
MEASURING CROP RESIDUE BURNING IN PAKISTAN: METHODS, ANALYSIS, AND NEXT STEPS*

B. Kelsey Jack, Muhammad Usman Naeem, Ben Moscona, Michael Greenstone

Preliminary White Paper

1 Background

1.1 Crop Residue Burning

In Pakistan, agricultural emissions are responsible for an estimated 20% of air pollution in Punjab Province [1]. The emissions are primarily a result of the burning of rice stubble after harvest by farmers engaged in rice-wheat cropping. Mechanization with combine harvesters allows farmers to harvest rice more efficiently than traditional methods, but leaves rice straw, or stubble, which must be managed or removed during a narrow two- to three-week time window between rice harvest and wheat planting [2]. Farmers commonly burn the stubble, as other methods are more costly, time-consuming, and labor-intensive. However, stubble burning releases harmful pollutants, and may negatively impact soil quality and productivity through increased erosion and soil nutrient depletion [3]. Burning stubble emits particulate matter (PM10 and PM2.5) and other air pollutants harmful to human health, and black carbon and greenhouse gases that contribute to climate change [4]. The Punjab government banned stubble burning in 2019, but enforcement is weak and the practice persists, due in part to a lack of accessible and cost-effective alternatives.

1.2 Equipment Subsidies in Pakistan

In response to the damaging effects of rice stubble burning on air and soil quality, a new generation of mechanized equipment has been developed. These machines, called Rice Straw Shredders and Happy Seeders (which we refer to as RSS-HS), shred and incorporate rice stubble into soil. Research has shown that these machines can be effective at mitigating the externalities presented by rice stubble burning. For example, a study on crop burning in India found that RSS-HS reduces GHG emissions per hectare by more than 78% relative to burning [5]. Over the medium to long term, use of the machines can increase productivity via improvements in soil fertility [6], as well as through the improvement of sowing techniques [7].

For the machines to mitigate externalities or improve productivity, they have to be used. Adoption rates for RSS-HS machines by Punjab farmers are very low. This can be explained in part by the high upfront cost of purchasing the

*We thank Precision Development, Centre for Economic Research in Pakistan (CERP), Punjab Agriculture Department (AgriPunjab), and Agricultural Robotics Lab at the Lahore University of Management Sciences (LUMS) for collaboration on data collection and measurement. Fatima Azmat, Danae Roumis, Awaid Yasin, Leslie Hodel, Daniel Posmik, and Colton Lapp at the Energy & Environment Lab (E&E) at the University of Chicago provided excellent research assistance. Funding was provided by the Agricultural Technology Adoption Initiative (ATAI), the International Growth Centre (IGC), and the Development Economics Research Fund (DERF) at the University of Chicago. Olga Rostapshova provided invaluable project leadership. Please contact Kelsey Jack (kelseyjack@ucsb.edu) or Muhammad Usman Naeem (usman.naeem@tufts.edu) with questions or feedback.

machines [5]. Additionally, the fact that most of the benefits of adoption accrue to others further undermines adoption. Accordingly, the Punjab Agriculture Department (AgriPunjab) launched the “Mechanised Management of Rice Crop Residue” (MMRCR) subsidy program in 2020 to encourage and assist farmers with the purchase of the machines. This program provides 80% subsidies for the purchase of the two RSS-HS machines. The estimated market price of a Rice Straw Shredder at the time was U\$2,040 (and subsidy was capped at U\$1,632) and that of a Happy Seeder was U\$2,122 (with subsidy capped at U\$1,698). Program eligibility criteria included a commitment to rent the equipment to other farmers for at least four years. Subsidies were awarded through lotteries, with winners randomly chosen on a stratified geographical level (markaz). The program trained lottery winners to use the equipment, while farmers were responsible for repair and maintenance. We leverage the randomization of the lottery recipients to evaluate the impact of the equipment subsidy program on crop burning, with an emphasis at this stage of the project on measurement of this primary outcome of interest.

AgriPunjab has initiated two phases of the MMRCR equipment subsidy program to date. In this paper, we focus on the first phase (Phase 1) which began in 2020. Within Phase 1, AgriPunjab administered two separate application waves, both of which involved lotteries conducted at the district level. In the first wave, only one farmer was selected from each markaz (an administrative unit below the district level). In Wave 2, however, there were between two and three farmers per markaz selected. Each wave consisted of a total of 15 lotteries, with each lottery corresponding to one of 15 unique rice farming districts in Punjab. These districts represent 80% of the 4.6 million acres of rice cultivation in Punjab. In wave 1, 134 farmers won the lottery out of the 938 total applicants. In wave 2, 366 of the 1391 applicants won. Together, the two lotteries resulted in 495 winning farmers and 1834 plausibly comparable farmers who did not win the lottery. The very high application rates are *prima facie* evidence of high demand for the equipment among farmers in Punjab.

The lotteries were advertised through local AgriPunjab field wing offices, newspapers, and the radio. Farmers could apply for Rice Straw Shredder, Happy Seeder, or both. The lottery process itself was conducted in-person and was open to the public. Up to six farmers were drawn for each machine pair per markaz and recorded as draw numbers 1-6. The rationale behind this was that if a winning farmer later refused to purchase the machine(s) then AgriPunjab would use one of the back-ups (i.e., if draw number one refused, then draw number two was contacted/offered the machine(s)). Table 1 shows the only observable covariate available in the administrative data – land area – is mostly balanced in Waves 1-2, i.e., treatment farmers (Draw 1, the omitted category) are not systematically different from control farmers (Draws 2-6).² All lottery winning farmers from Phase 1 received subsidy allotment letters which they used to purchase machinery between September 2020 and January 2021 (Wave 1) or between September 2021 and January 2022 (Wave 2). Following the completion of the first phase of the program, AgriPunjab launched a second phase of the program which extended the program period another 5 years.

Whether these subsidies were sufficient to induce adoption and use of the machines – and to reduce burning – is an open question. A small literature investigates the determinants of farmers’ decisions to burn, including the availability of local labor (e.g., [8]; [9]). Other studies have attempted to understand the impacts of government subsidies on promoting technology adoption ([10]; [11]), as well as which factors influence take-up of the subsidies ([12]; [2]). This project attempts to develop and apply the measurement necessary to evaluate the impact of AgriPunjab’s subsidy program.

²A key difference in lottery between the two waves was that in Wave 1 only one machine pair was allocated per markaz meaning there were up to six draws per machine pair or markaz (1 machine pair per markaz x 6 draws per machine pair), whereas in Wave 2 between two and three machine pairs were kept per markaz that resulted in up to 18 draws across machine pairs per markaz (2-3 machine pairs per markaz x 6 draws per machine pair).

	(1) Area	(2) Area
Wave 1 Draw 2	11.808 (18.868)	
Wave 1 Draw 3	2.203 (26.212)	
Wave 1 Draw 4	-19.115 (29.626)	
Wave 1 Draw 5	9.068 (23.758)	
Wave 1 Draw 6	-5.478 (30.520)	
Wave 2 Draw 2		-56.224* (29.281)
Wave 2 Draw 3		13.785 (40.297)
Wave 2 Draw 4		-24.767 (48.123)
Wave 2 Draw 5		-30.331 (73.929)
Wave 2 Draw 6		82.826 (57.387)
Constant	212.995*** (18.664)	225.375*** (11.731)
Observations	356	522

Table 1: Waves 1-2 Lottery Rank Balance

Notes: The table shows results from the regression: $Area_i = a + b * \mathbf{Draw}_i + c * Markaz_i + u_i$, where i depicts farmer, $Area$ measures farmer's baseline land size, $\mathbf{Draw}$ is a vector of dummies of farmer draw (e.g., farmer drawn as 2nd in lottery, etc.), and $Markaz$ captures markaz fixed effects as lottery was stratified by markaz. There were up to six draws in Wave 1 and up to 18 draws in Wave 2; and while we show all six draws for Wave 1 in the table, we exclude Draws 7-18 for Wave 2 to save space. Standard errors clustered at markaz are shown in parentheses.

* : ($p < 0.10$), ** : ($p < 0.05$), *** : ($p < 0.01$)

1.3 Project Scope and Objectives

This pilot project has two main goals: (1) test alternative approaches to developing predictive measures of farmer-level burning outcomes through better on-the-ground measurement of burning and (2) estimate preliminary treatment effects on burning from the equipment subsidies. Figure 1 shows the overall project timeline.




Figure 1: Project Timeline

2 Measuring Subsidy Impacts

2.1 Overview

Estimating treatment effects of an intervention on crop residue burning requires knowing burning outcomes for both treated and control units. In the program we study, subsidized equipment is randomly assigned at the farmer level. However, farmers could rent out the equipment, meaning that the treated unit might extend beyond the farmer's own fields to a catchment area corresponding to the rental market, likely relatively nearby. In many settings, including ours, asking farmers to self-report burning in a survey may not generate reliable data because residue burning is illegal. Additionally, surveying every farmer in the program, along with all farmers who may have rented from the treatment or control farmer, is time consuming and expensive. Remotely sensed data can help overcome both challenges, providing an objective measure of burning across relevant spatial extents. A major hurdle, however, is the collection of relevant training or ground truthing data to build a predictive model to interpret the remote sensing data. We discuss approaches to building a training dataset at the plot level and aggregate level in this section before presenting preliminary treatment effect estimates in the next section.

2.2 Plot Level Measurement of Burning

Leveraging satellite imagery for crop residue burning at the plot level requires development of new predictive models. Existing products that detect active fires and burn scars are of too low a resolution to pick up activities on the scale of a typical farmer's plot in our sample. However, new models requires burn and no-burn labels for training purposes. These can be obtained by manually analyzing high resolution imagery, however, actual ground-based measurements are preferable for differentiating between similar visual signals from e.g., tilling and burning ([13]; [14]). Difficulty in collecting reliable measurements on the ground is compounded by the fact that negative (no-burn) labels require observing that no burning occurs over the entire agricultural season (specifically the time period between harvesting and tilling, but that is often hard to observe). Observing that a field is not burnt at one moment in time does not rule out that it was burnt earlier (with short-lived burning indicators) or will be burnt later. Positive labels can be easier to collect since a field observed as burnt is a true positive label.

We explore three alternative approaches to collecting ground based measurement of burning: (1) in-person spot checks on a sample of farmers from Wave 1 of the lottery, (2) field based sensors, and (3) soil samples. (2) and (3) were conducted on test plots using controlled burns.

2.2.1 Ground Based Measurement of Burning

Spot Checks Two waves of spot checks were conducted by surveyors to determine whether or not farmers were burning rice stubble. Spot checks consisted of surveyors visiting lottery participants' homes and assessing whether or not crop burning had occurred. The first wave of spot checks took place in 2021 and targeted both treated and control farmers from the Wave 1 lottery. The second round of spot checks occurred in the subsequent harvest cycle in 2022 and also targeted treated and control farmers from the Wave 1 lottery. The 2022 spot checks used an updated survey protocol that addressed data collection issues from the first round of spot checks. Throughout the spot checks, surveyors answered a series of yes or no questions based on visual observation during each weekly visit to a farmer. For each plot, the observation checklist captured signs of burning, Rice Shredder or Happy Seeder equipment usage, other management techniques, or conditions indicating residue management had not yet occurred. The survey results were then used to construct binary burn or no-burn outcome variables for each farmer. This process was modeled on data collection from work in [13].

The 2021 spot check surveys were conducted between late November and December of 2021. Because these surveys were conducted so late in the harvest cycle, most farmers had completed tilling by the time the plot visits took place.

The 2022 spot check surveys were conducted earlier in the harvest cycle, between October and November of 2022. As a result, we rely on these surveys as our preferred ground validation of burning.

Objective

Surveyors were dispatched by Precision Development and instructed to survey a list of farmers provided by the researchers. Trusted local parties were contacted to determine the most likely time of rice harvest. Surveyors were then dispatched to farmers homes and, conditional on survey consent, were asked questions about their farming practices. During the first visit, surveyors inquired about the farmer's crops and cultivation timelines, and asked them to identify which plots they cultivate. If the opportunity to burn rice stubble had not yet occurred (rice had not been harvested yet or stubble remained), surveyors left and were instructed to attempt to return up to 5 more times. If rice had been harvested, visits were then made to farmer's plots in an attempt to assess whether or not burning had occurred. Surveyors looked for characteristics of burning such as ash and smoldering, as well as for characteristics of machine use such as equipment presence or shredded stubble. Surveyors then walked the perimeter of each plot, taking GPS points at each corner. These GPS data points were collected to be used in the satellite imagery predictive model.

Spot check visits were meant to be unannounced. Often times, however, surveyors had to call farmers ahead of the visit to receive additional directions to locate the farmer. In an attempt to maximize participation in spot check surveys, surveyors did not announce the specific purpose of their visit.

The two rounds of spot checks in 2021 and 2022 differed in a few ways. First, 2022 spot checks allowed for non-farmers (i.e. family members) to be interviewed in the event that farmers were not present. Second, the instructions for collecting GPS data were improved in 2022. Third, the 2022 spot checks were completed at a more appropriate time in the harvest cycle to more accurately be able to detect rice stubble burning. Finally, while spot checks 1 included a single plot visit, spot checks 2 increased plot visit frequency to up to five visits (one per week) and therefore had a higher probability of detecting burning.

The study obtained IRB approval prior to implementation. All enumerators underwent human subjects protections and data security training.

Approach and Data Collection

In the 2021 spot checks, 578 out of 614 farmers were successfully located and 550 consented to the surveys (95 treatment and 455 control). Due to poor timing of surveys in the 2021 spot checks, however, it was difficult to effectively detect crop burning. Of the 519 farmers who had planted rice, 100% had harvested rice more than 5 days ago. Due to this lag between the survey and the rice planting, it was unlikely for surveyors to observe indicators of burning. Regardless, this data collection exercise suggested statistically significant differences between visual indicators of equipment use by treatment and control farms (6.9% treatment farms indicated Happy Seeder use vs 2.3% control farms, and 7.5% treatment farms indicated Rice Straw Shredder use vs 1.9% control farms).

In the 2021 spot checks, surveyors also faced difficulties collecting GPS points for the plots, as some farmers did not allow them to walk the entire field perimeter or enter the field. Additionally, plots that were still flood irrigated posed obstacles to collecting GPS points. Compounding these issues were technical difficulties using survey devices to collect GPS information. Surveyors attempted to create plot boundaries by walking across a plot and taking coordinates every couple of steps, capturing 5-6 points per plot. However, the only reliable GPS information that was reliably collected for the 2021 spot checks was a single GPS point for most farms visited. These issues were remedied in the 2022 spot checks by instructing surveyors to pre-warm survey devices and also to set higher GPS accuracy threshold at 5 meters. As a result, the 2022 spot checks obtained GPS points for all vertices of each plot during every visit.

For the second round of spot checks in 2022, 93 lottery winners were targeted, located, and provided consent to be surveyed; additionally, 295 non-winning applicants were targeted and 294 provided consent to be interviewed. In total, 1,366 plot-visits were made throughout the 2022 spot checks. The final spot check data from 2022 is at the farmer level and can detect if a farmer was ever caught burning rice on any of its plots we visited.

Results

Using the spot check survey data, we created multiple aggregated variables that captured occurrences of burning with two objectives: (1) to serve as labels in a predictive model using remote sensing data and (2) as an outcome with which to measure treatment effects. In order to create labels with which to train (fit) the satellite-based monitoring model, we constructed binary burn or no burn classifications for each plot-visit from the set of spot check survey questions. We then further aggregated these weekly, plot-level observations to a plot-level cross-section.

We assigned a "Yes" burn label to plots where the following items were observed:

- Active burning
- Smoldering
- Smoke rising from a field
- Partially burnt vegetation
- Fully burnt vegetation
- Ash and straw removed from the soil

We assigned a "No" burn label to plots where the following items were observed:

- Straw baler machine
- Rice Shredder
- Happy Seeder
- Wheat growing in rows with rice straw between

To aggregate across weeks to the farmer level, we assigned a "Yes" burn label to plots if any of the weekly spot checks were labeled as a burn, and a "No" burn label if plots were never labeled as a burn. Figure 2 compares the proportion of spot checks in which burns were spotted: 2021 vs 2022. The main conclusion is that in 2021, almost all spot checks were burns on the first observation (we observed too late). In contrast, we did a better job observing burning in 2022 by starting earlier. Table 2 shows the percentage of plots by treatment status for the 2022 spot checks that were labeled burned or not burned using this approach. The number of observations for each category is noted in parentheses.

	Control	Treatment
Burned	0.82 (134)	0.88 (51)
Did not Burn	0.18 (29)	0.12 (6)

Table 2: Share of Aggregate Burn Labels by Treatment Status for Spot Checks 2022

Notes: The table shows percentage of plots that were burned or not burned in spot checks 2022. Figures outside parenthesis are percentages and figures inside parenthesis are number of observations.

This is quite a restrictive approach to assigning burn labels, and it particularly limits the fields labeled unburned because it requires the presence of machinery on a plot or the observation of wheat growing early in the season with rice straw between rows. We used a set of criteria that requires a positive identification of unburned plots, rather than just the absence of burning, because high quality labels are required for the satellite-based classification step, and absence of a clear sign of burning like active burning, smoldering, partially burnt vegetation, or fully burnt vegetation is not a clear sign that burning never occurred on a plot during the study period.

To investigate if there were any differences across treatment status for signs of burning across visits, we also created contingency tables that show by plot-visit what proportion (split by treatment status) of farmers displayed signs of burning, and the co-occurrence with potential survey items (like the presence of a Happy Seeder). Table 3 shows for a

given plot-visit, the proportion of control or treatment plots that showed signs of burning given the presence of a straw baler. The treatment and control groups are nearly identical in the observed proportions, and we never see signs of burning when a straw baler is present. Table 4 shows the same calculations, but for the presence of a Happy Seeder. The treatment and control groups are still very similar, but signs of burning did occur albeit infrequently when a Happy Seeder was present.

		Control	Treatment
No Straw Baler	No Signs of Burning	0.873	0.869
	Signs of Burning	0.123	0.128
Straw Baler	No Signs of Burning	0.004	0.003
	Signs of Burning	0.000	0.000

Table 3: Co-Occurrence by Visit of Straw Balers and Signs of Burning

Notes: The table shows the proportion of control or treatment plots that showed signs of burning given the presence of a straw baler for a given plot-visit.

		Control	Treatment
No Happy Seeder	No Signs of Burning	0.868	0.860
	Signs of Burning	0.121	0.125
Happy Seeder	No Signs of Burning	0.009	0.012
	Signs of Burning	0.002	0.003

Table 4: Co-Occurrence by Visit of Happy Seeders and Signs of Burning

Notes: The table shows the proportion of control or treatment plots that showed signs of burning given the presence of a happy seeder for a given plot-visit.

We also created contingency tables showing the proportion of plots (split by treatment status) that *ever* show clear signs of burning across visits given the survey question response on a plot-visit. For example, Table 5 shows that observing finely chopped crop residue is not a robust sign that burning never occurred. In the treatment group, 9% of plots were observed with finely chopped crop residue and obvious burning in a plot-visit, while 14% of plots with finely chopped crop residue never showed obvious burning. Tables 6, 7, 8 show this same pattern of weak separability between survey responses meant to be indicative of not burning and farmers actually never burning that plot. This weak separability in our inputs to our labels results in noise in the dependent variable (burning) in our satellite-based classification model. We address this noise by only including a cross-visit burn label when we see a clear sign that a plot was burned or never burned. However, this severely limits our training data for unburned fields. Furthermore, as seen in Table 2, very few plots showed no obvious signs of burning, and treatment does not appear to have a positive effect on farmers not burning their plots, even under this restrictive categorization scheme.

		Control	Treatment
No Finely Chopped Crop Res. Obs.	No Obvious Burning Obs.	0.550	0.537
	Obvious Burning Obs.	0.230	0.232
Layer of Finely Chopped Crop Res.	No Obvious Burning Obs.	0.167	0.143
	Obvious Burning Obs.	0.053	0.088

Table 5: Finely Chopped Crop Residue and Obvious Burning Ever Occurring

Notes: The table shows the proportion of plots (split by treatment status) that *ever* show clear signs of burning across visits given the survey question response on a plot-visit.

		Control	Treatment
No Happy Seeder	No Obvious Burning Obs.	0.709	0.674
	Obvious Burning Obs.	0.280	0.311
Happy Seeder	No Obvious Burning Obs.	0.008	0.006
	Obvious Burning Obs.	0.003	0.009

Table 6: Happy Seeder and Obvious Burning Ever Occurring

Notes: The table shows the proportion of plots (split by treatment status) that *ever* show clear signs of burning across visits given the survey question response on a plot-visit.

		Control	Treatment
No Straw Baler	No Obvious Burning Obs.	0.713	0.677
	Obvious Burning Obs.	0.283	0.320
Straw Baler	No Obvious Burning Obs.	0.004	0.003
	Obvious Burning Obs.	0.000	0.000

Table 7: Straw Baler and Obvious Burning Ever Occurring

Notes: The table shows the proportion of plots (split by treatment status) that *ever* show clear signs of burning across visits given the survey question response on a plot-visit.

		Control	Treatment
No Straw Obs.	No Obvious Burning Obs.	0.694	0.662
	Obvious Burning Obs.	0.276	0.305
Wheat in Rows with Straw Between	No Obvious Burning Obs.	0.023	0.018
	Obvious Burning Obs.	0.007	0.015

Table 8: Wheat in Rows with Rice Straw between and Obvious Burning Ever Occurring

Notes: The table shows the proportion of plots (split by treatment status) that *ever* show clear signs of burning across visits given the survey question response on a plot-visit.

Lastly, we present sample characteristics and balance for spot checks 2022 data at the farmer level in table 9. For balance we regress baseline cultivation area on treatment status and markaz (as lottery was stratified at markaz level). The area variable comes from the application data of AgriPunjab.

	Non-Spot Check		Spot Check		Diff.
	Mean	SD	Mean	SD	
Cultivation Area	13.49	21.01	15.08	29.52	3.891* (2.160)
Observations	397		334		

Table 9: Spot Checks 2022 Sample Characteristics and Balance

Notes: The Diff. column is the coefficient of a simple farmer level regression of baseline cultivation area on treatment status and markaz (stratification level) with standard errors clustered at the markaz level and shown in parenthesis. Stars indicate whether this difference is significant. *: ($p < 0.10$), **: ($p < 0.05$), ***: ($p < 0.01$)

Field-based Sensors A significant challenge with predicting crop residue burning at the plot level is obtaining high-resolution labels (Did a given plot burn or not burn during the study period?), as illustrated by the spot check results presented in the previous section. Scaling up a field experiment on crop burning makes it cost-prohibitive for surveyors to visit each plot in the study every day until the crop residue is managed, and less frequent visits substantially reduce the quality of the labels. In our 2022 spot checks, surveyors visited plots on a weekly basis over the study period. Moreover, this data collection and previous work on crop burning in India [14] suggest that obvious signs of plot burning were often only visible for 24-48 hours in the satellite imagery (and on the ground). We attempted to overcome these hurdles by testing our ability to capture high-fidelity burn signals using ground-based sensors. Ground-based sensors can monitor burning indicators continuously, preventing the gaps in temporal coverage presented by using in-person field visits. This is likely to result in higher quality labels if sensors are able to reliably detect burn events, which would in turn improve the accuracy of a statistical model like a random forest used to predict burning on other fields.

To test the potential for sensor-based burning detection, we designed an experiment on a large private farm in a rice growing region outside Lahore without a history of burning (a farm with a history of burning would have been preferable because it more closely resembles a typical rice field in Pakistan). The typical plot size on the farm was about 1 acre. We used four different sensors to measure air temperature, air humidity, soil temperature, and soil water content while burning dry crop residue on each plot. The sensors were locally assembled at a cost of ~PKR19,000 (~US\$68) per sensor. We also used a scrap of fabric attached to a pole as a low-tech sensor (i.e., if the fabric burned, the field burned), and collected pre/post burn soil samples on each of the plots (described in more detail below).

We considered a number of different field based “sensors” including buried scraps of fabric, thermal infrared cameras, and visible light cameras. These were discarded ex ante.³ We sought to identify sensors cheap enough to allow scaling up of label collection efforts, but robust enough to have a consistent signal of residue burning.

A concern with all sensors is the potential for tampering by farmers. For the four main sensors we used, a farmer could conceivably burn around the sensors if soil or air conductivity of temperature or moisture was low. For thermal or visible light cameras, a farmer could not burn the section of field in the field of view of that sensor. A scrap of fabric could easily be tilled or removed if the farmer was aware of its location. Etceteras. Therefore, the robustness of the burn signal from the sensor was key. If we could not observe a very robust signal in an experimental setting under different distances of the fire to the sensor, field deployment would be even less reliable.⁴

³Burying scraps of fabric as a burn indicator is challenging in that they must be very precisely marked by surveyors, and might be tampered by farmers. We avoided cameras for concerns about cost and observer (Hawthorne) effects.

⁴Controlled burns have the potential to generate labels, however, they may create a very different spectral signature than “natural” burns, hence the interest in sensor that could be deployed in natural settings.




Figure 2: Count of Spot Checks and Managed/Burned Plots by Year

Objective

We partnered with the Centre for Water Informatics & Technology (WIT) and its sister unit, the Agricultural Robotics Lab, at the Lahore University of Management Sciences (LUMS) that had already successfully designed and tested sensors that could measure soil and air temperature and moisture levels and relay real-time information to an online dashboard. We anticipated that in the event of fire, soil and air temperature and moisture levels should increase and decrease, respectively. We procured four types of sensors – soil temperature sensor, soil moisture sensor, air temperature sensor, air moisture sensor – and Internet of Things (IoT) Gateway for Global System for Mobiles (GSM) connectivity (to transmit data to the server), as well as Gateway Power Bank to overcome power blackouts (typical for Pakistan). We also partnered with a local farmer – facilitated by the World Wildlife Fund (WWF)-Pakistan – who rented to us a small piece of his farm where we could undertake the activity: burn a plot in a controlled way and study sensor data for temperature and moisture to see if fire incidence could be detected.

The private farm was situated in a rice growing region around 45 minutes from Lahore. The farm had some workforce available to aid the research team on the ground (mainly the field RA) with basic activities, such as preparing the sub-plots, managing fire, etc.

Approach and Data Collection

Our sensors were comprised of five “poles” (each ~ 50 cm tall; see Figure 3), one Gateway, and a virtual Dashboard. Each pole had one soil temperature sensor + one soil moisture sensor and one air temperature sensor + one air moisture sensor attached to it. Soil sensors were attached at the bottom of the pole and air sensors were attached at the top of the pole. The sensors were detachable and not all four were required to be attached to the same pole; altogether we used four poles for the exercise (and kept the 5th pole as back-up in case any of the other four got damaged from fire). Additionally, each pole had a transmitter attached to it that allowed it to transmit real-time data in flexible time-pockets (e.g., 20-min, 5-min, 1-min, etc.) to the Gateway. The poles had to be placed inside the plot to be able to pick up and share readings with the Dashboard via the Gateway. The Gateway could receive data from any number of poles within a 1 km radius and was placed in a secure location outside the plot to avoid tampering. The Dashboard provided real-time access to the research team of (i) raw data and (ii) graphs based on the data received by the Gateway.



Figure 3: Sensors (Air) and Transmitter Attached to the Pole

The field protocol for the sensors is shown in the appendix.

Results

To assess the outputs of the four main sensors, we first created time-series plots of their signals (see Figures 4 and 5), normalized by the time the field was burned and analyzed them for a trend break when burning began. In Figure 4, we see a defined trend break when the fire is lit (red vertical line). Air temperature increases, percent humidity in the air decreases, soil temperature increases, and soil water content increases. These results are what we would expect, and they are consistent across sensors. Unfortunately, this was not the typical output we observed from sensors. A more typical example of the ground sensor output is Figure 5, where there is no obvious trend break when the fire is lit. Only two out of 17 trials resulted in figures with a clear trend break. We suspect that this is a result of different distances from burning crop residue significantly affecting sensor readings; moreover, soil and air conductivity is low enough that the low-intensity fires used to manage crop residue are not hot enough to be picked up by sensors unless the sensors are extremely close to the fire. For instance, in fire round nine the sensors were inserted at a depth of two inches into the soil and in fire round six the sensors were inserted at a depth of four inches into the soil; the former gave a stronger signal.




Figure 4: Ideal Example of Ground Sensor Time Series

To see if statistical analysis could detect patterns more readily than visual inspection, we used a random forest model to predict pre or post burn period for a given plot. The model was too weakly predictive of this simple task to be useful for labeling purposes, so we trained a second, more sophisticated featurization (created different RHS inputs) of the random forest that used a moving average to impute observations across continuous time, and then predicted if the plot was in a pre or post burn period. This model was also a very weak predictor of burning. Next, we posed an easier question: how did elapsed time since the fire was lit affect our sensor outputs? We ran a regression using the same moving average feature set from the second random forest to find the effect of the elapsed time since the fire was started on each of the rolling means (air temp, air humidity, soil temp, soil water content) from the sensor data. The standard errors were clustered by fire round and elapsed time in minutes since the fire, with 0 as the reference period. Ideally, we




Figure 5: Typical Example of Ground Sensor Time Series

would have compared burned fields with controls at the same points in time; since we lacked contemporaneous control fields, we used fixed effects with the pre-burn period as a baseline.

Figure 6 displays confidence intervals for the effect of time on rolling means of sensor output at every minute pre- and post-burning. The estimates are very noisy and capture zero at nearly every point in time, showing that there is too much noise in ground sensor detection to offer a reliable burning measure. We conclude that ground sensors are unlikely to resolve measurement challenges, even under ideal circumstances.




Figure 6: Fixed Effects Confidence Intervals of Burning Effects on Sensor Measures

Soil Sampling

Objective

We also collected soil samples during the plot burning exercise detailed above. The field protocol for this component is also shown in the appendix. The objective for collecting pre- and post-burn soil samples at the sub-plots was to understand how fire affects soil chemicals. Soil measurement may offer an attractive alternative to spot checks or sensors in that it provides – in theory – a more durable signal. If burning leaves persistent changes in chemical content, then soil samples would only need to be collected at the end of the season, reducing the temporal burden of accurate measurement.

The set of chemicals chosen for the study, namely Organic Matter, Phosphorus, Potassium, Zinc, Copper, Iron, Magnesium, and Boron, were shortlisted with input from soil scientists at AgriPunjab; more specifically, those chemicals were chosen that had a larger probability of being affected by fire and were of high research interest for the department.

Approach and Data Collection

We divided the plot at the farm into 10 sub-plots, each with dimensions of 30 by 30 feet² (~ 9 by ~ 9 meters²). Subsequently, sub-plot # 1 was burnt on day 1, sub-plot # 2 was burnt on day 2, and so on until sub-plot # 10 was burnt on day 10. One pre-burn soil sample was collected from each sub-plot; and four post-burn soil samples were collected from each sub-plot: 1st on the burning day, 2nd three days after burning day, 3rd six days after burning day, and 4th 18 days after burning day. Since, each soil sample consisted of two sub-samples, 1st at a depth of six inches and 2nd from

the same point but at a depth of 12 inches, we ended up with two pre-burn and eight post-burn sub-samples per sub-plot, resulting in 100 observations altogether.

Results

Event-study style plots for each of the eight chemicals using raw data are shown in the appendix (see Figure A.8).

Largely due to budgetary constraints we were unable to collect data on control sub-plots. This leaves us with just a treatment group (10 sub-plots) and two time periods (before and after burning). Accordingly, we are only able to implement before-after comparisons. The specification does not distinguish between sample depths (6 and 12 inches) and pools the full set of 100 observations.

We use the following specification to estimate the effects of fire on soil chemicals:

$$Chemicals_{it} = a + b * Pre_Post_t + u_{it} \quad (1)$$

where i and t denote sub-plot and day, respectively, $Chemicals_{it}$ is a continuous level measure for each of the eight chemicals (Organic Matter, Phosphorus, Potassium, Zinc, Copper, Iron, Magnesium, and Boron), and Pre_Post is a dummy variable that is equal to 1 if the sub-plot is in post-burning period and 0 otherwise. We cluster standard errors at the plot level.

Table 10 shows results from this regression. As can be seen, other than for Organic Matter (column 1), none of the chemicals are significantly effected by fire. For Organic Matter, fire is associated with a 0.039 pp decrease in the amount of Organic Matter present in the soil.

	(1) OM	(2) Phosphorus	(3) Potassium	(4) Zinc	(5) Copper	(6) Iron	(7) Magnesium	(8) Boron
Before-After	-0.039* (0.018)	-0.363 (0.447)	3.200 (3.031)	-0.154 (0.099)	0.006 (0.025)	-0.035 (0.076)	0.018 (0.017)	-0.038 (0.024)
Constant	0.376*** (0.019)	8.450*** (0.511)	75.200*** (3.317)	1.812*** (0.168)	0.462*** (0.053)	2.432*** (0.224)	0.429*** (0.014)	0.559*** (0.025)
Observations	100	100	100	100	100	100	100	100

Table 10: Effects of Fire on Chemicals

Notes: Standard errors in parentheses. * : ($p < 0.10$), ** : ($p < 0.05$), *** : ($p < 0.01$)

2.3 Aggregate Measurement of Burning

Farmer applications to the MMRCR program were accompanied with an affidavit where the farmer agreed to comply with: “*I will not refuse to lease my machinery to the farmers of my Markaz who want to make the paddy residue efficient.*”. Accordingly, farmers who received the subsidy may loan or rent machines to other farmers in their villages. Since both Rice Straw Shredder and Happy Seeder are machines without an engine and wheels of their own (they are mounted on the back of a tractor for use), it is reasonable to think that the machines would be rented in or around treatment villages. Therefore, to investigate more aggregate measures of burning, corresponding to a rental catchment, we created a Union Council (UC) satellite based burning measure. A UC is geographic unit above a village but below a markaz (average population of a Punjabi UC is $\sim 28,000$ [15]; [16]).

The administrative lottery data that includes treatment (Draw 1) and control (remaining draws) farmers (across Waves 1 and 2) does not include farmer GPS coordinates. Instead, we used Punjab UC Shapefiles from the Institute of Development and Economic Alternatives (IDEAS) and assigned farmers in the lottery data to UCs by extracting UC information across the following variables present in the lottery data: (a) an incomplete farmer UC variable; (b) farmer

address in English; (c) farmer address in Urdu and; (d) farmer village name in Urdu. This labor intensive matching was performed by an English speaking and both English and Urdu speaking member of the research team to maximise match rate as well as minimise any data entry errors in the manual process. We end up with match rates of 48.96% (469/958) and 39.74% (484/1218) for Waves 1 and 2, respectively. Table 11 further shows that treatment status is not predictive of match rate for both the waves.

	(1) Wave 1: Match Found	(2) Wave 2: Match Found
Treatment	-0.021 (0.046)	-0.055 (0.033)
Constant	0.493*** (0.017)	0.411*** (0.016)
Observations	356	522

Table 11: Lottery-UC Match Rate Balance by Wave

Notes: The table shows results from the regression: $Treatment_i = a + b * Match_i + u_i$, where i depicts farmer, $Treatment$ is a dummy variable equal to 1 if farmer's lottery draw was 1 and 0 otherwise, and $Match$ is a dummy variable equal to 1 if farmer in lottery data is assigned to UC and 0 otherwise. Standard errors are shown in parentheses. *: ($p < 0.10$), **: ($p < 0.05$), ***: ($p < 0.01$)

At the UC level, a treatment variable was constructed based on the number of lottery winners in a UC. We also created a variable for the total number of lottery winners in the UC, to control for local equipment demand. We then make a few adjustments to this data. We merge the UC level administrative dataset with a measure of fire activity using VIIRS fire counts during the burning season (October to December) for 2021 and 2022.

VIIRS sensors have been interpreted as an active fire product, measured approximately daily. For fire counts we used the daily 1km scale Suomi National Polar-Orbiting Partnership NASA VIIRS Thermal Anomalies/Fire (VNP14A1) Version 1 data product. The measurement, however, is still imperfect: smoke and haze can interfere with fire detection and residue fires typically last only a few hours. In addition, crop residue fires are relatively low intensity, so even if a satellite collects clear images at the time that a field is burning, if it is the only fire in a pixel, it may not register as a fire. Therefore, treatment effects on burning using VIIRS should be interpreted with some caution. In the UC-VIIRS data, we observe 2.04 and 0.90 average fires per UC in 2021 and 2022, respectively.

3 Subsidy Impacts

3.1 Spot Check 2022 Based Treatment Effects

Here we use the 2022 spot checks data to estimate treatment effects for whether or not farmers ever burned. As described under Spot Checks in subsection 2.2.1, we aggregated the weekly, plot-level observations from the spot checks to a plot-level cross-section and to aggregate across weeks to the farmer level, we assigned a "Yes" burn label to plots if any of the weekly spot checks were labeled as a burn, and a "No" burn label if plots were never labeled as a burn. We use the following specification to estimate treatment effects.

$$Burned_i = a + b * Treatment_i + c * \mathbf{X}_i + u_i \quad (2)$$

where i denotes farmer, $Burned$ is a dummy variable equal to 1 if farmer burned and 0 otherwise, and $\mathbf{X}$ is a vector of controls including baseline cultivation area and markaz fixed effects (since lottery was stratified at markaz level). We cluster standard errors at the markaz level.

Results from this regression are shown in Table 12. As can be seen from the table, we do not see any statistically significant effect of the treatment on whether or not the farmer ever burned (95% CI = -0.06, 0.27).

	(1)
	Farmer Ever Burned
Treatment	0.105 (0.084)
Area	0.001* (0.000)
Constant	0.823*** (0.103)
Observations	386

Table 12: Effects of Lottery Winning on Farmer Level Fire Incidence in Spot Checks 2022

Notes: Standard errors in parentheses. *: ($p < 0.10$), **: ($p < 0.05$), ***: ($p < 0.01$)

3.2 Aggregate Treatment Effects Using Satellite Data

Next, we use the UC-level VIIRS-based dataset described in subsection 2.3. This dataset reports daily fire count for the UCs in which treatment and control farmers for lottery Wave 1 and 2 in years 2021 and 2022. We estimate the effects of the number of lottery winners on the number of fires at the UC level, conditioning on non-zero lottery applicants. For robustness we also use share of lottery winners instead of number of lottery winners. We run the following regression specification for years 2021 and 2022 for Wave 1 and only for year 2022 for Wave 2 (controlling for 2021 fires).

$$Fires_i = a + b * Winners_i + c * X_i + u_i \quad (3)$$

where i denotes UC, $Fires$ is a continuous variable measuring total number of fires in a UC in a year, $Winners$ is a continuous variable separately measuring total number of lottery winners in a UC and share of lottery winners in a UC, and X is a vector of controls including a continuous measure of number of applicants in a UC, fires in 2021 (only for Wave 2 regressions), and markaz fixed effects (since lottery was stratified at markaz level). We include only UCs with a positive number of lottery applicants and cluster standard errors at the markaz level.

Table 13 presents results for the above regressions. Column 1 shows that a one-farmer increase in the number of lottery winners in UCs with non-zero applicants *increases* fires by 1.026. This effect is statistically significant at the 10% level. The direction of the effect appears counter-intuitive in the sense that one would expect the number of fires to go down as the number of winners goes up in a UC either through increased personal usage of machinery by the winners or through renting of equipment to other farmers in the UC. The negative coefficient on number of applicants presumably reflects demand; more applicants is correlated with a greater interest in alternatives to burning. Using share of winners instead of number of winners in column 2 shows similar patterns. In column 3, we analyze the effect on fires in the following year. The slightly smaller, though still positive and statistically significant estimate implies that the positive relationship persists even after farmers have time to learn about the machines and the rental market.

For Wave 2, columns 5-6, coefficients are negative but not significantly different from zero.

3.3 Conclusion and Next Steps

In this work, we have explored alternative methods for ground truthing crop residue burning, including spot checks, sensors, and soil testing. All methods present challenges. While sensors seemed most promising *ex ante*, our field work highlighted that readings are not sufficiently sensitive or reliable. Spot checks on the other hand provide lower

	Wave 1				Wave 2	
	(1) Fires in 2021	(2) Fires in 2021	(3) Fires in 2022	(4) Fires in 2022	(5) Fires in 2022	(6) Fires in 2022
Number of Winners	1.026* (0.583)		0.821* (0.477)		-0.228 (0.252)	
Share of Winners		1.885** (0.889)		1.892* (0.971)		-0.234 (0.472)
Number of Applicants	-0.398*** (0.125)	-0.355*** (0.117)	-0.275*** (0.071)	-0.240*** (0.065)	0.156 (0.160)	0.130 (0.159)
Fires in 2021					0.705*** (0.079)	0.704*** (0.080)
Constant	0.796*** (0.250)	0.711*** (0.235)	0.551*** (0.141)	0.480*** (0.129)	-0.312 (0.320)	-0.260 (0.318)
Observations	668	668	668	668	745	745

Table 13: Effects of Lottery Winning on Union Council Level Fire Incidence

Notes: Standard errors in parentheses. * : ($p < 0.10$), ** : ($p < 0.05$), *** : ($p < 0.01$)

quality labels and a greater risk of Hawthorne effects. Therefore, realistically speaking, spot checks can be used for less-continuous monitoring (relative to sensors), but still end up producing noisy fire outcomes as we learned from our fairly intensive spot check (2022) activity. Soil testing deserves further investigation. The potential to measure once, at the end of the season, is attractive, and a set of pure control fields might reveal additional signal beyond what we were able to observe. Repeated burning over long periods of time would also likely have a different effect on soil fertility.

We use our measurement and the administrative data to test for effects of machinery subsidies on burning activities. If anything, we find some indication of a positive relationship – while most coefficients are noisily estimated, all coefficients with positive sign are significantly different from zero. More importantly, we see no indication of a large effect of machinery subsidies on burning, suggesting that complementary policies to encourage use will be necessary for this initial investment to realize its potential.

References

- [1] Food and Agriculture Organization. Remote sensing for spatio-temporal mapping of smog in punjab and identification of the underlying causes using gis techniques (r- smog). Technical report, 2018.
- [2] Tanvir Ahmed, Bashir Ahmad, and Waseem Ahmad. Why do farmers burn rice residue? examining farmers' choices in punjab, pakistan. *Land use policy*, 47:448–458, 2015.
- [3] Gaurav Kumar Porichha, Yulin Hu, Kasanneni Tirumala Venkateswara Rao, and Chunbao Charles Xu. Crop residue management in india: Stubble burning vs. other utilizations including bioenergy. *Energies*, 14(14):4281, 2021.
- [4] Niveta Jain, Arti Bhatia, Himanshu Pathak, et al. Emission of air pollutants from crop residue burning in india. *Aerosol and Air Quality Research*, 14(1):422–430, 2014.
- [5] Priya Shyamsundar, NP Springer, Heather Tallis, Stephen Polasky, Mangi Lal Jat, Harminder S Sidhu, PP Krishnapriya, Natalya Skiba, William Ginn, Vikram Ahuja, et al. Fields on fire: Alternatives to crop residue burning in india. *Science*, 365(6453):536–538, 2019.

- [6] Sushil Kharia, H Thind, Sandeep Sharma, H Sidhu, M Jat, and Yadvinder Singh. Tillage and rice straw management affect soil enzyme activities and chemical properties after three years of conservation agriculture based rice-wheat system in north-western india. *International Journal of Plant & Soil Science*, 15(6):1–13, 2017.
- [7] HS Sidhu, Manpreet Singh, Yadvinder Singh, John Blackwell, Shiv Kumar Lohan, Elizabeth Humphreys, ML Jat, Vicky Singh, and Sarbjeet Singh. Development and evaluation of the turbo happy seeder for sowing wheat into heavy rice residues in nw india. *Field Crops Research*, 184:201–212, 2015.
- [8] A Patrick Behrer. Earth, wind and fire: The impact of anti-poverty efforts on indian agriculture and air pollution. *Unpublished manuscript, Stanford University, Stanford, CA*, 2019.
- [9] Teevrat Garg, Maulik Jagnani, and Hemant K Pullabhotla. Agricultural labor exits increase crop fires. 2021.
- [10] Jeetendra Prakash Aryal, Tek B Sapkota, Ritika Khurana, Arun Khatri-Chhetri, Dil Bahadur Rahut, and Mangi Lal Jat. Climate change and agriculture in south asia: Adaptation options in smallholder production systems. *Environment, Development and Sustainability*, 22(6):5045–5075, 2020.
- [11] Longlong Xia, Shuwei Wang, and Xiaoyuan Yan. Effects of long-term straw incorporation on the net global warming potential and the net economic benefit in a rice–wheat cropping system in china. *Agriculture, ecosystems & environment*, 197:118–127, 2014.
- [12] Alwin Keil, PP Krishnapriya, Archisman Mitra, Mangi L Jat, Harminder S Sidhu, Vijesh V Krishna, and Priya Shyamsundar. Changing agricultural stubble burning practices in the indo-gangetic plains: is the happy seeder a profitable alternative? *International Journal of Agricultural Sustainability*, 19(2):128–151, 2021.
- [13] B Kelsey Jack, Seema Jayachandran, Namrata Kala, and Rohini Pande. Money (not) to burn: Payments for ecosystem services to reduce crop residue burning. Technical report, National Bureau of Economic Research, 2022.
- [14] Kendra Walker, Ben Moscona, Kelsey Jack, Seema Jayachandran, Namrata Kala, Rohini Pande, Jiani Xue, and Marshall Burke. Detecting crop burning in india using satellite data. *arXiv preprint arXiv:2209.10148*, 2022.
- [15] Punjab Bureau of Statistics. Punjab development statistics 2022.
- [16] Pakistan Bureau of Statistics. Population by single year age, sex and rural/urban, pakistan population census 2017.

A Appendix

A.1 Sensor (and Soil Sampling) Field Protocol

The following steps were taken (in the same order):

Note: There were minor instances where the field RA deviated from the above protocol given ground realities / requirements

1. We prepared 10 small (manageable) plots, in square shape. Preparation entailed defining plot boundary, ploughing plot soil, and sprinkling thin layer of dry crop residue on the plot. Dimensions of plot were logged. Field RA took input on height of residue from workforce available at the farm to mimic real-fire setting
 - (a) Baseline data: The poles were allowed to capture baseline data (i.e., without any fire) for one hour (frequency of every min) – time at which the one hour started and ended was logged
2. The four poles were inserted into the soil - see pattern statement in 6. below - the 5th pole was kept as back-up (e.g., in case any pole got destroyed from fire, etc.). Pole depth, i.e., how deep the pole is inserted into the soil, was logged in a survey. Pole location was also logged
3. Crop residue was carefully lit using kerosene and matchbox. Time when fire was lit and when it ended was logged
4. Soil samples (essentially to establish fire incidence through changes in soil chemicals - details in next sub-section)
 - (a) Day 1.1: Before each burning round: Collected 2 x 0.5 kg soil from same point on the plot; marked following points on the plot: one center, one 1 ft from the edge, one midpoint in between to check for gradient; then chose the midpoint for the 2 x 0.5 kg sample; 1st sample from the point was taken at 6 inches depth and 2nd from another 6 inches at the same point before burning, and placed it in sealable bag
 - (b) Day 1.2: Repeated after burning [same point ~ 1 ft away (because soil had been perturbed on same point)]. Placed soil into a plastic bag, tied / sealed, and marked date and time and location: YYMMDD-hh:mm-M (M = midpoint), e.g. 220823-13:13-M



Figure A.1: LHS: Soil Sampling Points on the Plot; RHS: Plot Dimensions

- (c) To look for potential time-variant soil signals, collected soil samples at the same plot and locations at **Days 3, 6, 18⁵**
5. Steps 1-4 (above) made 1 complete fire round
 6. Steps 1-4 (above) were repeated with varying plot surface area (30 x 30, 60 x 60 ft² = 9 x 9, 18 x 18 m²), pole location along the plot boundary, pole depth (10, 30 cm), quantity of crop residue used (1, 2x), quantity of kerosene used (1, 2x normal x L per kg residue), and points of ignition (corner or center) – on pole location, for simplicity, the poles were kept equidistant from one another along the plot boundary, e.g., if a pole was inserted at the midpoint of side 1 of the square plot, then the remaining 3 poles were also inserted at the midpoints of their respective sides. Overall, because of the limited days, we could only capture 1 fire round data for each permutation (no replica; or consider 2 replicas of each set)

Note: Soil sampling was only done each day only on 30 x 30 ft² = 9 x 9 m² plot

⁵Assume signal might decay ~ exponentially over time: more samples at earlier times, then slowly taper off

- (a) Day 1: **10 cm** depth for all, plot area **30 x 30 ft²**, crop residue amount **1x** normal, kerosine **1x** normal, **ignition point (center)**, pole **pattern 1**: center c, midpoint, corner, 1 ft away from corner outside plot
- i. If we found outside corner not to give great signal, shifted to **pattern 2** for the rest of the trials: center, spaced diagonally 5 ft 8 in (= 4 ft vertical / horizontal) each towards corner




Figure A.2: LHS: Day 1 Pattern 1; RHS: Day 1 Pattern 2

Repeated for plot area **60 x 60 ft²**

- (b) Day 2: **10 cm** depth for all, plot area **30 x 30 ft²**, crop residue amount **1x** normal, kerosine **1x** normal, **ignition point (corner far away from poles)**, pole **pattern 1**




Figure A.3: Day 2 Pattern 1

Repeated for plot area **60 x 60 ft²**

- (c) Day 3: **10 cm** depth for all, plot area **30 x 30 ft²**, crop residue amount **1x** normal, kerosine **1x** normal, **ignition point (corner at poles)**, pole **pattern 1**




Figure A.4: Day 3 Pattern 1

Repeated for plot area **60 x 60 ft²**

- (d) Day 4: **10 cm** depth for all, plot area **30 x 30 ft²**, crop residue amount **1x** normal, kerosine **1x** normal, **ignition point (corner away from poles)**, pole **pattern 1**




Figure A.5: Day 4 Pattern 1

Repeated for plot area **60 x 60 ft²**

- (e) Day 5: **30 cm** depth for all, plot area **30 x 30 ft²**, crop residue amount **1x** normal, kerosine **1x** normal, ignition point (**center**), pole **pattern 1**

Repeated for plot area **60 x 60 ft²**

- (f) Day 6: **10 cm** depth for all, plot area **30 x 30 ft²**, **crop residue amount 2x** normal, kerosine **1x** normal, ignition point (**center**), pole **pattern 1**

Repeated for plot area **60 x 60 ft²**

- (g) Day 7: **10 cm** depth for all, plot area **30 x 30 ft²**, **crop residue amount 1x** normal, **kerosine 2x** normal, ignition point (**center**), pole **pattern 1**

Repeated for plot area **60 x 60 ft²**

- (h) Day 8: **10 cm** depth for all, plot area **30 x 30 ft²**, crop residue amount **1x** normal, kerosine **2x** normal, ignition point (**center**), **pole pattern 2**




Figure A.6: Day 8 Pattern 2

Repeated for plot area **60 x 60 ft²**

- (i) Days 9-10: Any replicates for weird data OR **pattern 1** ignite **two corners**

Repeated for plot area **60 x 60 ft²**




Figure A.7: Days 9-10 Pattern 1

A.2 Soil Sampling Event-Study Raw Plots




Figure A.8: Raw Event-Study Plots for Soil Sampling

Notes: Each dot represents average of 6 and 12 inches samples for a given sub-plot on a given event-day; the vertical dotted line represents the event (fire)