest spatial discontinuities due to potentially substantial cross-state transport.

Together, these four regimes offer a transparent sequence of identification environments, ranging from maximal atmospheric isolation (stringent) to fully inclusive conditions (all days). The pattern of RD estimates across these regimes provides a diagnostic of whether the observed cross-border differences in $\text{PM}_{2.5}$ reflect local emission behaviour rather than wind-driven dispersion.

5.4 Wind-Pollution Integration

For each year, daily $\text{PM}_{2.5}$ readings from GHAP and LGHAP are averaged across October–November (the core period of the Kharif rice-residue burning season) separately under each wind regime. This yields wind-regime-specific seasonal means

$$\{\text{PM}_{c,t}^r\} \in \{\text{stringent, less stringent, parallel, all}\},$$

where, for example, $\text{PM}_{c,t}^{\text{stringent}}$ denotes mean burning-season $\text{PM}_{2.5}$ in cell c and year t under stringent wind conditions. These variables are used as outcomes in the spatial RD design that compares air quality on either side of the Punjab–Haryana border under comparable meteorological conditions. This multi-source, wind-conditioned framework ensures that estimated policy discontinuities in $\text{PM}_{2.5}$ capture local differences in emission behaviour rather than artefacts of atmospheric transport or measurement scale.

5.5 Harmonisation of Pollution Sources

As mentioned in Section 3, because GHAP is available only from 2017 onward while LGHAP provides a longer historical series, I combine the two data products to study pollution effects over 2015–2022. The two datasets differ in calibration, so I harmonise them to express all $\text{PM}_{2.5}$ measures on a common scale.

Over the overlapping years 2017–2021, and separately for each wind regime (stringent, less

stringent, parallel, all), I estimate an affine mapping from LGHAP to GHAP:

$$\text{PM}_{c,t}^{\text{GHAP},r} = \alpha_r + \beta_r \text{PM}_{c,t}^{\text{LGHAP},r} + \gamma_t + \varepsilon_{c,t}, \quad t \in [2017, 2021], \quad (2)$$

where γ_t are year fixed effects. In the all-days sample, regressing GHAP on LGHAP yields an intercept of about $22 \mu\text{g}/\text{m}^3$ and a slope of about 0.82 (both statistically significant at the 1% level), with R^2 around 0.8. Across wind regimes, the corresponding mappings display slopes in a narrow band around 0.8 and similarly high R^2 values, indicating a stable affine relationship across meteorological conditions. See Appendix 10 and Figure 13 for details.

I then construct a harmonised $\text{PM}_{2.5}$ series by using GHAP where available and rescaling LGHAP to GHAP units before 2017:

$$\text{PM}_{c,t}^{\text{Harmonised},r} = \begin{cases} \text{PM}_{c,t}^{\text{GHAP},r}, & \text{if } t \geq 2017, \\ \hat{\alpha}_r + \hat{\beta}_r \text{PM}_{c,t}^{\text{LGHAP},r}, & \text{if } t < 2017. \end{cases} \quad (3)$$

Unless otherwise noted, all $\text{PM}_{2.5}$ outcomes are expressed in this harmonised GHAP scale.

5.6 Empirical Strategy

The empirical design for pollution mirrors the spatial RD framework used for burning outcomes (Equation 1).⁹ Haryana began enforcing crop-residue burning regulations more stringently in 2018, whereas Punjab continued to apply weaker measures for several additional years. Identification exploits the sharp administrative border separating the two states and compares local averages of pollution immediately on either side.

A key feature of the pollution analysis is that rice-burning-season $\text{PM}_{2.5}$ outcomes are not constructed from all daily observations, but rather averaged from the subset of days belonging to specific wind regimes. For each grid cell and year, I compute harmonised $\text{PM}_{2.5}$ means over October–November separately for the *stringent*, *less stringent*, *parallel-only*, and *all-days* regimes. The preferred specification uses the *stringent* regime, which isolates conditions with minimal cross-border dispersion; the more inclusive regimes progressively relax these conditions and serve as robustness analyses by allowing increasing amounts of atmospheric transport. Formally, for grid cell c ($2\text{km} \times 2\text{km}$), signed distance D_c from the border, year t , and wind regime r , I estimate

⁹The pollution strategy is parallel to the burning analysis except for two timeline/construction differences: (i) average $\text{PM}_{2.5}$ outcomes are constructed over the core period October–November (whereas fires are counted over September–November); and (ii) $\text{PM}_{2.5}$ data begin in 2015, so the pre-policy period in the pollution analysis consists only of 2015 (whereas the burning analysis starts in 2012).

$$\text{Harmonised PM}_{c,t}^r = \alpha_t^r + \tau_t^r \text{Haryana}_c + \lambda_{1t}^r D_c + \lambda_{2t}^r D_c \times \text{Haryana}_c + \gamma_{s(c)}^r + \varepsilon_{c,t}^r, \quad (4)$$

where Harmonised PM_{ct}^r is harmonised $\text{PM}_{2.5}$ in cell c and year t under regime r , D_c and Haryana_c are defined as in equation 1, and $\gamma_{s(c)}^r$ are border-segment fixed effects. As in the burning analysis, D_c and $D_c \times \text{Haryana}_c$ implement side-specific local-linear trends in distance to the border, and the estimation window uses the optimal bandwidth selected following Calonico et al. (2014); for $\text{PM}_{2.5}$ this yields an optimal bandwidth of $h_t = 40$ km on each side of the border in the baseline specification. This reflects the smoother spatial structure and higher effective noise of satellite-based $\text{PM}_{2.5}$ relative to fire counts¹⁰. Within this window, I again apply a uniform kernel, assigning equal weight to all cells with $|D_c| \leq 40$ and zero weight otherwise¹¹. Cells that intersect the boundary are excluded via a donut-hole correction. Standard errors are two-way clustered on overlapping spatial grids, each of $10\text{km} \times 10\text{km}$, allowing for general spatial correlation patterns across cells and years.

The parameter of interest, τ_t^r , captures the discontinuity in $\text{PM}_{2.5}$ across the border under wind regime r in year t . I first estimate τ_t^r separately for each year to construct the event-study profiles in Figure 7. To summarise effects over broader regulatory phases while keeping the unit of observation at the cell-year level, similarly to burning outcomes, I then estimate the same RD specification on four period-specific samples: pre-policy (2015), transition (2016–2017), Haryana treated (2018–2021), and both states treated (2022). A negative and statistically significant τ_t^r indicates that seasonal $\text{PM}_{2.5}$ is lower on the Haryana side of the border, consistent with the view that Haryana’s strengthened enforcement reduced burning-related emissions.

5.7 Main Results

Table 8 reports the estimated discontinuities in harmonised $\text{PM}_{2.5}$ concentrations under the *stringent* wind regime, which—by restricting the sample to days with winds that are both nearly parallel to the border and below 1.5 m/s—provides the cleanest isolation of local emission differences. The estimates for the periods preceding Haryana’s regulatory ramp-up, namely 2015 and 2016–2017, are small in magnitude and statistically indistinguishable from zero. This absence of any systematic cross-border gap prior to 2018 strengthens the

¹⁰Satellite-based $\text{PM}_{2.5}$ is spatially smoother and measured with non-trivial error, and wind-conditioning reduces the effective signal per cell-year by restricting the set of days entering the seasonal mean. Consistent with these features, the procedure selects a wider window for $\text{PM}_{2.5}$ (40 km) than for fire counts (15 km), for which the outcome is more granular and exhibits sharper local heterogeneity in distance to the border.

¹¹As in the burning specification, I use a uniform kernel within the estimation window (with border-intersecting cells excluded via the donut-hole). Results are robust to using a triangular kernel (see Section 5.7).

credibility of the design by confirming that the two states evolved similarly in pollution levels before the onset of differential agricultural & environmental enforcement.

A sharp and persistent discontinuity emerges once Haryana strengthens enforcement in 2018. During the 2018–2021 period, pollution levels are approximately $9 \mu\text{g}/\text{m}^3$ lower on the Haryana side of the border, and the effect remains sizeable in 2022, with an estimated reduction of $8.8 \mu\text{g}/\text{m}^3$. These magnitudes are economically meaningful: relative to the contemporaneous mean $\text{PM}_{2.5}$ concentration measured in the adjacent Punjab border area (roughly $88.8 \mu\text{g}/\text{m}^3$), the discontinuity corresponds to about a 10% reduction. The temporal alignment between these decreases and the documented fall in rice-residue fires in Section 4 provides evidence that Haryana’s enforcement reduced burning, thereby effectively improving ambient air quality.

Figure 7 complements the period-aggregated estimates by plotting yearly RD coefficients. The event-study profile reveals a flat and statistically insignificant pattern in all pre-2016 years, followed by a downward shift from 2017 onward. The yearly estimates fluctuate more widely than the multi-year averages; nevertheless, the shape of the profile mirrors the qualitative pattern in the fire outcomes, showing a break at the moment when Haryana intensifies enforcement.

The estimated decline of $9 \mu\text{g}/\text{m}^3$ during the burning season represents an improvement in air quality. When averaged over the year, this corresponds to an annualised reduction of approximately $1.50 \mu\text{g}/\text{m}^3$ in long-run exposure. The effect sizes are consistent with existing estimates of the contribution of agricultural burning to national $\text{PM}_{2.5}$ exposure in India, as reported by Lan et al. (2022) and Dipoppa and Gulzar (2024). The magnitude of the discontinuity implies a persistent cross-border gap in seasonal pollution attributable to differences in environmental enforcement, aligned both with the timing of Haryana’s regulatory ramp-up and with the contemporaneous decline in fire counts documented in Section 4.

5.8 Robustness checks

I assess the robustness of the pollution RD estimates to alternative wind-selection rules and the weighting scheme used within the RD window, while holding the core specification fixed.

Relaxing the wind restrictions progressively attenuates the estimated effects. Under the *less stringent* regime in Table 9, which admits slightly stronger winds while maintaining near-parallel orientation, and the *parallel-only* regime in Table 10, which imposes no restriction on wind speed, the effect remains negative and statistically significant but falls to roughly $7 \mu\text{g}/\text{m}^3$. The *all-days* specification in Table 11 produces estimates that are smaller still and

imprecisely estimated. This monotonic attenuation is consistent with atmospheric transport: as the wind criteria are relaxed and more days enter the seasonal mean—including days with greater scope for cross-border dispersion—regional mixing increasingly blurs the spatial discontinuity, reducing the share of variation plausibly attributable to local enforcement. Lastly, Table 12 re-estimates Table 8 using a triangular kernel; point estimates and inference are unchanged.

Taken together, the evidence across tables and event-study estimates shows a coherent and policy-aligned pattern. Prior to 2018, pollution levels are symmetric around the border; immediately after Haryana escalates enforcement, $\text{PM}_{2.5}$ concentrations fall relative to Punjab; and the effect is strongest under those meteorological regimes that most effectively isolate local emissions. The pollution analysis thus mirrors the fire-based RD results: under meteorological conditions most conducive to separating local from regional pollution, seasonal $\text{PM}_{2.5}$ is lower on the Haryana side after the enforcement escalation, and this pattern is robust to alternative kernel choices.

5.9 Health impacts and monetary implications of Haryana’s crop-burning enforcement

This section translates the RDD-based reduction in $\text{PM}_{2.5}$ induced by Haryana’s crop-residue burning enforcement into health impacts and monetary equivalents. The calculations combine three ingredients: (i) the estimated change in exposure from the spatial RD, (ii) externally estimated concentration-response relationships linking long-run $\text{PM}_{2.5}$ exposure to mortality (Ayres, 2010; Chen and Hoek, 2020) and to early-life mortality (Dipoppa and Gulzar, 2024), and (iii) demographic and disease-burden baselines from official population projections (National Commission on Population and Ministry of Health & Family Welfare, 2019; National Health Systems Resource Centre, 2021) and the Global Burden of Disease (GBD 2018) (Institute for Health Metrics and Evaluation – IHME, 2025). Table 13 summarises Haryana’s demographic data, and Appendix 11 documents each transformation used to obtain the figures below.

The spatial RDD implies that, from 2018 onward, Kharif-season $\text{PM}_{2.5}$ is about $9 \mu\text{g}/\text{m}^3$ lower on the Haryana side of the border than on the Punjab side; expressed as an annual-average exposure change, this corresponds to roughly $1.5 \mu\text{g}/\text{m}^3$. Applying the log-linear all-cause mortality CRFs in Ayres (2010) and Chen and Hoek (2020) implies a 0.9–1.2% reduction in all-cause mortality, corresponding to about 1,600–2,100 avoided deaths per year in Haryana (7,800–10,400 over 2018–2022). Beyond the aggregate effect, the age profile is policy-relevant. Using the early-life elasticities in Dipoppa and Gulzar (2024), the same exposure decline implies approximately 187 fewer infant deaths and 272 fewer under-five

deaths per year.¹² These magnitudes imply that a non-trivial share (13%–17%) of mortality reductions accrues in the first years of life, consistent with heightened vulnerability to air pollution during pregnancy and early childhood, and with the idea that enforcement-induced reductions in burning can generate immediate distributional gains for households with young children.

To summarise mortality changes in a single health-burden metric that accounts for age at death and non-fatal health loss, I use Global Burden of Disease (GBD 2018) estimates for Haryana (Institute for Health Metrics and Evaluation – IHME, 2025) to convert avoided deaths into disability-adjusted life years (DALYs). One DALY corresponds to one lost year of full health, combining years of life lost due to premature mortality and years lived with disability. In Haryana, each PM_{2.5}-attributable death corresponds to about 31 DALYs for all ages and about 90 DALYs for young children, reflecting the larger remaining life expectancy when mortality occurs early in life. Combining these ratios with the avoided-death estimates implies that the policy gains from reducing pollution concentration correspond to about 0.7 additional healthy days per resident per year; the implied gains are about 10.4 healthy days per infant per year and about 3.0 healthy days per under-five child per year.

Two assumptions of the calculation push the estimates towards conservatism. First, the exposure reduction is identified at the Haryana–Punjab *border*. Extending this effect to the entire state assumes that the change in enforcement is broadly uniform across districts. Existing political-economy work consistently shows that the quality of governance, public-good provision, and the impact of national institutions tend to deteriorate with distance from the political centre: isolated regions exhibit weaker accountability and enforcement, and the effect of national institutions on local development decays with distance to the capital (Helland and Whitford, 2003; Michalopoulos and Papaioannou, 2014; Provenzano, 2024). Under this empirical regularity, the border discontinuity is likely to *understate* the average improvement experienced within Haryana, implying that the mortality gains reported above should be viewed as lower bounds. Second, the CRFs are estimated in populations that, on average, face substantially lower baseline pollution than Northern India; if marginal health gains from reducing PM_{2.5} are larger at high exposure levels—that is, if the pollution-health relationship is convex—then applying these CRFs to Haryana again yields lower-bound estimates of the true benefits. The avoided deaths, DALYs, and healthy days reported above therefore represent conservative estimates of the health gains causally attributable to Haryana’s crop-burning enforcement. I next attach monetary values to these gains to characterise the implied scale of the policy’s health benefits.

I first express these health gains in monetary terms using a DALY-based valuation. Because

¹²By construction, the under-five category (ages 0–4) includes infants (ages 0–1). Infant estimates are therefore a subset of under-five estimates and should not be added to them when aggregating effects.

India does not have an official cost-effectiveness threshold, following Bhat et al. (2023) I use the commonly applied WHO-CHOICE rule of thumb of 1–3 times GDP per capita per DALY averted.¹³ India’s GDP per capita in 2018 was about 2,000 USD (current US\$; World Development Indicators) (The World Bank, 2025), implying a benchmark range of roughly 2,000–6,000 USD per DALY. Applied to the midpoint estimate of 56,700 DALYs saved annually, this implies monetary benefits on the order of 110–340 million USD per year from Haryana’s enforcement. Relative to the 650,900 acres estimated to have been burned in the Kharif season of 2018 in Haryana (Kumar et al., 2019), this corresponds to per-acre benefits of about 170–520 USD. Any policy that achieves the estimated reduction in burning at a cost below this range per acre would, under the WHO-CHOICE rule of thumb, fall within conventional cost-effectiveness thresholds.¹⁴

As a complementary benchmark, I also compute a value-of-a-statistical-life (VSL)-based valuation. Using a conservative VSL for India of 688,000 USD (Jack et al., 2025; Majumder and Madheswaran, 2018) and the estimated 1,600–2,100 avoided deaths per year yields monetary benefits of about 1.10–1.45 billion USD annually. Normalizing by the same 2018 burned acreage (Kumar et al., 2019) implies roughly 1,700–2,200 USD¹⁵ per acre.¹⁶

These monetary valuations should be interpreted as short-run benchmarks. The RDD estimates are identified from the first four years following Haryana’s de facto enforcement (2018–2022) and thus capture the initial response of burning and emissions to the policy. Over a longer horizon, farmers’ behaviour and production choices may adjust further—through changes in compliance, technology adoption, and cropping decisions—so that the share of the burning-related health burden mitigated by enforcement could increase. Analysing these longer-run dynamics is beyond the scope of this paper, but the results indicate that even in the early years of the policy, Haryana’s enforcement offsets a material fraction of the health damages associated with crop-residue burning.

¹³While the health impacts are expressed in DALYs, I use GDP-per-capita multiples as an order-of-magnitude benchmark for valuing a healthy life-year in the Indian context; closely related benchmarks are often applied to QALYs in the applied health-economics literature.

¹⁴These figures should be read as benchmark valuations of the estimated health gains. A full cost-effectiveness comparison requires credible measurement of the policy’s costs, including administrative, monitoring, and compliance costs, which are not fully observed here.

¹⁵The VSL-based numbers exceed the DALY-based valuations because they reflect willingness-to-pay for mortality risk reductions rather than health-system benchmark thresholds.

¹⁶For comparison, Lan et al. (2022) estimate that crop-residue burning causes 44,000–98,000 premature deaths annually across India, attributing 7.8–14% of this burden to Haryana—implying 6,000–10,000 deaths per year from the state’s burning alone. Jack et al. (2025) estimate VSL-valued mortality damages of about 6,400 USD per acre of residue burned in Punjab. These studies quantify the *total* mortality burden of burning; by contrast, the estimates in this paper capture the *marginal* health and monetary gains from Haryana’s 2018 enforcement, not the full burden.

6 Conclusion

This paper asks whether subnational environmental regulation can curb crop-residue burning, a salient, privately cheap, and spatially mobile source of $\text{PM}_{2.5}$ in low- and middle-income countries. Focusing on Northern India and the Punjab–Haryana border, I study the effects of Haryana’s regulatory bundle—which combines prohibitions, penalties, monitoring, and technology subsidies—on crop-residue burning, $\text{PM}_{2.5}$, and health. The empirical strategy exploits the sharp administrative boundary and the staggered rollout of enforcement, comparing neighbouring locations just inside Haryana and Punjab that share similar agro-climatic conditions, cropping patterns, and pre-policy socio-economic characteristics. Fire outcomes are measured using 2012–2023 VIIRS active fire detections on a 2×2 km grid. Pollution outcomes use harmonised satellite-based $\text{PM}_{2.5}$ conditioned on wind regimes that limit cross-border transport. I then combine the estimated exposure reductions with standard concentration-response functions and demographic data to quantify health impacts and monetary equivalents.

Three main sets of findings emerge. First, on farmers’ behaviour, Haryana’s enforcement changes crop-residue management practices. During the 2018–2021 “Haryana treated” period, cells on the Haryana side of the border experience roughly 2 fewer rice-residue fires per Kharif season relative to adjacent Punjab cells, a reduction of about one-third relative to the Punjab-side mean in the same period. At the extensive margin, the probability that a cell is fire-free rises by about 14 percentage points from a 17% baseline on the Punjab side, implying an 80% increase in the likelihood of zero burning. These changes are absent in the pre-2016 period and attenuate once Punjab begins implementing a comparable residue-management plan in 2022. Taken together, the results indicate that state-level regulation can reduce burning intensity and shift a substantial mass of locations from persistent burning to complete compliance.

Second, the reductions in burning translate into improvements in local air quality. Conditioning on near-calm, border-parallel winds, the spatial RD estimates indicate that burning-season $\text{PM}_{2.5}$ on the Haryana side of the border is about $9 \mu\text{g}/\text{m}^3$ lower than in Punjab after 2018, roughly a 10% reduction relative to contemporaneous Punjab means. The temporal pattern in pollution mirrors the fire results: no systematic gap before 2018 and a discrete decline during the period of asymmetric enforcement. The estimated discontinuity is largest under wind regimes that most tightly limit cross-border transport and attenuates as wind restrictions are relaxed, consistent with atmospheric mixing blurring local pollution levels. These findings show that even for pollutants as mobile as fine particulates, differential enforcement can generate sharp discontinuities in exposure near jurisdictional boundaries.

Third, the implied health gains are economically large. Applying standard concentration-

response functions for all-cause mortality to the estimated exposure reductions yields roughly 1,600–2,100 avoided premature deaths per year in Haryana. Of these, about 187 are infant deaths and 272 are deaths among children under five. Translating these effects into disability-adjusted life years, the enforcement policy saves roughly 0.7 additional healthy days per resident annually and approximately 10.4 and 3 additional healthy days per year for the average infant and under-five child, respectively. Valued using DALY-based thresholds commonly applied in the Indian context, these health gains imply annual monetary benefits of roughly 110–340 million USD; expressed on a per-acre basis, this corresponds to a benchmark valuation of about 170–520 USD per acre of residue burning avoided under the WHO–CHOICE rule of thumb. Under a conservative value-of-statistical-life approach, the estimated mortality reductions correspond to approximately 1.10–1.45 billion USD in annual benefits, or about 1,700–2,200 USD per acre of burning abated. These valuations underscore that stricter enforcement of burning bans is not only environmentally effective but also potentially highly cost-effective.

Beyond quantifying these effects, the paper contributes to broader debates on environmental governance in low- and middle-income countries. A substantial literature documents that strong environmental standards often coexist with weak *de facto* enforcement, particularly in contexts such as India, leading to widespread pessimism about the scope for regulation to curb pollution. This view is especially salient in the crop-residue burning crisis, where previous studies have found modest or fragile impacts. By contrast, the Haryana experience shows that when regulation is credibly enforced and supported by appropriate incentives and technology, it can effectively reduce emissions and improve public health in densely populated, rapidly developing settings. By documenting that a credibly enforced state-level regulation can deliver large and persistent reductions in burning and exposure, this paper adds a further piece to the policy puzzle of how to confront what has long been treated as an intractable environmental problem.

References


- ABDURRAHMAN, M. I., S. CHAKI, AND G. SAINI (2020): “Stubble burning: Effects on health & environment, regulations and management practices,” *Environmental Advances*, 2, 100011.
- ASHER, S., T. LUNT, R. MATSUURA, AND P. NOVOSAD (2021): “Development Research at High Geographic Resolution: An Analysis of Night-Lights, Firms, and Poverty in India Using the SHRUG Open Data Platform,” *The World Bank Economic Review*, 35, 845–871.
- AYRES, J. G. (2010): *The mortality effects of long-term exposure to particulate air pollution in the United Kingdom: a report*, England?: Health Protection Agency.
- BAI, K., K. LI, L. SHAO, X. LI, C. LIU, Z. LI, M. MA, D. HAN, Y. SUN, Z. ZHENG, R. LI, N.-B. CHANG, AND J. GUO (2024): “LGHAP v2: a global gap-free aerosol optical depth and PM_{2.5} concentration dataset since 2000 derived via big Earth data analytics,” *Earth System Science Data*, 16, 2425–2448.
- BEHRER A. PATRICK (2023): “Man or Machine? Environmental Consequences of Wage Driven Mechanization in Indian Agriculture,” *World Bank Policy Research Working Paper No. 10376*.
- BHARADWAJ, P., J. FENSKE, N. KALA, AND R. A. MIRZA (2020): “The Green revolution and infant mortality in India,” *Journal of Health Economics*, 71, 102314.
- BHAT, B., J. DE QUIDT, J. HAUSHOFER, V. PATEL, G. RAO, F. SCHILBACH, AND P.-L. VAUTREY (2023): “The Long-Run Effects of Psychotherapy on Depression, Beliefs, and Economic Outcomes,” .
- BIKKINA, S., A. ANDERSSON, E. N. KIRILLOVA, H. HOLMSTRAND, S. TIWARI, A. K. SRIVASTAVA, D. S. BISHT, AND O. GUSTAFSSON (2019): “Air quality in megacity Delhi affected by countryside biomass burning,” *Nature Sustainability*, 2, 200–205.
- BONDARENKO, M., R. PRIYATIKANTO, N. TEJEDOR GARAVITO, W. ZHANG, T. MC-KEEN, A. CUNNINGHAM, T. WOODS, J. HILTON, D. CIHAN, B. NOSATIUK, T. BRINKHOFF, A. TATEM, AND A. SORICETTA (2025): “Constrained estimates of 2015-2030 total number of people per grid square at a resolution of 3 arc (approximately 100m at the equator) R2025A version v1.” .
- BURGESS, R., F. COSTA, AND B. A. OLKEN (2023): “National Borders and the Conservation of Nature,” .
- CALONICO, S., M. D. CATTANEO, AND R. TITIUNIK (2014): “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs: Robust Nonparametric Confidence Intervals,” *Econometrica*, 82, 2295–2326.
- CAMERON, A. C., J. B. GELBACH, AND D. L. MILLER (2011): “Robust Inference With Multiway Clustering,” *Journal of Business & Economic Statistics*, 29, 238–249.
- CHEN, J. AND G. HOEK (2020): “Long-term exposure to PM and all-cause and cause-specific mortality: A systematic review and meta-analysis,” *Environment International*, 143, 105974.
- CHEN, Y., A. EBENSTEIN, M. GREENSTONE, AND H. LI (2013): “Evidence on the impact of sustained exposure to air pollution on life expectancy from China’s Huai River policy,” *Proceedings of the National Academy of Sciences of the United States of America*, 110, 12936–12941.
- COPERNICUS CLIMATE CHANGE SERVICE (2024): “ERA5 post-processed daily statistics on single levels from 1940 to present,” .
- DAVIS, L. (2008): “The Effect of Driving Restrictions on Air Quality in Mexico City,” *Journal of Political Economy*, 116, 38–81.
- DIPOPPA, G. AND S. GULZAR (2024): “Bureaucrat incentives reduce crop burning and child mortality in South Asia,” *Nature*, 634, 1125–1131.
- DUFLO, E., M. GREENSTONE, R. PANDE, AND N. RYAN (2013): “Truth-telling by

- Third-party Auditors and the Response of Polluting Firms: Experimental Evidence from India*,” *The Quarterly Journal of Economics*, 128, 1499–1545.
- FALCON, W., G. HADIWIDJAJA, R. EDWARDS, M. HIGGINS, R. NAYLOR, AND S. SUMARTO (2022): “Using Conditional Cash Payments to Prevent Land-Clearing Fires: Cautionary Findings from Indonesia,” *Agriculture*, 12, 1040.
- FENSKE, J., M. HASEEB, AND N. KALA (2023): “How Rules and Compliance Impact Organizational Outcomes: Evidence from Delegation in Environmental Regulation,” *NBER Working Paper No. 31991*.
- GARG, T., M. JAGNANI, AND H. K. PULLABHOTLA (2024): “Rural Roads, Farm Labor Exits, and Crop Fires,” *American Economic Journal: Economic Policy*, 16, 420–450.
- GELMAN, A. AND G. IMBENS (2019): “Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs,” *Journal of Business & Economic Statistics*, 37, 447–456.
- GOVERNMENT OF INDIA AND MINISTRY OF AGRICULTURE & FARMERS WELFARE (2019): “Report of the Committee Review of the Scheme ”Promotion of Agricultural Mechanisation for in-Situ Management of Crop Residue in States of Punjab, Haryana, Uttar Pradesh and NCT of Delhi”,” Tech. rep., New Delhi.
- (2021): “Agricultural Mechanization for In-Situ Management of Crop Residue,” .
- GREENSTONE, M. AND R. HANNA (2014): “Environmental Regulations, Air and Water Pollution, and Infant Mortality in India,” *American Economic Review*, 104, 3038–3072.
- HE, G., T. LIU, AND M. ZHOU (2020): “Straw burning, PM2.5, and death: Evidence from China,” *Journal of Development Economics*, 145, 102468.
- HELLAND, E. AND A. B. WHITFORD (2003): “Pollution incidence and political jurisdiction: evidence from the TRI,” *Journal of Environmental Economics and Management*, 46, 403–424.
- HOLMES, T. J. (1998): “The Effect of State Policies on the Location of Manufacturing: Evidence from State Borders,” *Journal of Political Economy*, 106, 667–705.
- INSTITUTE FOR HEALTH METRICS AND EVALUATION – IHME (2025): “GBD Results,” .
- JACK, B. K., S. JAYACHANDRAN, N. KALA, AND R. PANDE (2025): “Money (Not) to Burn: Payments for Ecosystem Services to Reduce Crop Residue Burning,” *American Economic Review: Insights*, 7, 39–55.
- JAGNANI, M. AND M. MAHADEVAN (2025): “Women leaders improve environmental outcomes: Evidence from crop fires in India,” *Journal of Public Economics*, 248, 105443.
- JALALI, U. (2023): “Pollution pangs: No solution in sight for Delhi,” *The New Indian Express*, 23 October 2023.
- JETHVA, H., O. TORRES, R. FIELD, A. LYAPUSTIN, R. GAUTAM, AND V. KAYETHA (2019): “Connecting Crop Productivity, Residue Fires, and Air Quality over Northern India,” *Scientific Reports*, 9, 16594.
- KHUNDRAKPAM, P. AND J. K. SARMAH (2023): “Stubble Burning in India: Politics and Policy Responses,” *Indian Journal of Public Administration*, 69, 303–316.
- KUMAR, P. AND L. JOSHI (2015): *Socioeconomic and Environmental Implications of Agricultural Residue Burning*.
- KUMAR, P., S. K. RAJPOOT, V. JAIN, S. SAXENA, NEETU, AND S. S. RAY (2019): “Monitoring of Rice Crop in Punjab and Haryana with Respect to Residue Burning,” *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-3-W6, 31–36.
- LAN, R., S. D. EASTHAM, T. LIU, L. K. NORFORD, AND S. R. H. BARRETT (2022): “Air quality impacts of crop residue burning in India and mitigation alternatives,” *Nature Communications*, 13, 6537.
- LIPSCOMB, M. AND A. M. MOBARAK (2017): “Decentralization and Pollution Spillovers: Evidence from the Re-drawing of County Borders in Brazil*,” *The Review of Economic*

- Studies*, 84, 464–502.
- LIU, T., L. J. MICKLEY, S. SINGH, M. JAIN, R. S. DEFRIES, AND M. E. MARLIER (2020): “Crop residue burning practices across north India inferred from household survey data: Bridging gaps in satellite observations,” *Atmospheric Environment: X*, 8, 100091.
- MAGRUDER, J. (2013): “Can minimum wages cause a big push? Evidence from Indonesia,” *Journal of Development Economics*, 100, 48–62.
- MAJUMDER, A. AND S. MADHESWARAN (2018): “Value of statistical life in India: A hedonic wage approach,” *Working Papers*.
- MICHALOPOULOS, S. AND E. PAPAIOANNOU (2014): “National Institutions and Subnational Development in Africa *,” *The Quarterly Journal of Economics*, 129, 151–213.
- MINISTRY OF RURAL DEVELOPMENT, GOVERNMENT OF INDIA (2011): “Socio Economic and Caste Census,” .
- NATIONAL COMMISSION ON POPULATION AND MINISTRY OF HEALTH & FAMILY WELFARE (2019): “Population Projections for India and States, 2011–2036,” Tech. rep.
- NATIONAL GREEN TRIBUNAL (2017): “Order dated 18 December 2017 in Vardhaman Kaushik v. Union of India & Ors.” .
- NATIONAL HEALTH SYSTEMS RESOURCE CENTRE (2021): “Health Dossier 2021: Reflections on Key Health Indicators – Haryana,” Tech. rep.
- NIAN, Y. (2023): “Incentives, penalties, and rural air pollution: Evidence from satellite data,” *Journal of Development Economics*, 161, 103049.
- OLIVA, P. (2015): “Environmental Regulations and Corruption: Automobile Emissions in Mexico City,” *Journal of Political Economy*, 123, 686–724.
- PINKOVSKIY, M. L. (2017): “Growth discontinuities at borders,” *Journal of Economic Growth*, 22, 145–192.
- PROVENZANO, S. (2024): “Accountability failure in isolated areas: The cost of remoteness from the capital city,” *Journal of Development Economics*, 167, 103214.
- RANGEL, M. A. AND T. S. VOGL (2019): “Agricultural Fires and Health at Birth,” *The Review of Economics and Statistics*, 101, 616–630.
- SARKAR, S., R. P. SINGH, AND A. CHAUHAN (2018): “Crop Residue Burning in Northern India: Increasing Threat to Greater India,” *Journal of Geophysical Research: Atmospheres*, 123, 6920–6934.
- SRINIVASAN, A. (2023): “Choked: Where Is New Delhi’s Response to Air Pollution?” *The Diplomat*, 21 April 2023.
- STEVENS, F. R., A. E. GAUGHAN, C. LINARD, AND A. J. TATEM (2015): “Constrained estimates of 2015–2030 total number of people per grid square at a resolution of 3 arc (approximately 100m at the equator) R2025A version v1. Global Demographic Data Project - Funded by The Bill and Melinda Gates Foundation (INV-045237). WorldPop - School of Geography and Environmental Science, University of Southampton.” .
- THE WORLD BANK (2022): “Urgent Action Needed in South Asia to Curb Deadly Air Pollution,” .
- (2025): “World Development Indicators | DataBank,” .
- VENKATESH, S., J., I. KUKRETI, K. PANDEY, D. G. NIYOGI, AND P. MUKERJEE (2017): “India’s burning issue of crop burning takes a new turn,” *Down To Earth*.
- WEI, J., Z. LI, A. LYAPUSTIN, J. WANG, O. DUBOVIK, J. SCHWARTZ, L. SUN, C. LI, S. LIU, AND T. ZHU (2023): “First close insight into global daily gapless 1 km PM2.5 pollution, variability, and health impact,” *Nature Communications*, 14, 8349.
- WHO (2023): “Ambient (outdoor) air pollution,” .
- WORLD METEOROLOGICAL ORGANIZATION (1970): “The Beaufort Scale of Wind Force: Technical and Operational Aspects,” Reports on Marine Science Affairs No. 3, Geneva: Secretariat of the World Meteorological Organization.


7 Figures

Figure 1: Policy Timeline – Government of India, Haryana and Punjab




Notes: This figure summarises the timeline evolution of the national and state-level regulation of crop-residue burning in Northern India from 1981 to 2022. State-level annotations distinguish Haryana-specific context (above the timeline) from Punjab-specific context (below the timeline). See Section 2 for further institutional details.

Figure 2: Spatial Distribution of Rice-Residue Burning Around the Punjab–Haryana Border




Notes: This figure maps rice-residue burning activity in a 30 km band around the Punjab–Haryana border. Each panel corresponds to one Kharif (rice-residue burning) season in 2016, 2018, 2020, and 2022—years that mark distinct policy phases in the regulatory timeline—and shows 2×2 km grid cells coloured by the total number of detected fires during September–November; darker colours indicate more burning events. The solid orange line marks the administrative boundary used in the spatial RD design. See Section 3 for a full description of data sources; for details on the RD specification, see Section 4.1 and Section 5.6.

Figure 3: Spatial RD Plots for Rice-Residue Burning (2016–2022)




Notes: This figure presents spatial RD plots for rice-residue burning along the Punjab–Haryana border. Each panel corresponds to one Kharif (rice-residue burning) season in 2016, 2018, 2020, and 2022—years that mark distinct policy phases in the regulatory timeline. The vertical axis reports the mean number of rice-residue fires per 2×2 km cell (fires per 4 km^2), while the horizontal axis shows the signed distance to the border (negative values denote Punjab; positive values Haryana). Grey dots show binned cell averages; solid lines show local linear fits on each side of the border using the preferred bandwidth. These plots provide a non-parametric visualisation of the spatial discontinuity underlying the RD estimates. See Section 3 for data sources; for details on the RD specification, see Section 4.1.

Figure 4: Event-Study Estimates of the Spatial RD Effect on Rice-Residue Fires



Notes: This figure reports event-study RD estimates of the effect of being located in Haryana (versus Punjab) on the number of rice-residue fires per 2×2 km cell during the Kharif rice-residue burning seasons, 2012–2023. Points indicate yearly RD coefficients; vertical bars show 95% confidence intervals based on two-way spatial clustering. Vertical dashed lines mark key changes in the policy environment, with labels at the bottom indicating corresponding policy phases. See Section 3 for data sources; for RD details, see Section 4.1 and Section 5.6.


Figure 5: Event-Study Estimates of the RD Effect on Per-Capita Fire Exposure




Notes: This figure displays event-study RD estimates for exposure to rice-residue fires, expressed in per-capita terms. The vertical axis reports the difference in fires per capita between Haryana and Punjab within the RD bandwidth; the horizontal axis shows Kharif rice-residue burning seasons, 2012–2023. Points display yearly RD coefficients with 95% confidence intervals based on two-way clustering. Fires per capita are computed by dividing the number of rice-residue fires per 2×2 km cell by the local population count assigned to that cell. See Section 3 for data sources; for RD details, see Section 4.1 and Section 5.6.

Figure 6: Seasonality of the Regulation Effect on Agricultural Burning

(a) Rabi Wheat-Residue Burning Season




(b) Non-Harvest Months



Notes: This figure reports event-study RD estimates outside the Kharif (rice-residue burning) season. Panel (a) reports estimates for the Rabi (wheat-residue burning) season, April–May; Panel (b) reports estimates for months with no major harvesting activity, December–March and June–August. Vertical axes report the RD coefficient (difference in fires per 2×2 km cell across the border); horizontal axes report calendar years. See Section 3 for data sources; for RD details, see Section 4.1 and Section 5.6.

Figure 7: Event-Study Estimates of the Regulation Effect on PM_{2.5} Under Stringent Wind Conditions



Notes: This figure presents event-study RD estimates of the impact of Haryana’s regulation on PM_{2.5} concentrations under stringent wind-selection rules that limit cross-border transport. The vertical axis reports the estimated concentration difference (Haryana minus Punjab, $\mu\text{g}/\text{m}^3$); the horizontal axis reports yearly Kharif rice-residue burning seasons. PM_{2.5} is measured using harmonised GHAP–LGHAP satellite products and aggregated to seasonal means within each 2×2 km cell. See Section 3 for a detailed description of PM_{2.5} and wind data construction; for details on the pollution RD specification, see Section 5.6.

8 Tables

Table 1: Main Table: Policy Impact on Number of Fire Events per cell (4 km²)

	N. of Fires			
	(1) Pre 2016	(2) 2016-2017	(3) 2018-2021	(4) Post 2022
Haryana=1	0.873 [0.740]	0.338 [0.817]	-1.979*** [0.598]	-1.665*** [0.402]
Distance to Border	-0.058 [0.055]	-0.065 [0.061]	-0.082 [0.060]	-0.055 [0.040]
Haryana=1 × Distance to Border	-0.120 [0.086]	-0.126 [0.099]	-0.020 [0.070]	0.005 [0.045]
Border-Segments FE	Yes	Yes	Yes	Yes
R^2	0.295	0.316	0.357	0.307
Mean dep. var. Punjab (border area)	4.981	6.138	6.014	3.732
Observations	9608	4804	9608	4804

Notes: This table reports spatial RD estimates of the effect of environmental regulation (Haryana=1) on the number of rice-residue fires per 2×2 km cell during the Kharif (rice-residue burning) seasons. Columns group yearly seasons into four policy periods: pre-2016 (no effective enforcement), 2016–2017 (early Haryana implementation), 2018–2021 (mature Haryana enforcement), and post-2022 (both states implementing additional measures). Each column reports the coefficient on the Haryana indicator, the distance-to-border term, and its interaction with Haryana, estimated with linear polynomials in signed distance to the border and border-segment fixed effects within the preferred RD bandwidth. The dependent variable is the total number of detected rice-residue fires per cell; the mean for Punjab cells in the bandwidth is reported for reference. Standard errors in brackets are two-way clustered on overlapping 100 km² grids. Stars indicate significance in the usual way. A negative coefficient on Haryana=1 indicates fewer fires in Haryana than in nearby Punjab cells within the RD window. See Section 3 for a full description of data sources; for details on the RD specification and clustering, see Section 4.1 and Section 5.6.

Table 2: Policy Impact on Complete Fire Elimination: Probability of No Fires events per cell (4 km²)

	Prob No Fires			
	(1) Pre 2016	(2) 2016-2017	(3) 2018-2021	(4) Post 2022
Haryana=1	-0.050	0.021	0.137**	0.168***
	[0.063]	[0.057]	[0.056]	[0.056]
Distance to Border	0.013**	0.009*	0.007	0.010**
	[0.005]	[0.005]	[0.005]	[0.005]
Haryana=1 × Distance to Border	0.002	0.007	0.003	0.004
	[0.009]	[0.008]	[0.008]	[0.007]
Border-Segments FE	Yes	Yes	Yes	Yes
R^2	0.277	0.281	0.304	0.264
Mean dep. var. Punjab (border area)	0.251	0.168	0.174	0.249
Observations	9608	4804	9608	4804

Notes: This table presents spatial RD estimates of the effect of being in Haryana (Haryana=1) on the probability that a 2×2 km cell experiences complete environmental compliance—no rice-residue fires during the yearly Kharif (rice-residue burning) seasons. The dependent variable is an indicator equal to one if the cell records zero fires in a given Kharif season. Columns group yearly Kharif seasons into the same four policy periods as in Table 1. All specifications include linear distance-to-border terms and border-segment fixed effects within the optimal RD bandwidth. Standard errors in brackets are two-way clustered on overlapping 100 km² grids. Stars indicate significance in the usual way. The mean probability of zero fires on the Punjab side is reported at the bottom of each column. A positive coefficient on Haryana=1 implies that cells in Haryana are more likely to be fire-free than comparable cells in Punjab. See Section 3 for a full description of data sources and outcome construction; for details on the RD specification, see Section 4.1 and Section 5.6.

Table 3: Summary Statistics (Socio-Economic census - agricultural data within 40km of the border)

Variable	Punjab		Haryana	
	N	Mean	N	Mean
Total count of households	41772	233.64	39996	223.55
Share of households that own land	41364	0.43	39468	0.41
Share of households whose main source of income is cultivation	41364	0.36	39468	0.36
Mean number of dwelling rooms owned by household	41364	2.75	39468	2.70
Land with assured irrigation for two crops per household (in acres)	41364	5.02	39468	4.67
Share of households that own mechanized agri equipment	41364	0.16	39468	0.21
Share of households that own irrigation equipment	41364	0.25	39468	0.27
Share of households that own agricultural equipment	41364	0.28	39468	0.30

Notes: This table reports balance tests for pre-regulation socio-economic characteristics drawn from the 2011 Socio-Economic and Caste Census. Each row corresponds to a baseline covariate. Columns report average values at the Punjab–Haryana border within a 40 km bandwidth. The lack of large and systematic discontinuities suggests that, prior to the policy, household characteristics were broadly comparable across the border, supporting the validity of the RD design. See Section 3 for details on the census data; for the RD specification, see Section 4.1 and Section 5.6.

Table 4: Summary Statistics (FAO variables within 50km of the border)

Variable	Punjab within 50km from the border		Haryana within 50km from the border	
	N	Mean	N	Mean
Mean Rice Harvested km2 (2010)	221	0.23	223	0.18
Mean Rice Yield tons/km2 (2010)	221	0.04	223	0.04
Mean Wheat Harvested km2 (2010)	221	0.40	223	0.40
Mean Wheat Yield tons/km2 (2010)	221	0.04	223	0.04
Suitability index Wheat (0 – 10000)	221	4184.62	223	4231.51
Suitability index Rice (0 – 10000)	221	2492.05	223	2544.41

Notes: This table reports balance tests for pre-policy agricultural fundamentals using FAO-GAEZ and related crop data for 2010/2011. Rows cover measures of rice and wheat harvested area, yields, and agro-climatic suitability indices; columns present average values at the Punjab–Haryana border within a 50 km bandwidth. The absence of large, systematic breaks in crop area, productivity, or suitability indicates that agricultural potential and land use were broadly comparable across the border before Haryana’s regulation. See Section 3 for data sources and variable definitions; for the RD design, see Section 4.1 and Section 5.6.

Table 5: Robustness Policy Impact on Number of Fire Events - 20x20km unique spatial cluster

	N. of Fires			
	(1) Pre 2016	(2) 2016-2017	(3) 2018-2021	(4) Post 2022
Haryana=1	0.873 [0.726]	0.338 [0.785]	-1.979*** [0.572]	-1.665*** [0.395]
Distance to Border	-0.058 [0.056]	-0.065 [0.063]	-0.082 [0.059]	-0.055 [0.039]
Haryana=1 $\times$ Distance to Border	-0.120 [0.083]	-0.126 [0.090]	-0.020 [0.067]	0.005 [0.041]
Border-Segments FE	Yes	Yes	Yes	Yes
R^2	0.295	0.316	0.357	0.307
Mean dep. var. Punjab (border area)	4.047	5.144	5.307	3.199
Observations	9608	4804	9608	4804

Notes: This table assesses the robustness of the main fire-count RD estimates to an alternative spatial clustering scheme. It replicates Table 1 but clusters standard errors at the 20 km $\times$ 20 km grid level instead of using two-way 10 km $\times$ 10 km grids. As in the baseline, each column reports the coefficient on Haryana=1, distance-to-border terms, and their interaction for different policy periods defined over yearly Kharif (rice-residue burning) seasons. The similarity of point estimates and standard errors across the two clustering schemes indicates that the baseline two-way 10 km clustering captures the relevant spatial dependence in fire outcomes around the Punjab–Haryana border. Stars indicate significance in the usual way. See Section 3 for fire data sources; for the clustering design and RD specification, see Section 4.1 and Section 5.6.

Table 6: Robustness Policy Impact on Complete Fire Elimination - 20x20km unique spatial cluster

	Prob NO fires			
	(1) Pre 2016	(2) 2016-2017	(3) 2018-2021	(4) Post 2022
Haryana=1	-0.050	0.021	0.137**	0.168***
	[0.059]	[0.063]	[0.061]	[0.060]
Distance to Border	0.013**	0.009*	0.007	0.010*
	[0.006]	[0.005]	[0.005]	[0.006]
Haryana=1 × Distance to Border	0.002	0.007	0.003	0.004
	[0.008]	[0.007]	[0.007]	[0.007]
Border-Segments FE	Yes	Yes	Yes	Yes
R^2	0.277	0.281	0.304	0.264
Mean dep. var. Punjab (border area)	0.251	0.168	0.174	0.249
Observations	9608	4804	9608	4804

Notes: This table provides an analogous robustness check for the probability-of-zero-fires outcome. It reproduces the specifications of Table 2, which estimate the effect of being in Haryana (Haryana=1) on the probability that a 2×2 km cell records no rice-residue fires during the yearly Kharif season, but clusters standard errors at the 20 km×20 km grid level. Coefficients on the Haryana indicator, distance-to-border, and their interaction remain similar in magnitude and precision, suggesting that the baseline clustering scheme already accommodates broader spatial correlation. Stars indicate significance in the usual way. See Section 3 for fire data sources; for details on the RD specification and clustering, see Section 4.1 and Section 5.6.

Table 7: Robustness Policy Impact on Number of Fire Events - Triangular Kernel

	N. of Fires			
	(1) Pre 2016	(2) 2016-2017	(3) 2018-2021	(4) Post 2022
Haryana=1	1.050	0.469	-1.884***	-1.755***
	[0.759]	[0.824]	[0.604]	[0.404]
Distance to Border	-0.082	-0.068	-0.090	-0.030
	[0.073]	[0.079]	[0.070]	[0.042]
Haryana=1 $\times$ Distance to Border	-0.113	-0.152	-0.029	-0.029
	[0.100]	[0.110]	[0.081]	[0.047]
Border-Segments FE	Yes	Yes	Yes	Yes
R^2	0.324	0.351	0.360	0.287
Mean dep. var. Punjab (border area)	4.981	6.138	6.014	3.732
Observations	9188	4594	9188	4594

Notes: This table assesses the robustness of the main fire-count RD estimates to an alternative kernel choice. It replicates Table 1 but estimates the local-linear RD using a triangular kernel rather than the uniform kernel used in the baseline specification. All other features of the design are held fixed, including the symmetric 15 km bandwidth, the donut-hole exclusion of border-intersecting cells, border-segment fixed effects, and the baseline two-way overlapping spatial clustering on 10 km $\times$ 10 km grids. As in the baseline, each column reports the coefficient on Haryana=1, distance-to-border terms, and their interaction for different policy periods defined over yearly Kharif (rice-residue burning) seasons. The similarity of point estimates and standard errors across kernels indicates that the main results are not sensitive to the weighting scheme within the estimation window. Stars indicate significance in the usual way. See Section 3 for fire data sources; for the clustering design and RD specification, see Section 4.1 and Section 5.6.


Table 8: Policy Impact on Pollution Concentration - Wind Regime: Stringent (4 km²)

	Harmonised PM25			
	(1) 2015	(2) 2016-2017	(3) 2018-2021	(4) 2022
Haryana=1	-1.247 [5.679]	-2.922 [3.562]	-9.120** [4.066]	-8.830** [3.926]
Distance to Border	0.172 [0.165]	-0.125 [0.139]	0.211* [0.118]	0.011 [0.149]
Haryana=1 × Distance to Border	-0.212 [0.220]	-0.012 [0.177]	0.096 [0.178]	0.091 [0.193]
Border-Segments FE	Yes	Yes	Yes	Yes
R^2	0.307	0.426	0.248	0.368
Mean dep. var. Punjab (border area)	113.710	113.710	88.832	65.820
Observations	5845	11798	23418	5514

Notes: This table reports spatial RD estimates of the effect of environmental regulation (Haryana=1) on PM_{2.5} concentrations under the “stringent” wind-selection rule used in Figure 7. The dependent variable is average PM_{2.5} ($\mu\text{g}/\text{m}^3$) during the yearly Kharif (rice-residue burning) season for days that satisfy strict conditions on wind direction and speed, measured in 2×2 km cells. Columns group yearly Kharif seasons into the four policy periods used in the fire tables (pre-2016, 2016–2017, 2018–2021, post-2022). Each specification includes linear distance-to-border terms and border-segment fixed effects. Standard errors in brackets are two-way clustered on overlapping 100 km² grids. Stars indicate significance in the usual way. A negative coefficient on Haryana=1 indicates lower PM_{2.5} levels on the Haryana side of the border relative to Punjab under wind conditions most conducive to cross-border pollution transport. See Section 3 for details on the PM_{2.5} and wind data; for the pollution RD specification, see Section 5.6.


9 Additional Figures and Tables

Figure 8: RD Plots – Rabi Wheat-Residue Burning Season: April–May




Notes: This figure presents spatial RD plots for residue burning during the Rabi (wheat-residue burning) season, April–May. Each panel corresponds to one Rabi season in 2016, 2018, 2020, or 2022. The vertical axis shows the average number of wheat-residue fires per 2×2 km cell; the horizontal axis reports signed distance to the Punjab–Haryana border (km), with negative values in Punjab and positive values in Haryana. Grey dots display binned averages, and solid lines show separate local linear fits on either side of the border; the vertical dashed line marks the border. The absence of sizeable discontinuities supports the interpretation that the main policy effect operates through the Kharif rice-residue burning season. See Section 3 for fire data and season definitions; for the RD specification, see Section 4.1 and Section 5.6.

Figure 9: RD Plots – Non-Harvest Months: December–March; June–August




Notes: This figure reports RD plots for fires detected outside the main rice and wheat harvesting seasons, pooling non-harvest months December–March and June–August. Each panel (2016, 2018, 2020, 2022) plots the mean number of fires per 2×2 km cell against signed distance to the Punjab–Haryana border, with binned averages and separate local linear fits on each side; the vertical dashed line marks the border. The lack of systematic discontinuities during these months provides a placebo check that the main RD effects are specific to crop-residue burning seasons. See Section 3 for fire data and seasonal definitions; for the RD design, see Section 4.1 and Section 5.6.

Figure 10: RD Plots – 2011 Socio-Economic Census




Notes: This figure shows spatial RD diagnostics for key pre-policy socio-economic variables from the 2011 Socio-Economic and Caste Census. Each panel plots a different variable (e.g., share of households with cultivation as main income, land ownership, dwelling size, irrigated land, agricultural and mechanised equipment) against signed distance to the Punjab–Haryana border. Grey dots denote binned averages; solid lines are local linear fits on each side of the border; the vertical dashed line indicates the border. The absence of large, systematic discontinuities provides visual evidence of baseline balance in household conditions across the border. See Section 3 for data construction; for the RD design, see Section 4.1 and Section 5.6.

Figure 11: RD Plots – FAO-GAEZ 2010/2011 Rice and Wheat Indicators



Notes: This figure reports spatial RD plots for pre-policy agricultural indicators derived from FAO-GAEZ and related crop data for 2010/2011. Panels show rice and wheat harvested area, yields, and agro-climatic suitability indices plotted against signed distance to the Punjab–Haryana border. The vertical axis reports the outcome level (e.g., hectares, tons/ha, suitability index); the horizontal axis reports distance to the border in kilometres (negative values in Punjab, positive values in Haryana). Binned means and separate local linear fits are shown on each side of the border, with the vertical dashed line at distance zero. The absence of large discontinuities indicates that underlying agricultural potential and crop allocation were broadly similar across the border prior to the policy. See Section 3 for data construction; for the RD design, see Section 4.1 and Section 5.6.

Figure 12: Event-Study Estimates of the Regulation Effect on $PM_{2.5}$ Under Different Wind Conditions



Notes: This figure shows event-study RD estimates of the impact of Haryana’s crop-burning regulation on surface $PM_{2.5}$ concentrations under alternative wind-selection rules. Each panel reports yearly Kharif (rice-residue burning) season RD coefficients for $PM_{2.5}$, with the vertical axis in $\mu g/m^3$ and the horizontal axis labelling Kharif seasons by calendar year. Panels differ in the subset of days used: less stringent wind filters, parallel-only winds, and all days irrespective of wind. $PM_{2.5}$ is measured using harmonised satellite products and aggregated to seasonal means within each 2×2 km cell. Wind variables are constructed from ERA5; 95% confidence intervals are based on two-way spatial clustering. See Section 3 for $PM_{2.5}$ and wind data sources; for the pollution RD specification, see Section 5.6.

Table 9: Policy Impact on Pollution Concentration - Wind Regime: Less Stringent (4 km²)

	Harmonised PM25			
	(1) 2015	(2) 2016-2017	(3) 2018-2021	(4) 2022
Haryana=1	0.622 [4.246]	-2.092 [4.194]	-7.212*** [2.600]	-6.088* [3.643]
Distance to Border	0.081 [0.123]	-0.114 [0.131]	0.059 [0.073]	0.001 [0.119]
Haryana=1 × Distance to Border	-0.254 [0.167]	0.172 [0.174]	0.132 [0.122]	0.063 [0.173]
Border-Segments FE	Yes	Yes	Yes	Yes
R^2	0.415	0.543	0.345	0.105
Mean dep. var. Punjab (border area)	112.096	112.096	85.432	54.327
Observations	5998	11953	23628	5901

Notes: This table mirrors Table 8 but relaxes the wind-selection criteria (“less stringent” condition). The dependent variable is PM_{2.5} during the yearly Kharif (rice-residue burning) season in 2×2 km cells on days that satisfy looser thresholds for wind direction and speed. Columns group yearly Kharif seasons into policy periods. All specifications include linear distance-to-border terms and border-segment fixed effects. Standard errors in brackets are two-way clustered on overlapping 100 km² grids. Stars indicate significance in the usual way. Comparing estimates across the “stringent” and “less stringent” samples assesses the sensitivity of the pollution discontinuity to alternative definitions of days with upwind transport. See Section 3 for PM_{2.5} and wind data; for the RD specification, see Section 5.6.

Table 10: Policy Impact on Pollution Concentration - Wind Regime: Parallel (4 km²)

	Harmonised PM25			
	(1) 2015	(2) 2016-2017	(3) 2018-2021	(4) 2022
Haryana=1	0.633 [4.252]	-2.048 [4.194]	-7.072*** [2.600]	-6.071* [3.643]
Distance to Border	0.079 [0.123]	-0.114 [0.131]	0.062 [0.073]	0.001 [0.119]
Haryana=1 × Distance to Border	-0.257 [0.167]	0.170 [0.174]	0.122 [0.121]	0.062 [0.173]
Border-Segments FE	Yes	Yes	Yes	Yes
R^2	0.414	0.542	0.344	0.105
Mean dep. var. Punjab (border area)	112.104	112.104	85.216	54.327
Observations	5998	11953	23628	5901

Notes: This table reports spatial RD estimates for PM_{2.5} when restricting to days in which the wind blows approximately parallel to the Punjab–Haryana border (“parallel only” condition). The dependent variable is mean PM_{2.5} during the yearly Kharif (rice-residue burning) season in 2×2 km cells. Columns correspond to policy periods, and the specification matches that in the other PM_{2.5} tables, including linear distance-to-border terms and border-segment fixed effects. Standard errors in brackets are two-way clustered on overlapping 100 km² grids. Stars indicate significance in the usual way. This therefore serves as a robustness and placebo check for the pollution RD effect. See Section 3 for PM_{2.5} and wind data; for methodological details, see Section 5.6.

Table 11: Policy Impact on Pollution Concentration - Wind Regime: All Days (4 km²)

	Harmonised PM25			
	(1) 2015	(2) 2016-2017	(3) 2018-2021	(4) 2022
Haryana=1	-0.958 [1.398]	-2.582 [2.209]	-3.393 [2.092]	-1.980 [1.362]
Distance to Border	0.009 [0.037]	-0.064 [0.056]	0.016 [0.048]	-0.030 [0.042]
Haryana=1 $\times$ Distance to Border	-0.022 [0.059]	0.031 [0.092]	0.041 [0.095]	0.018 [0.070]
Border-Segments FE	Yes	Yes	Yes	Yes
R^2	0.741	0.689	0.569	0.623
Mean dep. var. Punjab (border area)	109.204	109.204	85.348	54.854
Observations	6041	12097	24224	6043

Notes: This table presents RD estimates for PM_{2.5} using all Kharif (rice-residue burning) season days irrespective of wind direction or speed (“all days” condition). The dependent variable is mean PM_{2.5} in 2 $\times$ 2 km cells; columns are organised by policy period. Standard errors in brackets are two-way clustered on overlapping 100 km² grids. Stars indicate significance in the usual way. Comparing these coefficients with those from the wind-conditioned samples helps to separate the contribution of local burning from broader regional pollution dynamics. See Section 3 for PM_{2.5} data construction; for the pollution RD design, see Section 5.6.

Table 12: Robustness Policy Impact on Pollution Concentration - Triangular Kernel

	Harmonised PM25			
	(1) 2015	(2) 2016-2017	(3) 2018-2021	(4) 2022
Haryana=1	-1.101	-3.610	-8.491**	-8.440**
	[5.658]	[3.462]	[3.908]	[3.735]
Distance to Border	0.184	-0.074	0.220**	-0.004
	[0.148]	[0.124]	[0.106]	[0.129]
Haryana=1 $\times$ Distance to Border	-0.239	-0.056	0.064	0.111
	[0.208]	[0.162]	[0.166]	[0.174]
Border-Segments FE	Yes	Yes	Yes	Yes
R^2	0.286	0.436	0.257	0.380
Mean dep. var. Punjab (border area)	113.710	113.710	88.832	65.820
Observations	7156	14430	28585	6771

Notes: This table assesses the robustness of the main PM_{2.5} concentration RD estimates to an alternative kernel choice. It replicates Table 8 but estimates the local-linear RD using a triangular kernel rather than the uniform kernel used in the baseline specification. The dependent variable is average PM_{2.5} ($\mu\text{g}/\text{m}^3$) during the yearly Kharif (rice-residue burning) season for days that satisfy strict conditions on wind direction and speed, measured in 2×2 km cells. Columns group yearly Kharif seasons into the four policy periods used in the fire tables (pre-2016, 2016–2017, 2018–2021, post-2022). Each specification includes linear distance-to-border terms and border-segment fixed effects. Standard errors in brackets are two-way clustered on overlapping 100 km² grids. The similarity of point estimates and standard errors across kernels indicates that the results are not sensitive to the weighting scheme within the estimation window. Stars indicate significance in the usual way. A negative coefficient on Haryana=1 indicates lower PM_{2.5} levels on the Haryana side of the border relative to Punjab under wind conditions most conducive to cross-border pollution transport. See Section 3 for details on the PM_{2.5} and wind data; for the pollution RD specification, see Section 5.6.

Table 13: Baseline Demographic and Disease Burden Indicators for Haryana

Variable	Source	Value	Reference year
Panel A: Demography (Government Sources)			
Population	National Commission on Population, Population Projections for India and States 2011–2036	29.5 million	2021
Crude birth rate (per 1,000)	Health Dossier 2021: Haryana	20.1	2019
Crude death rate (per 1,000)	Health Dossier 2021: Haryana	5.9	2019
Infant mortality rate, IMR (per 1,000 live births)	Health Dossier 2021: Haryana	27	2019
Under-five mortality rate, U5MR (per 1,000 live births)	Health Dossier 2021: Haryana	36	2019
Panel B: Burden of Disease (GBD 2018, Haryana)			
Total PM _{2.5} -attributable deaths (all ages)	IHME, GBD 2018	42,700	2018
Total PM _{2.5} -attributable DALYs (all ages)	IHME, GBD 2018	1.32 million	2018
Infant PM _{2.5} -attributable deaths (under age 1)	IHME, GBD 2018	3,440	2018
Infant PM _{2.5} -attributable DALYs (under age 1)	IHME, GBD 2018	309,000	2018
Under-five PM _{2.5} -attributable deaths (age 0–4)	IHME, GBD 2018	3,950	2018
Under-five PM _{2.5} -attributable DALYs (age 0–4)	IHME, GBD 2018	352,000	2018

Notes: Panel A compiles demographic indicators from official government sources, providing the population baseline used in the main analysis. Panel B reports Global Burden of Disease (GBD 2018) estimates for Haryana, including PM_{2.5}-attributable deaths and disability-adjusted life years (DALYs). These raw GBD values form the epidemiological inputs for the health impact calculations developed in Appendix 11.

10 Detailed Construction of Harmonised Pollution Data

This appendix provides additional detail on the construction of the harmonised PM_{2.5} series used in the main analysis. GHAP is available from 2017 onward, whereas LGHAP provides a longer record. To place the two products on a common scale, I estimate a linear mapping from LGHAP to GHAP over their overlapping years.

For each wind regime $r \in \{\text{stringent, less stringent, parallel, all}\}$ and for $t \in [2017, 2021]$, I estimate

$$\text{PM}_{it}^{\text{GHAP},r} = \alpha_r + \beta_r \text{PM}_{it}^{\text{LGHAP},r} + \gamma_t + \varepsilon_{it}, \quad (5)$$

where $\text{PM}_{it}^{\text{GHAP},r}$ and $\text{PM}_{it}^{\text{LGHAP},r}$ are cell-year means of PM_{2.5} for grid cell i and year t under regime r , and γ_t are year fixed effects. The coefficients $(\hat{\alpha}_r, \hat{\beta}_r)$ summarise the level and scale differences between the two products.

Figure 13 documents the fit in the all-days sample. Panel (a) plots binned means of GHAP against LGHAP; the points lie close to both the fitted line and the 45° line, indicating tight co-movement. Panel (b) reports the corresponding regression:

$$\widehat{\text{PM}}^{\text{GHAP}} = 21.97 + 0.822 \times \text{PM}^{\text{LGHAP}},$$

with both coefficients statistically significant at the 1% level and R^2 around 0.8. Across wind regimes, the estimated slopes cluster near 0.8 and the R^2 values remain high, so within each regime LGHAP and GHAP share similar spatial and temporal patterns up to a constant shift and rescaling.


These estimates imply that, conditional on year fixed effects, LGHAP can be interpreted as a noisy but linearly related measure of the same underlying PM_{2.5} field captured by GHAP. Because $\hat{\beta}_r$ is close to one in all regimes, the transformation preserves the ordering and the magnitude of cross-border differences while placing all years on a common, GHAP-consistent scale.

Using the estimated coefficients, I construct a harmonised PM_{2.5} series that uses GHAP when available and the mapped LGHAP values otherwise:

$$\text{PM}_{it}^{\text{Harmonised},r} = \begin{cases} \text{PM}_{it}^{\text{GHAP},r}, & \text{if } t \geq 2017, \\ \hat{\alpha}_r + \hat{\beta}_r \text{PM}_{it}^{\text{LGHAP},r}, & \text{if } t < 2017. \end{cases} \quad (6)$$

All pollution outcomes in the main text are expressed in this harmonised PM_{2.5} scale.

Figure 13: Affine Mapping from LGHAP to GHAP Scale in the Overlap (2017–2021)



Notes: Panel (a) reports a binned scatter comparison between GHAP and LGHAP PM_{2.5}. Each point shows the mean GHAP and LGHAP PM_{2.5} within the bin. The dashed line is the 45-degree line; the solid line shows the affine mapping $\widehat{\text{GHAP}} = \hat{\alpha} + \hat{\beta} \text{LGHAP} + \gamma_t$ estimated with year fixed effects. Panel (b) reports the corresponding regression coefficients. Robust standard errors in brackets. Stars indicate significance in the usual way. Together, the panels document an affine relationship between the two products, motivating the harmonisation procedure used to construct a single PM_{2.5} series for 2015–2022.

11 Detailed Back-of-the-Envelope Calculations for Health Impacts

This appendix documents, step by step, all back-of-the-envelope transformations used to obtain avoided deaths, DALYs, healthy days, and monetary valuations in Section 5.9. Each reported quantity is a mechanical function of (i) the RDD-implied change in $\text{PM}_{2.5}$ exposure, (ii) baseline demographic quantities, and (iii) external CRFs or elasticities.

C.1 Avoided deaths

i. Demographic baselines

Haryana's total population is fixed at

$$N = 29.5 \text{ million.}$$

Crude birth and death rates per 1,000 population (SRS 2018) are:

$$\text{CBR} = 20.1, \quad \text{CDR} = 5.9.$$

Infant and under-five mortality rates per 1,000 live births are:

$$\text{IMR} = 27, \quad \text{U5MR} = 36.$$

Births and deaths.

$$B = N \times \frac{\text{CBR}}{1,000} = 29.5 \times 10^6 \times 0.0201 \approx 592,950.$$

$$D = N \times \frac{\text{CDR}}{1,000} = 29.5 \times 10^6 \times 0.0059 \approx 174,050.$$

Infant and under-five deaths.

$$D_{\text{inf}} = B \times \frac{27}{1,000} \approx 16,010, \quad D_{\text{u5}} = B \times \frac{36}{1,000} \approx 21,346.$$

ii. Regulation-induced $\text{PM}_{2.5}$ changes

The spatial RDD estimates imply:

$$\Delta PM_{\text{season}} = 9 \mu\text{g}/\text{m}^3 \quad (\text{Kharif season}).$$

Since the Kharif season lasts approximately two months:

$$\Delta PM_{\text{annual}} = \Delta PM_{\text{season}} \times \frac{2}{12} = 9 \times \frac{1}{6} = 1.5 \mu\text{g}/\text{m}^3.$$

iii. Mortality impacts

Functional form. For all-cause mortality, I implement the log-linear CRF derived from Ayres (2010), Chen and Hoek (2020). If the relative risk is RR_{10} per 10 $\mu\text{g}/\text{m}^3$, then the proportional change in mortality for an exposure change ΔPM is approximated by

$$\% \Delta M \approx \frac{\text{RR}_{10} - 1}{10} \times \Delta PM,$$

which is accurate for the small exposure change considered here. Avoided deaths equal baseline deaths multiplied by $\% \Delta M$.

(a) All-cause mortality, all ages

Using a conservative CRF ($RR = 1.06$ per $10 \mu\text{g}/\text{m}^3$):

$$\% \Delta M_{\text{cons}} = \frac{1.06 - 1}{10} \times 1.5 = 0.006 \times 1.5 = 0.009,$$

i.e. a 0.9% mortality reduction.

Using a central CRF ($RR = 1.08$ per $10 \mu\text{g}/\text{m}^3$):

$$\% \Delta M_{\text{cent}} = \frac{1.08 - 1}{10} \times 1.5 = 0.008 \times 1.5 = 0.012,$$

i.e. a 1.2% mortality reduction.

Avoided deaths (all ages).

$$\Delta D_{\text{cons}} = 0.009 \times 174,050 \simeq 1,566, \quad \Delta D_{\text{cent}} = 0.012 \times 174,050 \simeq 2,089.$$

(b) Infant and under-five mortality

Elasticities from Dipoppa and Gulzar (2024):

- +7.8% infant mortality per $10 \mu\text{g}/\text{m}^3$ $\text{PM}_{2.5}$ - +8.5% under-five mortality per $10 \mu\text{g}/\text{m}^3$

Scaling to $\Delta PM_{\text{annual}} = 1.5$:

$$\% \Delta M_{\text{inf}} = 0.078 \times \frac{1.5}{10} = 0.0117,$$

$$\% \Delta M_{\text{u5}} = 0.085 \times \frac{1.5}{10} = 0.01275.$$

Avoided infant and under-five deaths.

$$\Delta D_{\text{inf}} = 0.0117 \times 16,010 \simeq 187, \quad \Delta D_{\text{u5}} = 0.01275 \times 21,346 \simeq 272.$$

C.2 Gains in health days

i. DALYs per death

GBD Haryana 2018:

$$\theta_{\text{all}} = \frac{1.32 \times 10^6}{42,700} \simeq 31 \quad (\text{all ages}),$$

$$\theta_{\text{child}} = \frac{309,000}{3,440} \simeq 90 \quad (\text{infants/under-five}).$$

ii. DALYs saved

(a) All ages

$$\Delta \text{DALY}_{\text{all,cons}} = 31 \times 1,566 \simeq 48,500, \quad \Delta \text{DALY}_{\text{all,cent}} = 31 \times 2,089 \simeq 64,800.$$

$$\Rightarrow \quad \Delta \text{DALY}_{\text{all}} \approx 48,000\text{--}65,000, \quad \overline{\Delta \text{DALY}}_{\text{all}} \simeq 56,700.$$

(b) Infants and under-five children

$$\Delta\text{DALY}_{\text{inf}} = 90 \times 187 \simeq 16,900, \quad \Delta\text{DALY}_{\text{u5}} = 90 \times 272 \simeq 24,500.$$

iii. Healthy days gained

(a) All ages

$$\begin{aligned} \frac{56,700}{29.5 \times 10^6} &= 0.00192 \text{ DALYs/person/year,} \\ 0.00192 \times 365 &\simeq 0.70 \text{ days per person per year.} \end{aligned}$$

(b) Infants

Infant population = $N_{\text{inf}} = 592,950$:

$$\begin{aligned} \frac{16,900}{592,950} &= 0.0285 \text{ DALYs/infant/year,} \\ 0.0285 \times 365 &\simeq 10.4 \text{ healthy days per infant per year.} \end{aligned}$$

(c) Under-five children

Under-five population = $5B = 2.96$ million:

$$\begin{aligned} \frac{24,500}{2.96 \times 10^6} &= 0.00828 \text{ DALYs/child/year,} \\ 0.00828 \times 365 &\simeq 3.0 \text{ healthy days per under-five child per year.} \end{aligned}$$

C.3 Monetary valuation

i. DALY-based valuation (WHO–CHOICE thresholds)

Let $\Delta\text{DALY}_{\text{all}}$ denote annual DALYs saved for all ages. Using $\overline{\Delta\text{DALY}}_{\text{all}} = 56,700$ and cost-effectiveness thresholds $v \in \{2,000, 6,000\}$ USD per DALY (WHO–CHOICE benchmarks for India, as in Bhat et al. (2023)), annual monetary benefits are

$$B^{\text{DALY}}(v) = v \times \overline{\Delta\text{DALY}}_{\text{all}}.$$

Thus,

$$B^{\text{DALY}}(2,000) = 56,700 \times 2,000 \approx 113 \text{ million USD,}$$

$$B^{\text{DALY}}(6,000) = 56,700 \times 6,000 \approx 340 \text{ million USD.}$$

ii. VSL-based valuation

Let ΔD denote annual avoided deaths (all ages). Using a conservative India VSL of 688,000 USD (Jack et al., 2025; Majumder and Madheswaran, 2018), annual mortality benefits are

$$B^{\text{VSL}} = 688,000 \times \Delta D.$$

For $\Delta D \in [1,600, 2,100]$,

$$B^{\text{VSL}}(1,600) \approx 1.10 \text{ billion USD}, \quad B^{\text{VSL}}(2,100) \approx 1.45 \text{ billion USD}.$$

iii. Per-acre normalisation

Let $A = 650,900$ denote acres estimated to have been burned in Haryana in Kharif 2018 (Kumar et al., 2019). Benefits per acre are

$$b^{\text{DALY}}(v) = \frac{B^{\text{DALY}}(v)}{A}, \quad b^{\text{VSL}} = \frac{B^{\text{VSL}}}{A}.$$

Using the values above yields approximately 170–520 USD/acre under DALY-based valuation and about 1,700–2,200 USD/acre under VSL-based valuation.

School of Economics and Finance



This working paper has been produced by
the School of Economics and Finance at
Queen Mary University of London

Copyright © 2026 The Author(s). All rights reserved.

School of Economics and Finance
Queen Mary University of London
Mile End Road
London E1 4NS
Tel: +44 (0)20 7882 7356
Fax: +44 (0)20 8983 3580
Web: www.econ.qmul.ac.uk/research/workingpapers/